



HYCU R-Cloud Hybrid Cloud Edition v5.0.0

User Guide

Legal notices

Copyright notice

© 2024 HYCU. All rights reserved.

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

Trademarks

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

Acropolis and Nutanix are trademarks of Nutanix, Inc. in the United States and/or other jurisdictions.

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

Azure®, Microsoft®, Microsoft Edge™, Microsoft Entra™, Microsoft 365™, and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Cloudian and HyperStore are registered trademarks or trademarks of Cloudian, Inc.

Dell Technologies, Dell and other trademarks are trademarks of Dell Inc. or its subsidiaries.

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

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

NetApp®, NetApp Keystone®, NetApp Cloud Volume Services®, and ONTAP®, and are trademarks of NetApp, Inc. and are registered in the United States and/or other jurisdictions.

Red Hat Enterprise Linux is a registered trademark of Red Hat, Inc. or its subsidiaries in the United States and other countries.

VMware ESXi™, VMware Tools™, VMware vCenter Server®, VMware vSAN™, VMware vSphere®, VMware vSphere® Data Protection™, VMware vSphere® Virtual Volumes™, and VMware vSphere® Web Client are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions.

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	15
HYCU key features and benefits	16
Data protection environment overview	17
HYCU data protection	18
2 Deploying the HYCU virtual appliance	20
Sizing resources for your HYCU backup infrastructure	22
Adjusting firewall configuration	23
Targets	24
Virtual machine and volume group protection	25
Restore of individual files and application awareness	26
File share protection	27
Infrastructural services	32
User interaction and administration	33
SpinUp	34
Adjusting antivirus configuration	35
Deploying HYCU to a Nutanix AHV cluster	35
Deployment tasks	36
Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment ..	40
Deploying HYCU to an AWS GovCloud (US) environment	42
Deploying HYCU to an Azure Government environment	44
Deploying HYCU to Google Cloud	46
Deployment tasks	47
Signing in to HYCU	49
Setting the language	50

3 Establishing a data protection environment	52
Adding sources	53
Adding a Nutanix cluster	54
Adding a vCenter Server	57
Adding a cloud environment	58
Adding a file server	62
Adding a server	78
Setting up targets	78
Setting up an NFS target	80
Setting up an SMB target	82
Setting up a Nutanix target	85
Setting up a Nutanix Objects target	88
Setting up an iSCSI target	90
Setting up an Azure target	92
Setting up an Amazon S3 / S3 Compatible target	95
Setting up a Google Cloud target	99
Setting up a tape target	101
Setting up a Data Domain target	105
Defining your backup strategy	107
Taking advantage of predefined policies	108
Creating a custom policy	109
Setting a default policy	126
4 Protecting virtual machines	128
Planning virtual machine protection	129
Preparing your data protection environment	129
Preparing for disaster recovery	132

Preparing for the restore to a different source	135
Server specifics	140
SpinUp specifics	142
Enabling access to data	146
Setting up virtual machine backup configuration options	149
Backing up virtual machines	153
Restoring virtual machines	155
Restore options	157
Restoring a virtual machine	159
Cloning a virtual machine	171
Validating the virtual machine backup	188
Restoring virtual disks	193
Cloning virtual disks	195
Exporting virtual disks	197
Restoring individual files	202
5 Protecting applications	209
Enabling access to application data	209
Planning application protection	212
Setting application-specific options	218
Excluding databases from the SQL Server application backup	221
Backing up applications	222
Restoring whole applications	223
Restore options	225
Restoring a virtual machine	225
Cloning a virtual machine	238
Restoring SQL Server databases	252
Restoring Exchange Server databases, mailboxes, and public folders ..	256

Restoring Oracle database instances and tablespaces	260
6 Protecting file shares	264
Configuring file share backup options	264
Backing up file shares	266
Restoring file share data	268
7 Protecting volume groups	274
Backing up volume groups	274
Restoring volume groups	275
Restoring a volume group	276
Cloning a volume group	277
Exporting virtual disks	278
8 Recovering your data protection environment	280
Preparing for disaster recovery	281
Deploying a recovery HYCU backup controller	282
Importing targets	284
Performing disaster recovery	286
Restoring the HYCU backup controller to the original source	287
Restoring the HYCU backup controller to a different source	290
9 Performing daily tasks	296
Using the HYCU dashboard	297
Managing HYCU jobs	300
Managing HYCU events	301
Configuring event notifications	302
Setting up email notifications	302
Setting up webhook notifications	303
Enabling the purge of events and jobs	305

Using HYCU reports	306
Getting started with reporting	307
Viewing reports	309
Generating reports	310
Scheduling reports	311
Exporting and importing reports	312
Viewing entity details	313
Viewing the backup status of entities	315
Tier statuses	317
Filtering and sorting data	317
Filtering data in panels	318
Sorting data in panels	324
Exporting the contents of the panel	324
Managing targets	325
Viewing target information	325
Editing a target	327
Activating or deactivating a target	329
Deleting a target	330
Managing policies	330
Viewing policy information	330
Editing a policy	331
Deleting a policy	332
Performing a manual backup	333
Setting up a validation policy	334
Archiving data manually	339
Recreating snapshots	340
Adjusting the HYCU backup controller resources	341

Customizing your HYCU web user interface	342
10 Managing users	343
HYCU groups	343
User roles	345
Setting up a user environment	347
Creating a user	348
Adding a user to a group	352
Creating a self-service group	353
Setting ownership	354
Activating or deactivating users or self-service groups	358
Switching to another group	359
Updating your user profile	360
11 Administering	362
Adding a cloud account	363
Adding a cloud provider account	363
Adding a HYCU account	371
Configuring target encryption	372
Exporting an encryption key	373
Importing an encryption key	373
Integrating HYCU with identity providers	373
Adding an identity provider to HYCU	374
Managing HYCU instances	381
Creating a HYCU instance by using the HYCU web user interface	382
Viewing HYCU instance information	383
Deleting a HYCU instance	384
Setting the iSCSI Initiator secret	384

Licensing	385
Creating a license request	387
Requesting and retrieving licenses	388
Activating licenses	389
Setting up logging	389
Configuring your network	392
Changing network settings	392
Limiting network bandwidth	394
Managing data retention	396
Filtering restore point tiers	400
Expiring data automatically	401
Expiring data from the entity panel	402
Setting power options	403
Managing secrets	404
Adding a Conjur configuration	405
Editing a Conjur configuration	407
Deleting a Conjur configuration	407
Configuring an SMTP server	407
Upgrading HYCU	408
Upgrading HYCU on a Nutanix AHV cluster	410
Upgrading HYCU on a Nutanix ESXi cluster	412
Upgrading HYCU in a vSphere environment	417
Upgrading HYCU in an AWS GovCloud (US) environment	423
Upgrading HYCU in Google Cloud	424
Upgrading HYCU in an Azure Government environment	426
Applying HYCU updates	427
Applying an update by using the HYCU web user interface	429

Applying an update by using the shell script	431
Configuring SSL certificates	432
Creating a self-signed certificate	433
Creating a certificate signing request	434
Importing a custom certificate	435
Sharing telemetry data with HYCU	438
Removing HYCU	439
12 Tuning your data protection environment	442
Accessing the HYCU backup controller virtual machine by using SSH	443
Enabling HTTPS for WinRM connections	446
Configuring FIPS mode for HYCU	446
Enabling FIPS mode for HYCU	447
Disabling FIPS mode for HYCU	447
Setting up LDAPS authentication	448
Setting up two-factor authentication	448
Managing API keys	449
Generating an API key	450
Revoking an API key	450
Managing FIDO authenticators	450
Adding a new FIDO authenticator	451
Revoking a FIDO authenticator	451
Securing SMTP connections	451
Setting up HYCU to use multiple networks	452
Setting up HYCU to use multiple networks on a Nutanix AHV or Nutanix ESXi cluster	453
Setting up HYCU to use multiple networks in a vSphere environment	454

Increasing the size of the HYCU virtual disks	455
Assigning privileges to a vSphere user	456
Using the HYCU REST API Explorer	460
Using the command-line interface	460
Using the pre and post scripts	461
13 Monitoring data protection environments	463
Managing HYCU controllers	463
Adding a HYCU controller	464
Viewing information about HYCU controllers	468
Editing a HYCU controller	470
Deleting a HYCU controller	470
Using the HYCU Manager console	470
Viewing entity data	472
Viewing events	474
Performing administration tasks	474
14 Employing Nutanix Mine with HYCU	480
Registering HYCU with Nutanix Prism	480
Accessing HYCU from the Nutanix Prism web console	481
Viewing the Nutanix Mine with HYCU dashboard	482
15 Protecting data across on-premises and cloud environments	485
Protecting data across on-premises and AWS environments	485
Migrating virtual machines across different environments	486
Performing disaster recovery of data to AWS	493
Protecting data across on-premises and Google Cloud environments ..	496
Migrating virtual machines across different environments	497

Performing disaster recovery of data to Google Cloud	503
Protecting data across on-premises and Azure environments	505
Migrating virtual machines across different environments	506
Performing disaster recovery of data to Azure	513
Protecting data across on-premises and Azure Government environments	515
Migrating virtual machines to cloud	515
Performing disaster recovery of data to Azure Government	520
A Customizing HYCU configuration settings	523
Snapshot settings	525
Utilization threshold settings	525
Display settings	526
SQL Server application settings	526
Settings for aborting jobs	526
File server settings	527
Data rehydration settings	528
Disaster recovery settings	528
User management settings	529
B After restoring a virtual machine to a different source	530
After restoring a virtual machine to a Nutanix AHV cluster	531
After restoring a virtual machine to a Nutanix ESXi cluster	532

Chapter 1

About HYCU

HYCU R-Cloud Hybrid Cloud Edition (HYCU), formerly known as HYCU for Enterprise Cloud, is a high performing backup and recovery solution for Nutanix, VMware, AWS GovCloud (US), Azure Government, file servers, and servers. It is designed to make data protection as simple and cost-effective as possible, to improve your business agility, and to bring unified security, reliability, performance, and user experience across on-premises and cloud environments.

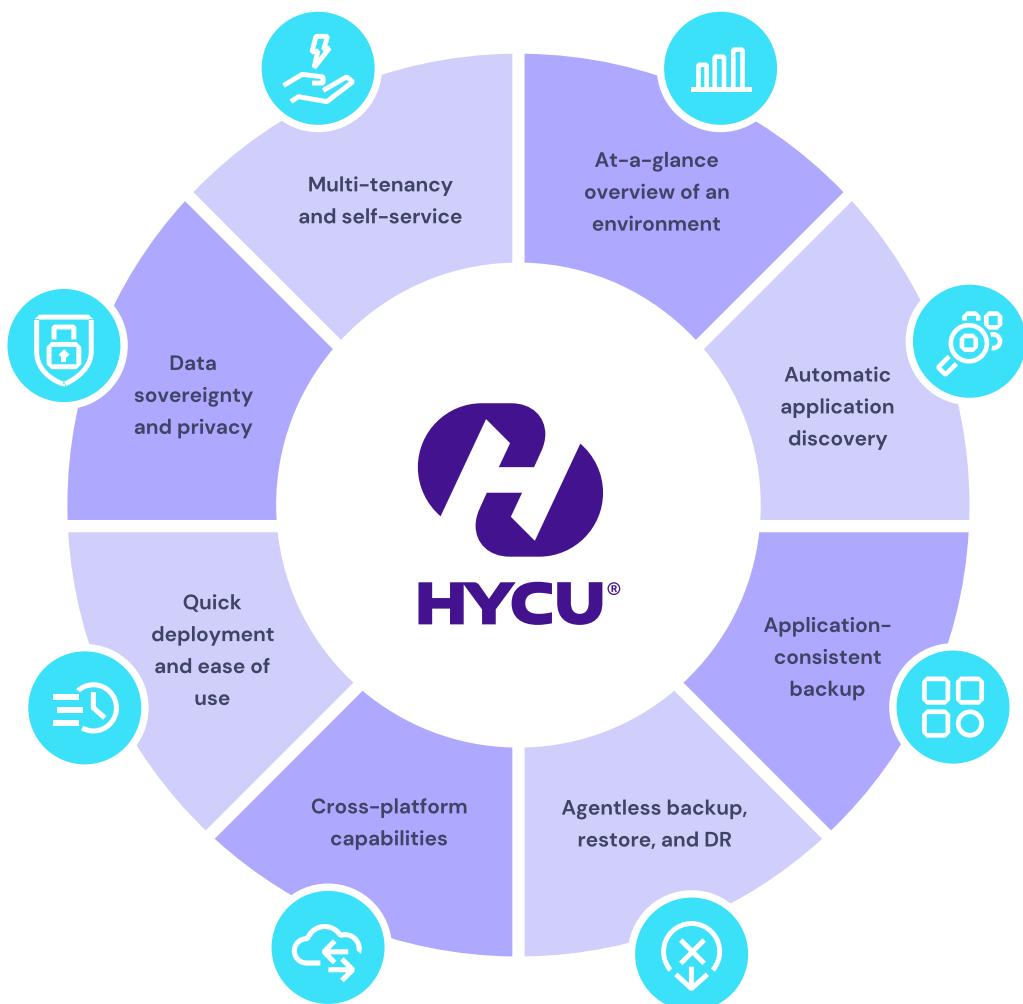


Figure 1-1: Introduction to HYCU

HYCU key features and benefits

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

- **Protects against data loss**

Delivers native and reliable data protection for mission-critical applications and data in hyperconverged environments, while ensuring data consistency and easy recoverability.

- **Simplifies deployment**

Deployment of the HYCU virtual appliance is performed through your native console.

- **Provides new-found visibility**

Discovery provides new-found visibility into virtual machines and servers, pinpointing where each application is running.

- **Protects data in a few minutes**

Data protection of virtual machines, applications, file shares, servers, volume groups, and virtual machine templates can be enabled in a few minutes after deployment.

- **Delivers predefined policies and provides opportunities for customization**

Predefined policies (Gold, Silver, and Bronze) that come with HYCU simplify the data protection implementation. However, if the needs of the data protection environment require it, a wide range of opportunities to customize policies is provided.

- **Schedules backups based on RPOs**

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

- **Discovers and protects applications**

In-built application awareness provides application discovery and application-specific backup and restore flow, ensuring that the entire application data is protected and recovered to a consistent state.

- **Lets you choose targets and sources**

Using data storage targets and sources is the administrator's choice.

- **Gives you an at-a-glance overview of your environment**

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

- **Offers a scalable backup for file servers**
Cuts down the time it takes to back up file shares, saves a significant amount of computing resources, and allows you to take more frequent backups, reducing the amount of data loss in case of a failure.
- **Provides business continuity of your data protection environment across different infrastructures**
The SpinUp functionality allows you to migrate protected data between the on-premises and cloud infrastructures (AWS, Google Cloud, global Azure, or Azure Government environments). In the event of a disaster, the SpinUp functionality provides disaster recovery of mission-critical data to cloud.
- **Provides an efficient ROBO data protection solution for Nutanix**
Backs up remote office/branch office (ROBO) data from data center replicas and enables a one-click restore within the data center or at any remote location.
- **Allows backup to become a service of the Nutanix platform**
Nutanix Mine with HYCU makes backup and recovery as a native service of the Nutanix platform and eliminates the need for isolated infrastructure for backup.

Data protection environment overview

The data protection environment consists of the following components:

HYCU backup controller	A virtual machine that processes data collected from sources and presents it in the web user interface.
HYCU interface	An interface for protecting entities and administering the data protection environment, available as the HYCU web user interface and the command-line interface (hyCLI).
Targets	Storage locations that HYCU uses for storing the protected data. Protected data can also be stored as snapshots.
Sources	Environments for which HYCU provides data protection—Nutanix clusters, vSphere environments,

	AWS GovCloud (US) environments, Azure Government environments, file servers, and servers.
Entities	Objects to which you can assign a policy and for which you therefore provide data protection—virtual machines, applications, file shares, servers, Nutanix volume groups, and vSphere virtual machine templates. Data is always protected at a granular level, allowing you to restore either the whole entities or their parts (disks and application items).

The following diagram shows the data protection environment and its most important components:

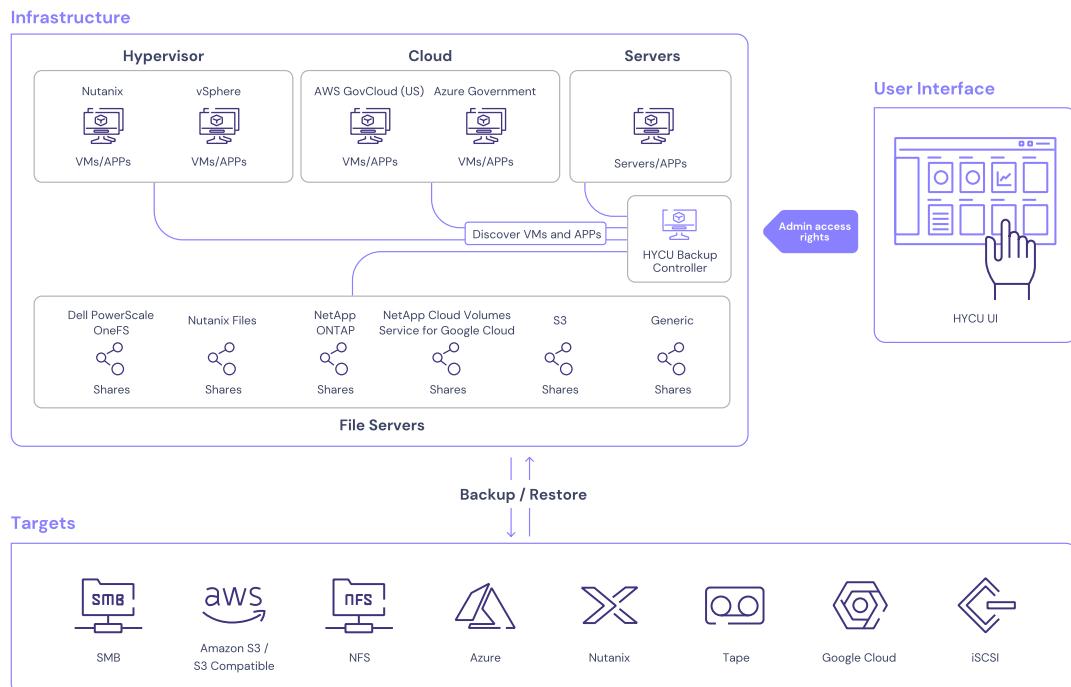


Figure 1-2: HYCU architecture

HYCU data protection

With the HYCU data protection solution, you can be confident that your business data is protected, which means that it is backed up in a consistent state, stored, can be restored, accessed, and is not corrupted.

HYCU enables you to protect virtual machines, applications running on them, file shares on file servers, servers, Nutanix volume groups, and vSphere virtual machine templates. After you establish your data protection environment (that is, add sources, set up targets, and, optionally, create policies), you can enable data protection. After the first backup is successfully completed, you can restore the data if it becomes damaged or corrupted.

Because HYCU is application-aware, when you set credentials for virtual machines or servers, it discovers if any applications are installed and running on them. In addition, it also detects details about the discovered applications such as their versions, the hosts where individual components for the discovered application are installed, and the role of each host.

After you deploy HYCU and establish your data protection environment, depending on what kind of data you want to protect, see one of the following sections:

- “Protecting virtual machines” on page 128
- “Protecting applications” on page 209
- “Protecting file shares” on page 264
- “Protecting volume groups” on page 274

Chapter 2

Deploying the HYCU virtual appliance

The HYCU virtual appliance is a preconfigured software solution that you can easily deploy to a Nutanix AHV cluster, a Nutanix ESXi cluster, a vSphere environment, an AWS GovCloud (US) environment, or an Azure Government environment for which you want to provide data protection. If you want to protect the file shares created by NetApp Cloud Volumes Service for Google Cloud, you can also deploy the HYCU virtual appliance to Google Cloud.

Deployment modes

Mode	Available for...	Description
HYCU Backup Controller	Nutanix AHV cluster	Enables you to protect virtual machines (including virtual machine templates), applications, file shares, servers, and volume groups.
	Nutanix ESXi cluster	
	vSphere environment	
	AWS GovCloud (US) environment	
	Azure Government environment	
	Google Cloud	A HYCU backup controller is a virtual machine that processes data collected from the sources and presents it in its web user interface.
HYCU Instance	Nutanix AHV cluster	Enables you to protect file shares.
	Nutanix ESXi cluster	
	vSphere environment	
	Google Cloud	A HYCU instance is a virtual machine that HYCU uses to perform data protection operations for file servers, taking the load off the HYCU backup controller.
HYCU Manager	Nutanix AHV cluster	Enables you to manage HYCU controllers.

Mode	Available for...	Description
	Nutanix ESXi cluster vSphere environment	HYCU Manager is a virtual machine residing in the source environment that collects data from all HYCU controllers in your on-premises and cloud data protection environments, and presents it in the web user interface.

Deployment tasks

Task	Instructions
1. Size the backup infrastructure for HYCU.	“Sizing resources for your HYCU backup infrastructure” on the next page
2. <i>Only if firewalls are configured on your network.</i> Open relevant ports in each involved firewall.	“Adjusting firewall configuration” on page 23
3. Customize antivirus settings.	“Adjusting antivirus configuration” on page 35
4. Deploy the HYCU virtual appliance to a source.	<ul style="list-style-type: none"> • “Deploying HYCU to a Nutanix AHV cluster” on page 35 • “Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment” on page 40 • “Deploying HYCU to an AWS GovCloud (US) environment” on page 42 • “Deploying HYCU to an Azure Government environment” on page 44 • “Deploying HYCU to Google Cloud” on page 46

After you successfully deploy the HYCU virtual appliance, you can access HYCU by using a supported web browser. For details on how to sign in to HYCU, see [“Signing in to HYCU” on page 49](#).

Sizing resources for your HYCU backup infrastructure

Before you deploy the HYCU virtual appliance, size the resources needed by your HYCU backup infrastructure as follows and ensure that other related requirements are met:

- HYCU virtual machine (HYCU backup controller, HYCU instance, HYCU Manager):
 - Network connection:
Make sure that you reserve an IP address for your virtual machine.
 - System requirements:
 - Minimum requirements are 4 CPU cores and 4 GiB of RAM.
 - The minimum data disk size is at least twice the amount of RAM and the data disk is larger than the OS disk.
 - *For deploying in the HYCU Backup Controller mode:* Keep in mind that aspects beyond the size of your data protection environment affect the system requirements. Performance of the sources, target efficiency, the chosen backup strategy, and backup data compression may all increase or decrease the need for specific resources. For example, if you plan to copy and archive backup data, the number of required targets increases. Similarly, if you specify a short RPO or a small backup threshold, the load on your backup infrastructure increases and HYCU requires more storage and compute resources.

Consider the following recommendations:

Number of VMs in the environment	System requirements				
	vCPU	Cores	Memory	OS disk	Data disk
Fewer than 50	8	1	8 GiB	10 GiB	32 GiB
50–200	8	2	16 GiB	10 GiB	32 GiB
200–500	16	2	32 GiB	10 GiB	50 GiB

Number of VMs in the environment	System requirements				
	vCPU	Cores	Memory	OS disk	Data disk
More than 500	The figures vary. Contact HYCU Support .				

- HYCU web user interface:

For a list of web browsers that you can use to access the HYCU web user interface, see the *HYCU Compatibility Matrix*.

 **Note** HYCU web user interface is designed to work with a screen resolution of at least 1280 × 720 pixels.

- Targets:

For deploying in the HYCU Backup Controller mode: Make sure that destinations you want to use for storing your protected data are available and accessible.

Adjusting firewall configuration

Each deployed HYCU virtual machine includes a firewall with all the necessary ports already open. However, other firewalls installed on your network may block network traffic between specific communication endpoints. For HYCU to operate properly, you must adjust the firewall rules and open the ports listed in the tables in the following sections.

Firewalls installed on the source endpoints see the traffic as outbound, whereas firewalls installed on the destination endpoints see the traffic as inbound. If firewalls are installed elsewhere, they must be adjusted to allow connections in both directions.

Depending on the area of data protection that is relevant for your data protection needs, see the following sections:

- “Targets” on the next page
- “Virtual machine and volume group protection” on page 25
- “Restore of individual files and application awareness” on page 26
- “File share protection” on page 27
- “Infrastructural services” on page 32

- “User interaction and administration” on page 33
- “SpinUp” on page 34

Targets

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup of data to an NFS v4 target	HYCU backup controller	NFS v4 server	2049	TCP UDP
	HYCU instance			
Backup of data to an NFS v3 target	HYCU backup controller	NFS v3 server	111 2049 mountd ^a	TCP UDP
	HYCU instance			
Backup of data to an SMB target	HYCU backup controller	SMB server	445	TCP
	HYCU instance			
Backup of data to an iSCSI target	HYCU backup controller	iSCSI server	3260	TCP
Backup of data to a cloud target	HYCU backup controller	Cloud server	443 ^b	TCP
	HYCU instance			
Archive of data to a QStar NFS target	HYCU backup controller	QStar server	111 2049 mountd ^a 18082 ^c	TCP
	HYCU instance			
Archive of data to a QStar SMB target	HYCU backup controller	QStar server	445 18082 ^c	TCP
	HYCU instance			
Backup of data to a Data Domain target	HYCU backup controller	Data Domain server	2049 ^d	TCP
	HYCU instance			

^a For details on the port number, see NFS server documentation.

^b Cloud targets may utilize multiple IP addresses. For details on IP ranges used by public clouds,

see respective cloud documentation.

^c This is the default port for HTTPS connection, but other ports can also be used. HTTP connection is also supported, but it is not recommended.

^d Data Domain servers by default use port 2049, but other ports can also be used. For instructions on how to change the port, see Dell documentation.

Virtual machine and volume group protection

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Data protection of VMs on a Nutanix cluster or volume groups ^a	HYCU backup controller	Cluster virtual server (cluster virtual IP address) ^b	3205 3260	TCP
		iSCSI target discovery portal (iSCSI Data Services IP address) ^c		
Backup of entities in a vSphere environment	HYCU backup controller	ESXi hosts	902	TCP
		vCenter Server	443	

^a HYCU accesses Nutanix Volumes.

^b Only if a cluster virtual IP address is specified for the Target Portal option in the iSCSI target configuration in HYCU.

^c Only if an iSCSI Data Services IP address is specified for the Target Portal option in the iSCSI target configuration in HYCU.

Restore of individual files and application awareness

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Restore of individual files and application awareness for Linux VMs	HYCU backup controller	VMs	22 ^a	TCP
Restore of individual files and application awareness for Windows VMs	HYCU backup controller	VMs	5985 5986	TCP
Restore from backups created with the Fast Restore policy option enabled	HYCU backup controller	Nutanix Controller VMs	3205	TCP
Restore of applications or files to a Windows VM	VMs	Nutanix iSCSI Data Services	860 3260	TCP
		HYCU backup controller		
Restore of applications or files to a Windows VM if the <code>flr.fast.disable</code> configuration setting is set to <code>true</code>	VMs	HYCU backup controller	445	TCP
Restore of applications or files to a Linux VM	VMs	HYCU backup controller	445	TCP
Restore of applications or files to a Linux VM if the <code>flr.linux.cifs.disable</code> configuration setting is set to <code>true</code>	HYCU backup controller	VMs	22	TCP

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Restore of files to an SMB file share	HYCU backup controller	System with an SMB file share	445	TCP
Restore of files to an NFS file share	HYCU backup controller	System with an NFS file share	NFS4: 2049 NFS3: 111, mountd ^b	TCP
Restore of files to the local machine	System where the HYCU interface is accessed	HYCU backup controller	8443	TCP

^a An SSH server must be installed and configured to use the TCP port 22 for the SSH communication.

^b For details on the port number, see NFS server documentation.

File share protection

General

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup and restore of file shares	HYCU backup controller	HYCU instance	8443	TCP
	HYCU instance	HYCU backup controller		

Nutanix Files

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup and restore of Nutanix Files shares	HYCU instance	Nutanix Files server	445 ^a 2049 ^b 9440	TCP
	HYCU backup controller	Nutanix File Server Virtual Machines (FSVM)	9440 445 ^a 2049 ^b mountd ^c nlockmgr ^c status ^c rquotad ^c	TCP
	HYCU instance		111 ^b	TCP UDP
Deployment and upgrade of HYCU instances ^d	HYCU backup controller	Cluster virtual server (cluster virtual IP address)	9440	TCP
		Nutanix Controller VMs		

^a Only if HYCU accesses file shares by using the SMB protocol.

^b Only if HYCU accesses file shares by using the NFS protocol.

^c *NFSv3 only*. The actual port numbers vary from one vendor to another. The port numbers are customizable. You can run the `rpcinfo -p <NFSServerIP/Hostname>` command on the HYCU backup controller to determine the currently configured ports. For details, see NFS server documentation.

^d HYCU uses the Nutanix REST API v3.

Dell PowerScale OneFS

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup and restore of Dell PowerScale OneFS file shares	HYCU instance	Dell PowerScale OneFS server	445 ^a 2049 ^b 8080	TCP
	HYCU backup controller	IP ranges of the SmartConnect zone assigned to the System zone in the Dell PowerScale OneFS cluster.	8080	TCP
	HYCU instance	IP ranges of the SmartConnect zones assigned to each of the Dell PowerScale OneFS access zones to be backed up.	445 ^a 2049 ^b mountd ^c nlockmgr ^c status ^c rquotad ^c	TCP
	HYCU backup controller	Cluster virtual server (cluster virtual IP address)	111 ^b	TCP UDP
	HYCU instance	Nutanix Controller VMs	9440	
Deployment and upgrade of HYCU instances on a Nutanix AHV cluster	HYCU backup controller	vCenter Server	443	TCP
Deployment and upgrade of HYCU instances on a Nutanix ESXi cluster or in a vSphere environment	HYCU backup controller			

^a Only if HYCU accesses file shares by using the SMB protocol.^b Only if HYCU accesses file shares by using the NFS protocol.^c NFSv3 only. The actual port numbers vary from one vendor to another. The port numbers are

customizable. You can run the `rpcinfo -p <NFSServerIP/Hostname>` command on the HYCU backup controller to determine the currently configured ports. For details, see NFS server documentation.

NetApp ONTAP

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup and restore of NetApp ONTAP file shares	HYCU backup controller	NetApp ONTAP cluster	443	TCP
	HYCU instance			
	HYCU backup controller	NetApp Storage Virtual Machine (SVM)	445 ^a 2049 ^b mountd ^c nlockmgr ^c status ^c rquotad ^c	TCP
	HYCU instance		111 ^b	
Deployment and upgrade of HYCU instances on a Nutanix AHV cluster	HYCU backup controller	Cluster virtual server (cluster virtual IP address)	9440	TCP
		Nutanix controller VM		
Deployment and upgrade of HYCU instances on a Nutanix ESXi cluster or in a vSphere environment	HYCU backup controller	vCenter Server	443	TCP

^a Only if HYCU accesses file shares by using the SMB protocol.

^b Only if HYCU accesses file shares by using the NFS protocol.

^c *NFSv3 only*. The actual port numbers vary from one vendor to another. The port numbers are customizable. You can run the `rpcinfo -p <NFSServerIP/Hostname>` command on the HYCU backup controller to determine the currently configured ports. For details, see NFS server documentation.

NetApp Cloud Volumes Service for Google Cloud

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup and restore of file shares created by NetApp Cloud Volumes Service for Google Cloud	HYCU backup controller	NetApp Cloud Volumes Service for Google Cloud	443	TCP
	HYCU instance		445 ^a 2049 ^b mountd ^c nlockmgr ^c status ^c rquotad ^c	TCP
	HYCU backup controller	NetApp Cloud Volume private IP	111 ^b	TCP UDP
	HYCU instance	NetApp Cloud Volume private IP		

^a Only if HYCU accesses file shares by using the SMB protocol.

^b Only if HYCU accesses file shares by using the NFS protocol.

^c *NFSv3 only*. The actual port numbers vary from one vendor to another. The port numbers are customizable. You can run the `rpcinfo -p <NFSServerIP/Hostname>` command on the HYCU backup controller to determine the currently configured ports. For details, see NFS server documentation.

S3 and generic file shares

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup and restore of S3 file shares	HYCU backup controller	S3 compatible file server	80 443	TCP
	HYCU instance			
Backup and restore of generic file shares	HYCU backup controller	Generic file server host(s)	445 ^a 139 ^a 2049 ^b mountd ^c nlockmgr ^c status ^c rquotad ^c	TCP
	HYCU instance		137 ^a 138 ^a	UDP

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
			111 ^b	TCP UDP

^a Only if HYCU accesses file shares by using the SMB protocol.

^b Only if HYCU accesses file shares by using the NFS protocol.

^c *NFSv3 only.* The actual port numbers vary from one vendor to another. The port numbers are customizable. You can run the `rpcinfo -p <NFSServerIP/Hostname>` command on the HYCU backup controller to determine the currently configured ports. For details, see NFS server documentation.

Infrastructural services

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Use of a DNS server	HYCU backup controller	DNS server	53	TCP UDP
	HYCU instance			
Use of an NTP server	HYCU backup controller	NTP server	123	UDP
	HYCU instance			
Use of an LDAP server	HYCU backup controller	LDAP server	LDAP: 389 LDAPS: 636	TCP
Use of an SMTP server for sending email notifications	HYCU backup controller	SMTP server	25 ^a	TCP
Sharing telemetry data with HYCU	HYCU backup controller	Telemetry host: callhome.hycu.com ^b	443	TCP
		Data host: telemetry-production-bucket.s3.eu-central-1.amazonaws.com ^c		

^a SMTP servers commonly use port 25, but other ports can also be used (for example, 587 or 465).

^b The host name is an alias and resolves to an IP address reported by the DNS server. Keep in mind that the IP address is not static and might change over time.

^c The host name is an alias and resolves to an IP address from an IP address set that is generated from ip-ranges (as published at <https://ip-ranges.amazonaws.com/ip-ranges.json>) filtered by the region (eu-central-1) and the service (S3). Keep in mind that the IP address changes regularly.

User interaction and administration

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Use of the HYCU web user interface	System where the HYCU interface is accessed	HYCU backup controller	8443	TCP
Access to the HYCU backup controller by using SSH	System where the HYCU interface is accessed	HYCU backup controller	22	TCP

SpinUp

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Migrating protected data across on-premises and AWS environments if you use HYCU R-Cloud				
Migrating protected data across on-premises and Google Cloud environments if you use HYCU R-Cloud	HYCU backup controller	HYCU R-Cloud services ^a	443	TCP
Migrating protected data across on-premises and AWS environment if you use HYCU for AWS	HYCU backup controller	HYCU for AWS ^b	443	TCP
Migrating protected data across on-premises and Google Cloud environments if you use HYCU for Google Cloud	HYCU backup controller	HYCU for Google Cloud ^c	443	TCP
Migrating protected data across on-premises and Azure environments	HYCU backup controller	HYCU R-Cloud services ^d	443	TCP

^a The following host names are used: authentication.r-cloud.hycu.com, registry.r-cloud.hycu.com, r-cloud.hycu.com, and the host name of the HYCU R-Cloud manager. For details, contact [HYCU Support](#).

^b The following host names are used: endpoints.aws.hycu.com, registry.aws.hycu.com, authentication.aws.hycu.com, and the host name of the HYCU for AWS manager. For details, contact [HYCU Support](#).

^c The following host names are used: endpoints.hycu.com and the host name of the HYCU for

Google Cloud manager. For details, contact [HYCU Support](#).

^d The following host names are used: registry.azure.hycu.com and the host name of the HYCU for Azure manager. For details, contact [HYCU Support](#).

Adjusting antivirus configuration

HYCU may require access to the files and configuration of the guest operating system to achieve backup and recovery goals of your data protection environment. In this case, the required binary programs and scripts are executed within the virtual machines and you must make sure that your antivirus program allows their execution.

For details on the data protection scenarios when HYCU must be given access to data, see [“Enabling access to data” on page 146](#).

Considerations

- Each time a binary program or a script is to be executed, a new copy of the file is used. Part of the file name is a UUID and a new UUID is generated each time.
- If the antivirus program interferes with HYCU operations, on Windows systems, exclude the HYCU files stored in %ProgramData%\hycu that have no extensions or have the following ones: .bat, .cmd, .exe, .json, .log, .ps1, .txt, or .xml.

Deploying HYCU to a Nutanix AHV cluster

The HYCU virtual appliance is distributed as a virtual disk image that you can easily deploy to a Nutanix AHV cluster by using the Nutanix Prism web console.

Prerequisite

The backup infrastructure is sized according to the requirements described in [“Sizing resources for your HYCU backup infrastructure” on page 22](#).

Consideration

The instructions for deploying HYCU to a Nutanix AHV cluster apply also to a Nutanix Mine cluster.

Deployment tasks

When deploying HYCU to a Nutanix AHV cluster, you must perform the following tasks:

Task	Instructions
1. Upload the HYCU virtual appliance image to a Nutanix AHV cluster.	“Uploading the HYCU virtual appliance image to a Nutanix AHV cluster” below
2. Create a virtual machine for HYCU deployment.	“Creating a virtual machine for HYCU deployment on a Nutanix AHV cluster” on the next page
3. Configure HYCU on the created virtual machine.	“Configuring HYCU on the virtual machine” on page 38

The following flowchart shows an overview of the HYCU deployment tasks:

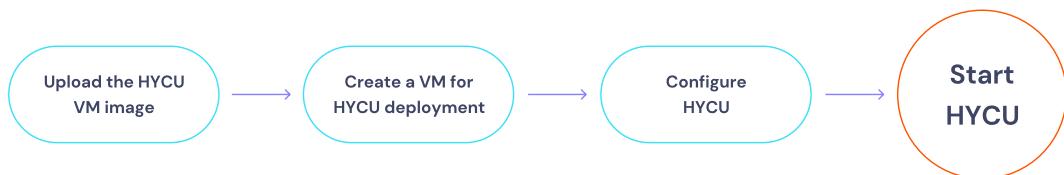


Figure 2-1: Overview of deployment tasks

Uploading the HYCU virtual appliance image to a Nutanix AHV cluster

Procedure

1. Sign in to the Nutanix Prism web console.
2. In the menu bar, click , and then select **Image Configuration**.
3. In the Image Configuration dialog box, click **Upload Image**.
4. In the Create Image dialog box, provide the following information:
 - a. Enter the HYCU image name in the format that should correspond to that of the HYCU image file you are uploading.

! Important The HYCU virtual appliance image must be uploaded to the Nutanix AHV cluster in the following format:

`hycu-<Version>-<Revision>`

For example: `hycu-5.0.0-3634`

If you enter the HYCU image name in a different format, you will not be able to use this image for an upgrade.

- b. *Optional.* Enter an annotation.
- c. From the Image Type drop-down menu, select **DISK**.
- d. From the Storage Container drop-down menu, select a storage container for the image to be uploaded.
- e. In the Image Source section, specify the location of the image file.

5. Click **Save**.

6. Click **Close** after the image is successfully uploaded.

Creating a virtual machine for HYCU deployment on a Nutanix AHV cluster

Procedure

1. In the menu bar in the Nutanix Prism web console, click **Home**, and then select **VM**.
2. Click **Create VM** at the upper right of the screen.
3. In the Create VM dialog box, provide the following information:
 - a. In the General Configuration section, do the following:
 - i. Enter a virtual machine name and, optionally, its description.
 - ii. Make sure that the time zone is set to UTC (the default value).
 - iii. Leave the Use this VM as an agent VM check box clear.
 - b. In the Compute Details section, enter the number of virtual CPUs and cores per virtual CPU, and the amount of memory to allocate to this virtual machine.
 - c. In the Disks section, click **Add New Disk**, and then, in the Add Disk dialog box, specify a system disk:
 - i. From the Type drop-down menu, select **DISK**.
 - ii. From the Operation drop-down menu, select **Clone from Image Service**.
 - iii. From the Bus Type drop-down menu, select **SCSI**.
 - iv. From the Image drop-down menu, select the image you uploaded.
 - v. In the Size (GiB) field, leave the default size of the system disk (10 GiB).

Note You can later increase the size of the system disk if needed. For details, see “[Increasing the size of the HYCU virtual disks](#)” on page 455.

- vi. Click **Add**.
- d. In the Disks section, click **Add New Disk**, and then, in the Add Disk dialog box, specify a data disk:
 - i. Leave the default values for the type of storage device, the device contents, and the bus type.
 - ii. From the Storage Container drop-down menu, select a storage container for the image to be uploaded.
 - iii. In the Size (GiB) field, enter 32.
- Note** You can later increase the size of the data disk if needed. For details, see “[Increasing the size of the HYCU virtual disks](#)” on page 455.
- iv. Click **Add**.
4. In the Network Adapters (NIC) section, click **Add New NIC**, and then select a VLAN and click **Add**.
5. Click **Save**.

Configuring HYCU on the virtual machine

Procedure

1. From the list of virtual machines in the Nutanix Prism web console, select the one you created, and then click **Power on**.
2. When the virtual machine is turned on, click **Launch Console**.
3. In the HYCU Mode Selection dialog box that opens, select one of the following deployment modes:
 - **HYCU Backup Controller**
 - **HYCU Instance**
 - **HYCU Manager**

For details on deployment modes, see “[Deployment modes](#)” on page 20.

4. Tab to **OK** and press **Enter**.
5. In the Network Configuration dialog box that opens, do the following:

a. Enter the values for the following:

- *Optional.* Host name for the virtual machine

The default host name is generated automatically during the HYCU virtual appliance deployment. If you want to use a custom host name, keep in mind the following:

- *Only if you selected the HYCU backup controller or HYCU Manager mode.* The host name should begin with a letter and may contain only letters, numbers, and hyphens (-).
- *Only if you selected the HYCU instance mode.* For host name naming conventions, see “[Managing HYCU instances](#)” on page 381.
- IPv4 address (for example, 10.1.100.1)
- Subnet mask (for example, 255.0.0.0)
- Default gateway (for example, 10.1.1.1)
- *Optional.* DNS server (for example, 10.1.1.5)
- *Optional.* Search domain (for example, domain.com)

 **Note** The domain name should begin with a letter and contain one or more periods. It may also contain only letters, numbers, and hyphens (-).

b. Tab to **OK** and press **Enter**.

The progress of the HYCU configuration displays.

6. *Only if deploying HYCU in the HYCU Instance mode.* In the HYCU Backup Controller dialog box that opens, enter the HYCU backup controller URL and the user name and password you use to access HYCU.

 **Important** If the HYCU backup controller host name cannot be resolved from the HYCU instance (for example, in environments that do not use DNS servers), make sure to use the IP address:

`https://<IPAddress>:<Port>`

The progress of the HYCU backup controller assignment displays.

7. After HYCU is configured, confirm the summary message by pressing **Enter**.

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see “[Licensing](#)” on page 385.

Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment

The HYCU virtual appliance is distributed as an OVF package that you can easily deploy to a Nutanix ESXi cluster or a vSphere environment by using the vSphere (Web) Client.

① Important You can use either the vSphere Web Client or the vSphere Client as the interface for performing the procedure described in this section, unless stated otherwise. As an example, you are guided through the steps that you must perform if you are using the vSphere Web Client.

Prerequisites

- The backup infrastructure is sized according to the requirements described in “[Sizing resources for your HYCU backup infrastructure](#)” on page 22.
- *For deploying HYCU in the HYCU Instance mode:* If your VMware vSphere version is 6.7 Update 3g or later, you can use the vSphere Web Client or the vSphere Client for deployment. Otherwise, the vSphere Web Client must be used.

Considerations

The following applies to the HYCU backup controller time synchronization:

- *For VMware vSphere versions 6.7 or later:* HYCU backup controller uses the time zone as configured on the vCenter Server.
- *For VMware vSphere version 6.5:* The time zone is set to UTC.

Procedure

1. Sign in to the vSphere Web Client.
2. Right-click your vCenter Server, and then select **Deploy OVF Template....** The Deploy OVF Template dialog box opens.
3. In the Select template section, specify the location of the OVF package:

URL	Specify a URL to the HYCU OVF package.
Local file	Browse your file system for the HYCU OVF package. ① Important When you are browsing your file

system, make sure to select both the .ovf file and the .vmdk file related to the OVF package.

Click **Next**.

4. In the Select name and location section, enter a name for the HYCU virtual machine and specify a location where you want to deploy it, and then click **Next**.
5. In the Select a resource section, select where to run the deployed package, and then click **Next**.
6. In the Review details section, verify the package details, and then click **Next**.
7. In the Select Configuration section, do the following:
 - a. Select a deployment configuration:
 - **HYCU Backup Controller**
 - **HYCU Instance**
 - **HYCU Manager**For details on deployment modes, see “Deployment modes” on page 20.
 - b. Click **Next**.
8. In the Select storage section, select where to store the files for the deployed package, and then click **Next**.
9. In the Select networks sections, leave the default values, and then click **Next**.
10. In the Customize template section, enter the values for the following:
 - *Optional.* Host name for the virtual machine
The default host name is generated automatically during the HYCU virtual appliance deployment. If you want to use a custom host name, keep in mind the following:
 - *Only if you selected the HYCU backup controller or HYCU Manager mode.* The host name should begin with a letter and may contain only letters, numbers, and hyphens (-).
 - *Only if you selected the HYCU instance mode.* For host name naming conventions, see “Managing HYCU instances” on page 381.
 - IPv4 address (for example, 10.1.100.1)
 - Subnet mask (for example, 255.0.0.0)
 - Default gateway (for example, 10.1.1.1)

- *Optional.* DNS server (for example, 10.1.1.5)
- *Optional.* Search domain (for example, domain.com)

 **Note** The domain name should begin with a letter and contain one or more periods. It may also contain only letters, numbers, and hyphens (-).

- *Only if deploying HYCU in the HYCU Instance mode.*
 - HYCU backup controller URL

 **Important** If the HYCU backup controller host name cannot be resolved from the HYCU instance (for example, in environments that do not use DNS servers), make sure to use the IP address:

`https://<IPAddress>:<Port>`

- HYCU backup controller user
- HYCU backup controller password

Click **Next**.

11. In the Ready to complete section, review data, and then click **Finish**.

 **Note** Creating the virtual machine may take a few moments. The Power On option is enabled only after the virtual machine is created.

12. From the list of virtual machines, right-click the newly created virtual machine, and then select **Power > Power On** to turn it on.

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see “[Licensing](#)” on page 385.

Deploying HYCU to an AWS GovCloud (US) environment

The HYCU virtual appliance is distributed as a virtual appliance image (a .vmdk file) that you can easily deploy to an AWS GovCloud (US) environment.

Prerequisite

The HYCU virtual appliance image must be uploaded to an Amazon S3 bucket.

Limitation

You can deploy HYCU only in the HYCU Backup Controller mode.

Consideration

When the HYCU backup controller is deployed in AWS GovCloud (US), changing network settings is prevented in HYCU.

Procedure

1. Sign in to the AWS GovCloud (US) console.
2. Import the HYCU virtual appliance image as a snapshot to AWS GovCloud (US). For instructions, see AWS documentation.
3. Create an Amazon Machine Image (AMI) from the snapshot:

```
aws ec2 register-image --name "<HYCUVirtualApplianceImage>"  
--block-device-mappings DeviceName="/dev/sda1",Ebs=  
{SnapshotId=<YourSnapshotID>} --root-device-name "/dev/sda1"
```

ⓘ Important Make sure that you use the following format for the HYCU virtual appliance image:

`hycu-<Version>-<Revision>`

For example: `hycu-5.0.0-363`.

If you enter the HYCU image name in a different format, you will not be able to use this image for an upgrade.

4. Navigate to **EC2**, select **Instances**, and then click **Launch Instances**.
5. On the Choose AMI page, select the HYCU AMI that you created, and then click **Select**.
6. Follow the wizard to create a HYCU backup controller virtual machine from the AMI, making sure that you do the following:
 - a. Add a new volume with the size of at least 32 GiB.
 - b. Create a new firewall rule to allow ingress network traffic on TCP port 8443 from the entire subnetwork to which the HYCU backup controller belongs.

You can leave the default values for the remaining options or adjust them to your needs. For details on all options, see AWS documentation.

7. Review the virtual machine details, and then click **Launch**.

8. Specify a key pair that will be used to connect to your HYCU backup controller. Click **Launch Instances**.
9. Sign in to the HYCU web user interface by specifying the following URL:

```
https://<IPAddress>:8443
```

In this instance, `<IPAddress>` is the external IP address of the newly deployed HYCU backup controller.

If you restart the HYCU backup controller, keep in mind that the IPv4 address is not retained and you must use a new value for `<IPAddress>` to be able to sign in to HYCU (you can find it as Public IPv4 address in the AWS GovCloud (US) console).

ⓘ Important The credentials you provided in AWS GovCloud (US) during virtual machine creation cannot be used to sign in to HYCU. For details on what credentials you can use to sign in to HYCU or to access the HYCU backup controller by using SSH, see “[Signing in to HYCU](#)” on page 49 or “[Accessing the HYCU backup controller virtual machine by using SSH](#)” on page 443.

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see “[Licensing](#)” on page 385.

Deploying HYCU to an Azure Government environment

The HYCU virtual appliance is distributed as a virtual appliance image that you can easily deploy to an Azure Government environment.

Prerequisite

The HYCU virtual appliance image must be uploaded to a storage container in Azure Government.

Limitation

You can deploy HYCU only in the HYCU Backup Controller mode.

Considerations

- When the HYCU backup controller is deployed in Azure Government, changing network settings is prevented in HYCU.
- After deploying HYCU, a warning message stating that the virtual machine agent status is not ready may be displayed in Azure Government. You can safely ignore this message.

Recommendation

It is recommended that you use Azure Storage Explorer to upload the HYCU virtual appliance image to Azure Government. For details, see Azure documentation.

Procedure

1. Sign in to Azure Government.
2. Create a managed image from the HYCU virtual appliance image:
 - a. In the Images navigation pane, click **Create**. In the Create an image menu that opens, make sure you specify the following:
 - In the Instance details section, in the Name field, enter the name of the HYCU virtual appliance image in the following format:
hycu-<Version>-<Revision>
For example, hycu-4.6.0-3634.
 - In the OS disk section, select the following:
 - OS type: **Linux**
 - VM Generation: **Gen 1**
 - b. Click **Review + Create** to review the information, and then click **Create** to create the managed image.
 3. Create a virtual machine from the managed image. Make sure the virtual machine is configured with an additional disk of 32 GiB in size. For details, see Azure documentation.
 4. *Only if you use a network security group.* Create a new firewall rule to allow ingress network traffic on TCP port 8443 from the entire subnetwork to which the HYCU backup controller belongs. For details, see Azure documentation.
 5. Sign in to the HYCU web user interface by specifying the following URL:

```
https://<IPAddress>:8443
```

In this instance, *<IPAddress>* is the external IP address of the newly deployed HYCU backup controller.

ⓘ Important The credentials you provided in Azure Government during virtual machine creation cannot be used to sign in to HYCU. For details on what credentials you can use to sign in to HYCU or to access the HYCU backup controller by using SSH, see “[Signing in to HYCU](#)” on page 49 or “[Accessing the HYCU backup controller virtual machine by using SSH](#)” on page 443.

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see “[Licensing](#)” on page 385.

Deploying HYCU to Google Cloud

The HYCU virtual appliance is distributed as a HYCU virtual appliance image (a `.tar.gz` file) that you can deploy to Google Cloud by using the Google Cloud console.

Limitations

- When HYCU is deployed in Google Cloud, you can only use it to protect file shares created by NetApp Cloud Volumes Service for Google Cloud. For details, see “[Adding NetApp Cloud Volumes Service for Google Cloud](#)” on page 74.
- You can deploy HYCU only in the HYCU Backup Controller mode and the HYCU Instance mode.

Consideration

When the HYCU backup controller is deployed in Google Cloud, changing network settings is prevented in HYCU.

Deployment tasks

When deploying HYCU to Google Cloud, complete the following tasks:

Task	Instructions
1. Create a custom image from the HYCU virtual appliance image.	“Creating a custom image” below
2. Install the HYCU backup controller in Google Cloud.	“Installing the HYCU backup controller in Google Compute Engine” on the next page
3. Create a HYCU instance in Google Cloud.	“Creating a HYCU instance by using the HYCU web user interface” on page 382

Creating a custom image

Prerequisites

- You must be signed in to the Google Cloud console.
- In the Google Cloud console, the project with the file shares created by NetApp Cloud Volumes Service for Google Cloud that you plan to protect must be selected.
- The HYCU virtual appliance image must be uploaded to a Google Cloud bucket. For details, see Google Cloud documentation.

Procedure

1. In the Google Cloud console, navigate to Google Compute Engine, and then click **Images**.
2. Click **Create image**.
3. Under Name, enter the HYCU image name in the following format: `hycu-<Version>-<Revision>`.

! Important The delimiting character in the `<Version>` name field must be a dash (-). For example, `hycu-4-9-0-3634`.
4. Under Source, select **Cloud Storage file**.
5. In the Cloud Storage file field, click **Browse** to select the HYCU virtual appliance image that is stored in your bucket.

6. Under Location, select a region for your custom image.
7. Click **Create** at the bottom of the page.

After you create the image, you can find it under Images in Google Compute Engine.

Installing the HYCU backup controller in Google Compute Engine

Procedure

1. In the Google Cloud console, navigate to Google Compute Engine, and then to VM instances.
2. Click **Create instance**.
3. Under Name, enter a name for your HYCU backup controller.
4. Under Machine configuration, do the following:
 - a. Select **General purpose**.
 - b. Under Series, select **E2**.
 - c. Under Machine Type, click the dropdown menu to select the machine type.
 - d. Select **Standard**, and then **e2-standard-2 (2 vCPU, 8GB memory)**.
5. Under Boot disk, click **Change**. In the Boot disk dialog box that opens, do the following:
 - a. Click **Custom images**.
 - b. Under Image, select your previously created custom image. For details, see “[Creating a custom image](#)” on the previous page.
 - c. Under Size, set the disk size to 10 GB.
 - d. Click **Select**.
6. Under Advanced options, click **Disks**, and then **Add new disk**. In the Add new disk dialog box that opens, set the required parameters:
 - a. Under Name, define the disk name.
 - b. Under Size, set the disk size to 32 GB.
 - c. Go to Attachment settings, and then to Deletion rule. Select **Delete disk** to automatically delete the disk if you later decide to delete the HYCU backup controller VM instance.
 - d. Click **Save**. The Add new disk dialog box closes.
7. On the Create an instance page, click **Create**.

Your HYCU backup controller is now running in Google Cloud. For details on how to sign into the HYCU backup controller, see “[Signing in to HYCU](#)” below.

⚠ Caution Because the HYCU backup controller is publicly accessible in Google Cloud, you should change the default password immediately.

Signing in to HYCU

After you successfully deploy the HYCU virtual appliance, you can access HYCU by using a supported web browser. For a list of supported web browsers, see the *HYCU Compatibility Matrix*.

Procedure

1. In a supported browser, enter the following URL:

```
https://<ServerName>:8443
```

In this instance, `<ServerName>` is the fully qualified domain name of the HYCU server.

For example:

```
https://hycu.example.com:8443
```

2. On the sign-in page, depending on how you want to sign in to HYCU, do one of the following:

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

You can use the default user name (admin) and password (admin) for initial access to HYCU. For security purposes, it is highly recommended that you change the default password. To change the password, click  at the upper right of the screen, and then select **Change Password**.

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

For details on how to integrate HYCU with identity providers, see “[Integrating HYCU with identity providers](#)” on page 373.

3. *Only if two-factor authentication is enabled for your account.* Enter the appropriate two-factor credentials:

- *For using time-based one-time passwords (OTP):* Enter the six-digit authentication code generated by your authentication application (for example, Google Authenticator or a compatible application).
When you sign in for the first time after two-factor authentication was enabled for your account, the OTP backup code is displayed. Scan the QR code with the chosen authentication application or enter the OTP backup code in the application manually and then enter the authentication code generated by your authentication application in the Authentication code field.
- *For using FIDO authenticators:* A security dialog box opens, requesting you to authenticate (for example, by inserting a key). Follow the instructions to authenticate your account.

When you sign in for the first time after two-factor authentication was enabled for your account, a security dialog box opens, prompting you to set up an authenticator (for example a security key or a fingerprint reader). The procedure depends on the selected authenticator and operating system. Follow the instructions to set up the authenticator. For details, see “[Managing FIDO authenticators](#)” on page 450.

 **Note** Keep in mind that the level of access depends on your user permissions. For details, see “[Managing users](#)” on page 343.

After you sign in to the HYCU web user interface, you can configure your environment to use also the HYCU command-line interface (hyCLI). For more information, see “[Using the command-line interface](#)” on page 460.

Setting the language

When you access the HYCU web user interface or the HYCU Manager console, the current browser language is detected and if it is one of the supported languages, the user interface is displayed in that language. If the browser language is not one of the supported languages, the user interface is displayed in English. For a list of supported languages, see the *HYCU Compatibility Matrix*.

Consideration

The HYCU REST API Explorer and the HYCU command-line user interface (hyCLI) are available only in English.

Procedures

- If you are an infrastructure or a self-service group administrator, you can set the preferred language for a user. For instructions, see “[Creating a user](#)” on page 348.
- You can set your preferred language by using the Update Profile option. For instructions, see “[Updating your user profile](#)” on page 360.
- You can set the preferred language for notifications that are sent when events occur. For instructions, see “[Configuring event notifications](#)” on page 302.

You can also change the user interface language by adding a LANG attribute to the URL that you use to access the HYCU web user interface or the HYCU Manager console. For example:

`https://hycu.example.com:8443/#!/login?lang=JA`

Chapter 3

Establishing a data protection environment

After you deploy the HYCU virtual appliance and sign in to HYCU, you must establish a data protection environment in which data will be effectively protected. Establishing the data protection environment involves adding sources, setting up targets, and if your environment requires custom policies, creating them.

The following flowchart explains the tasks you need to perform to establish your data protection environment:

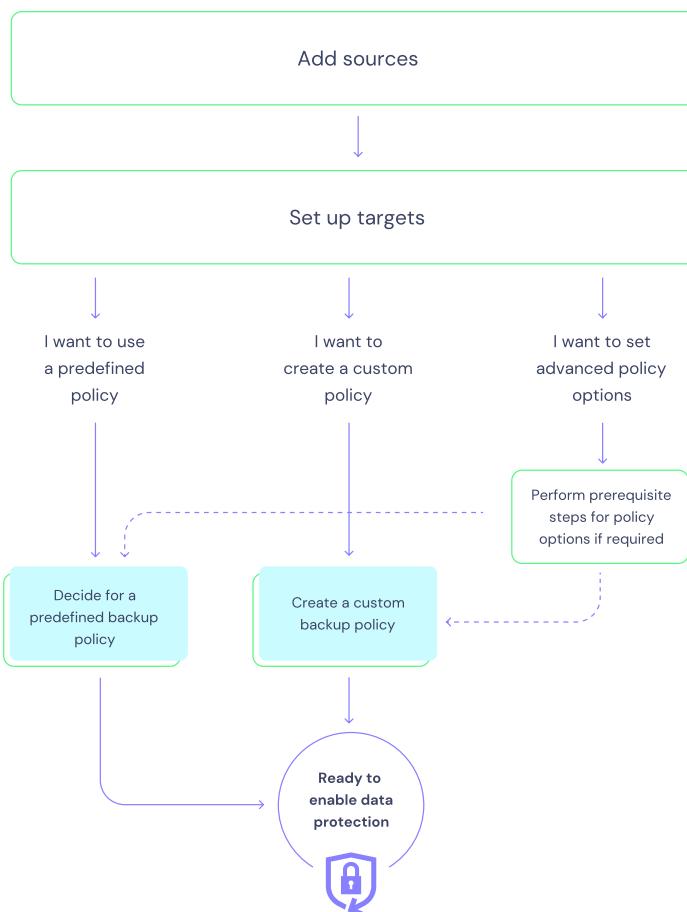


Figure 3-1: Establishing a data protection environment

The tasks that are required to establish a data protection environment can be performed only by an infrastructure group administrator and are as follows:

- “Adding sources” below
- “Setting up targets” on page 78

You can enable data protection by using predefined policies that come with HYCU. If you do not want to use any of them, make sure that you create your own policies. For details, see “[Creating a policy](#)” on page 110.

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

 **Note** Before you start protecting your data protection environment, make sure that the HYCU backup controller is protected. This way, you can quickly recover and resume your data protection activities in case of a disaster. For more information, see “[Preparing for disaster recovery](#)” on [page 132](#).

Adding sources

An environment for which HYCU provides data protection consists of one or more sources that you add to HYCU depending on what kind of data you want to protect—virtual machines or applications running on virtual machines on Nutanix clusters, in vSphere environments, in AWS GovCloud (US) environments, or in Azure Government environments, file shares on file servers, servers, or volume groups on Nutanix clusters. For instructions on how to add a specific source, see one of the following sections:

- “[Adding a Nutanix cluster](#)” on the next page
- “[Adding a vCenter Server](#)” on page 57
- “[Adding a cloud environment](#)” on page 58
- “[Adding a file server](#)” on page 62
- “[Adding a server](#)” on page 78

 **Important** To achieve the optimal performance of your data protection environment and ensure recoverability, make sure to add the source on which the HYCU backup controller is running to HYCU.

Adding a Nutanix cluster

A Nutanix environment consists of one or more Nutanix clusters that host entities (virtual machines on which applications are running and volume groups) for which HYCU provides data protection. Adding a Nutanix cluster to HYCU is the first step to protecting your data.

Prerequisites

- *For Nutanix ESXi clusters:*
 - Your cluster must be registered to the vCenter Server through the Prism web console. For details on how to do this, see Nutanix documentation.
 - A user with specific privileges for vCenter Servers must be specified. For details on which privileges must be assigned to a vSphere user, see [“Assigning privileges to a vSphere user” on page 456](#).
- *Only if you plan to set up automatic policy assignment.* The Nutanix AHV cluster that hosts virtual machines to which you want to automatically assign policies must be registered with Prism Central. For details on how to do this, see Nutanix documentation. For details on automatic policy assignment, see [“Setting up automatic policy assignment” on page 124](#).

Considerations

- *For Nutanix ESXi clusters:*
 - Make sure to use the Nutanix Prism web console to manage virtual machines.
 - Make sure to configure your Windows virtual machines to not go into sleep mode after a certain amount of time. Otherwise, the network settings are not recognized, and consequently such virtual machines cannot be protected by HYCU.
- For backing up virtual machines and volume groups from their replicas in remote office/branch office (ROBO) environments, you must add both the central site Nutanix cluster and the branch office site cluster.

Recommendation

For better performance, it is recommended that an iSCSI Data Service IP address is specified on the Nutanix cluster that you plan to add to HYCU. This automatically enables the Nutanix load balancing feature during data protection operations, which eliminates heavy I/O load on the Nutanix cluster and storage

containers. For details on how to specify an iSCSI Data Service IP address, see Nutanix documentation.

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, click the **Hypervisor** tab, and then click  **New**.
2. Enter the name of the Nutanix cluster in the following URL format:
`https://<ServerName>:<Port>`
3. Enter the user name and password of a user with cluster administrative rights.

 **Important** When adding a Nutanix cluster that has client authentication enabled, make sure you specify the local user.

4. *Only if client authentication is enabled on the Nutanix cluster that you are adding to HYCU.* Use the **Enable certificate authentication** switch, and then browse and upload the trusted CA certificate, the client certificate, and the client private key. Keep in mind the following:
 - The supported certificate file formats are PKCS#1 and PKCS#8.
 - The private key must not be encrypted.

 **Note** If you use Conjur for managing your HYCU secrets, you can enable the **Retrieve values from secrets manager** switch if you want to provide the secret instead of browsing for the file. For details on managing secrets, see “[Managing secrets](#)” on page 404.

By enabling certificate authentication, you allow HYCU to connect to the Nutanix cluster.

5. Click **Next**, and then, depending on the type of Nutanix cluster you are adding, do the following:

Type of Nutanix cluster	Instructions
Nutanix AHV cluster	If you plan to set up automatic policy assignment, in the New Prism Central Credentials dialog box, specify the URL of Prism Central with which your

Type of Nutanix cluster	Instructions
	<p>Nutanix AHV cluster is registered, and the user name and password of a user with cluster administrative rights. Otherwise, leave all the fields blank. Click Next.</p> <p>For details on automatic policy assignment, see “Setting up automatic policy assignment” on page 124.</p>
Nutanix ESXi cluster	<p>In the New vSphere Credentials dialog box, assign the vSphere credentials to the Nutanix ESXi cluster by specifying the URL of the vCenter Server to which the Nutanix ESXi cluster is registered, and the user name and password of a user with specific privileges for vCenter Servers. Click Next.</p> <p> Note After you add a Nutanix ESXi cluster, the  icon next to its type shows that it has the required vCenter Server permissions.</p>

6. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

You can also edit any of the existing Nutanix clusters (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**). When deleting a Nutanix cluster, consider the following:

- You can choose to delete or keep snapshots created by HYCU by using the **Delete snapshots** switch. Keep in mind that if Snapshot is defined as the backup target type in your policy and you choose to delete snapshots, all your backup data will be deleted.
- You can delete any Nutanix cluster, provided it does not have any dependencies. Therefore, it is not possible to delete a Nutanix cluster that is specified as the central site cluster in your policy or a Nutanix cluster that hosts the storage container that is specified in your validation policy until all its dependencies are removed.

Adding a vCenter Server

A vSphere environment consists of ESXi hosts that are managed by vCenter Servers. On each of these ESXi hosts that are included in one or more datacenters, a series of virtual machines running applications reside. Adding one or more vCenter Servers to HYCU is the first step to protecting your virtual machine data. When adding a vCenter Server to HYCU, you can select to add only datacenters that contain the virtual machines that you want to protect.

Prerequisite

A user with specific privileges for vCenter Servers must be specified. For details on which privileges must be assigned to a vSphere user, see [“Assigning privileges to a vSphere user” on page 456](#).

Limitation

Adding vCloud Director or a stand-alone ESXi host is not supported.

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, click the **Hypervisor** tab, and then click  **New**.
2. Enter the name of the vCenter Server in the following URL format:

`https://<vCenterServerFQDN>:<Port>`

The default port for the vCenter Server is 443.

ⓘ Important Make sure you configure the HYCU DNS settings in a way that allows HYCU to resolve this FQDN and, consequently, connect to the vCenter Server and ESXi hosts on which the virtual machines that you want to protect are running.

3. Enter the user name and password of a user with specific privileges for vCenter Servers, and then click **Next**.
4. From the list of all datacenters belonging to the vCenter Server that you are adding, select the datacenters that contain the virtual machines that you want to protect. All the available datacenters are preselected by default.

ⓘ Tip You can also search for a datacenter by entering its name, and then pressing **Enter** in the Search field.

5. Click **Next**.
6. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

You can later do the following:

- Edit any of the existing vCenter Servers. To do so, select the vCenter Server, click  **Edit**, make the required modifications (including the selection of the preferred datacenters), and then click **Save**.
- Delete the vCenter Servers that you do not need anymore as follows:
 1. Select the vCenter Server that you want to delete, and then click  **Delete**.
 2. *Only if you want to delete snapshots created by HYCU.* Enable the **Delete snapshots** switch.

! Important If Snapshot is defined as the backup target type in your policy and you choose to delete snapshots, all your backup data will be deleted.
 3. Click **Delete**.

Adding a cloud environment

Depending on your data protection needs in cloud, you can add the following as a source to HYCU:

Cloud source	I plan to...	Instructions
AWS GovCloud (US) region	Protect virtual machines and applications in an AWS GovCloud (US) environment.	“Adding an AWS GovCloud (US) region” on the next page
Google Cloud project	Manage HYCU instances and upgrade HYCU in Google Cloud as part of protecting file shares created by NetApp Cloud Volumes Service for Google Cloud.	“Adding a Google Cloud project” on page 60
Azure Government subscription	Protect virtual machines and applications in an Azure Government environment.	“Adding an Azure Government subscription” on

Cloud source	I plan to...	Instructions
		page 61

Adding an AWS GovCloud (US) region

An AWS GovCloud (US) environment consists of one or more AWS GovCloud (US) regions that contain virtual machines and applications running on virtual machines for which HYCU provides data protection. Adding one or more AWS GovCloud (US) regions to HYCU is the first step to protecting your data.

Prerequisite

An AWS GovCloud (US) account must be added to HYCU. For instructions, see [“Adding an AWS GovCloud \(US\) account” on page 366](#).

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, click the **Cloud** tab, and then click  **New**.
2. Select **AWS GovCloud (US) region**, and then click **Next**.
3. From the AWS GovCloud (US) account drop-down menu, select the account that has access to the virtual machines and applications that you want to protect.

 **Note** By default, the access key ID and the account ID of the selected AWS GovCloud (US) account are displayed.

4. From the Region drop-down menu, select the AWS GovCloud (US) region that you want to add to HYCU, and then click **Next**.
5. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

You can also do the following:

- View any of the existing AWS GovCloud (US) regions (click  **View**).
- Delete the AWS GovCloud (US) regions that you do not need anymore (click  **Delete**). When deleting an AWS GovCloud (US) region, you can choose to delete or keep snapshots created by HYCU by using the **Delete snapshots** switch.

⚠ Caution If Snapshot is defined as the backup target type in your policy and you choose to delete snapshots, all your backup data will be deleted.

Adding a Google Cloud project

As part of protecting file shares created by NetApp Cloud Volumes Service for Google Cloud, you must add one or more Google Cloud projects as sources to HYCU to be able to manage HYCU instances and upgrade HYCU in Google Cloud.

Prerequisite

A Google Cloud service account must be added to HYCU and it must have granted either the Compute Admin role or a custom role with the following permissions:

compute.addresses.list	compute.instances.deleteAccessConfig
compute.addresses.use	compute.instances.get
compute.disks.create	compute.instances.list
compute.disks.get	compute.instances.setMetadata
compute.disks.use	compute.instances.stop
compute.disks.useReadOnly	compute.machineTypes.get
compute.images.get	compute.subnetworks.get
compute.images.list	compute.subnetworks.use
compute.images.useReadOnly	compute.subnetworks.useExternalIp
compute.instances.create	compute.zoneOperations.get
compute.instances.delete	

For instructions on how to add a Google Cloud service account to HYCU, see [“Adding a Google Cloud service account” on page 367](#).

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, click the **Cloud** tab, and then click  **New**.
2. Select **Google Cloud project**, and then click **Next**.
3. From the Cloud account drop-down menu, select the Google Cloud service account that has access to the project where the HYCU backup controller and HYCU instances (if already created) reside.
After you select the service account, you can view the details about it (its type, project ID, private key ID, client e-mail, and client ID).
4. Click **Next**.
5. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

You can also view any of the existing Google Cloud projects (click  **View**) or delete the ones that you do not need anymore (click  **Delete**).

Adding an Azure Government subscription

An Azure Government environment consists of one or more Azure Government subscriptions that contain virtual machines and applications running on virtual machines for which HYCU provides data protection. Adding one or more Azure Government subscriptions to HYCU is the first step to protecting your data.

Prerequisite

An Azure Government service principal must be added to HYCU. For instructions, see “[Adding an Azure Government service principal](#)” on page 371.

Limitation

You cannot add a subscription whose state is Deleted or Disabled.

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, click the **Cloud** tab, and then click  **New**.
2. Select **Azure Government subscription**, and then click **Next**.

3. From the Service principal drop-down menu, select the service principal that has access to the virtual machines and applications that you want to protect.
4. From the Subscription drop-down menu, select the Azure Government subscription that you want to add to HYCU, and then click **Next**.
5. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

You can also view any of the existing Azure Government subscriptions (click  **View**) or delete the ones that you do not need anymore (click  **Delete**).

When deleting an Azure Government subscription, you can choose to delete or keep snapshots created by HYCU by using the **Delete snapshots** switch. Keep in mind that if Snapshot is defined as the backup target type in your policy and you choose to delete snapshots, all your backup data will be deleted.

Adding a file server

HYCU enables you to protect SMB, NFS, and S3 file shares. You can add specific file server types that are offered by storage providers, S3 compatible file servers, or generic file servers. Using the specific file server types or S3 compatible file servers is recommended due to the following advantages:

- The storage provider APIs enable HYCU to perform incremental backups.
- *For specific file server types:* HYCU can utilize snapshot capabilities to create consistent backups.

You can add one or more file servers that host file shares you want to include in the backup. For information on supported file servers, see the *HYCU Compatibility Matrix*.

Depending on the type of file server that you want to add, see one of the following sections:

- “Adding a Nutanix Files file server” on page 64
- “Adding a Dell PowerScale OneFS file server” on page 66
- “Adding a NetApp ONTAP file server” on page 70
- “Adding NetApp Cloud Volumes Service for Google Cloud” on page 74
- “Adding an S3 compatible file server” on page 76
- “Adding a generic file server” on page 77

 **Note** When adding an S3 compatible file server, keep in mind that you are adding cloud-based object storage in which the buckets are treated as file shares, and the objects as files.

Accessing the Sources dialog box

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

HYCU instances

For protecting file shares, a HYCU instance is introduced in the data protection environment. The HYCU instance is a virtual machine that HYCU uses for performing file server data protection operations, taking the load off the HYCU backup controller.

Before you can start protecting file shares, your HYCU backup controller should have at least one connected HYCU instance that will perform data protection operations. Having more than one HYCU instance is especially useful in environments with a large number of file shares in which HYCU instances can share the load among themselves when performing data protection operations.

You can create a HYCU instance by using one the following methods:

- By deploying the HYCU virtual appliance and selecting the HYCU instance mode. For details, see “[Deploying the HYCU virtual appliance](#)” on page 20.

 **Important** *For Google Cloud:* You cannot use this method when creating a HYCU instance.

- By using the HYCU web user interface. For details, see “[Creating a HYCU instance by using the HYCU web user interface](#)” on page 382.

If you later decide to remove any HYCU instance from your data protection environment, you can do so as described in “[Deleting a HYCU instance](#)” on page 384.

Considerations

- You can create a HYCU instance before or after adding a file server to HYCU.
- The created HYCU instance connects automatically to the corresponding HYCU backup controller.
- Each HYCU instance is by default created with 16 GiB of RAM, 1 CPU, 8 CPU cores, and the data disk size of 128 GiB. However, this can be

overridden by setting the `afs.instance.memory.mb`, `afs.instance.cpu`, `afs.instance.cores.per.cpu`, and `afs.instance.datadisk.size.gb` configuration settings to the preferred values. For details on how to customize the HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on page 523.

- If you change the host name or IP address of the HYCU backup controller, you must also change it on all connected HYCU instances. On each connected HYCU instance, update the `catalog.master.url` configuration setting in the `/hycudata/opt/grizzly/config.properties` file.
- *For HYCU instances running on Nutanix clusters:* When distributing the load among multiple HYCU instances, HYCU automatically prioritizes the HYCU instances that are running on the same Nutanix cluster as the HYCU backup controller and the file server. However, by changing the `afs.instance.afs.cluster.priority` or `afs.instance.bc.cluster.priority` configuration setting, you can adjust the load distribution process to your needs. For details on how to customize the HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on page 523.

 **Note** You can only configure the `afs.instance.afs.cluster.priority` setting when backing up a Nutanix Files server.

- *For HYCU instances running in Google Cloud:* Running multiple HYCU instances in Google Cloud prevents instance overloading and the resulting downtime by evenly distributing the load.

 **Important** To ensure all HYCU instances are instantly available, they must be constantly running. This may cause additional charges to your Google Cloud billing account.

Adding a Nutanix Files file server

Prerequisites

HYCU must have access to the added file server. For the relevant requirements, see “[Enabling HYCU to access a Nutanix Files server](#)” on the next page.

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, click the **File Servers** tab, and then click  **New**.
2. Select **Nutanix Files Server**, and then click **Next**.
3. *Optional.* Enter a custom display name for the file server.
4. Enter the host name or IP address of the file server in the following format:
`https://<Hostname/IP>:<Port>`
 Entering the port is optional if the default value is used (9440).
5.  **Important** If you are providing the host name, make sure the name is unique.
6. Click **Next**.
7. Enable the **Use SMB protocol for accessing shares** switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually.
8. Enable the **Use NFSv4/NFSv3 protocol for accessing shares** switch if you plan to protect NFS file shares.
9. Click **Save**.

You can later do the following:

- Edit any of the existing file servers (click  **Edit** and make the required modifications).
- Delete the file servers that you do not need anymore (click  **Delete**). When deleting a file server, you can choose to delete or keep snapshots created by HYCU by using the **Delete snapshots** switch.

Enabling HYCU to access a Nutanix Files server

To enable HYCU to access a Nutanix Files server, you must prepare the Nutanix Files environment to verify incoming REST API requests.

 **Note** Some versions of Nutanix Prism allow you to manage REST API access permissions through the Manage roles dialog box. For details, see Nutanix documentation.

If this dialog box is not available, create a new user to access the REST API. To do so, follow these steps:

1. Establish a connection to the Nutanix cluster:

```
ssh @<NutanixClusterHostname>
```

2. Run the `ncli fs list` command to list the UUID for the file server.
3. Create a new user:

```
ncli fs add-user uuid=<UUIDFromStep2> user=<Username>
password=<Password>
```

Adding a Dell PowerScale OneFS file server

Prerequisites

- HYCU must have access to the file server you are adding. For details, see [“Dell PowerScale OneFS user permissions” on page 69](#).
- *Only if you plan to protect SMB file shares:*
 - The SMB user must have the run as root permissions in each file share, and must also have the Backup and Restore privileges assigned. If the user belongs to an Active Directory domain, the user must be a member of the Backup Operators group.
 - The SMB options Dot Snap Accessible Child and Dot Snap Accessible Root must be set to Yes.
- *Only if you plan to protect NFS file shares using the NFSv4 protocol.* Root squashing must be disabled on the file server.
- To manage multiple access zones as a single source, the file shares in all access zones must be accessible using a common SMB account.

Considerations

- When adding a Dell PowerScale OneFS server, you can configure HYCU to manage the source in the following ways:
 - Manage individual access zones on your Dell PowerScale OneFS file server as separate sources. This is the advanced option for environments where each access zone uses a customized network and authentication configuration.
 - Manage all access zones or a selection of multiple access zones as a single source. This is the basic option for environments where all the access zones use the same network and authentication configuration.

Note If required, you can add the same Dell PowerScale OneFS file server as a source to the same HYCU backup controller several times, with different access zone network and authentication configurations.

- *Only if you plan to protect SMB file shares.* As part of adding a file server, HYCU tests the SMB credentials using a random SMB file share on the server. If a user does not have permissions for all file shares, the test may fail. The failed test causes the file server adding process to end with a reported error. You can skip the test by changing the `afs.skip.smb.test` configuration setting in the HYCU `config.properties` file to `false`. For details on how to customize the HYCU configuration settings, see [“Customizing HYCU configuration settings” on page 523](#).

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, click the **File Servers** tab, and then click  **New**.
2. Select **Dell PowerScale OneFS Server**, and then click **Next**.
3. Depending on how you want to manage your Dell PowerScale OneFS file server, do the following:

I want to...	Instructions
Use customized network and authentication configurations for each access zone.	<ol style="list-style-type: none"> Make sure the Manage multiple zones switch is disabled. <i>Optional.</i> Enter a custom display name for the file server. In the Zone URL field, enter the host name or IP address of the file server, or the Dell PowerScale OneFS access zone in the following format: <code>https://<Hostname/IP>:<Port></code> Entering the port is optional if the default value is used (8080). <p>Important If you are providing the host name, make sure the name is unique.</p>

I want to...	Instructions
	<p>d. Specify the user name and password of a user with administrative rights for REST API access on the file server, and then click Next.</p> <p>e. <i>Only if the current access zone is not the System zone.</i> Under System zone credentials, enter the System zone URL, and the System zone user name and password.</p> <p>f. Enable the Use SMB protocol for accessing shares switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually.</p> <p>g. Enable the Use NFSv4/NFSv3 protocol for accessing shares switch if you plan to protect NFS file shares.</p>
Use the same network and authentication configuration for all access zones or a selection of multiple access zones.	<p>a. Enable the Manage multiple zones switch.</p> <p>b. <i>Optional.</i> Enter a custom display name for the file server.</p> <p>c. In the System zone URL field, enter the host name or IP address of the Dell PowerScale OneFS System zone in the following format: <code>https://<Hostname/IP>:<Port></code> Entering the port is optional if the default value is used (8080).</p> <p>Important If you are providing the host name, make sure the name is unique.</p> <p>d. Specify the user name and password of a user with administrative rights for REST API access on the file server, and then click Next.</p> <p>e. Under Zones, depending on whether you want to add all access zones or only a selection of</p>

I want to...	Instructions
	<p>specific zones, do one of the following:</p> <ul style="list-style-type: none"> • <i>To add all access zones:</i> Select All zones. • <i>To add only specific access zones:</i> Select Specific zones, and then from the drop-down menu, select the zones that you want to add. <p>f. Enable the Use SMB protocol for accessing shares switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually.</p> <p>g. Enable the Use NFSv4/NFSv3 protocol for accessing shares switch if you plan to protect NFS file shares.</p>

4. Click **Save**.

You can later do the following:

- Edit any of the existing file servers (click  **Edit** and make the required modifications).
- Delete the file servers that you do not need anymore (click  **Delete**). When deleting a file server, you can choose to delete or keep snapshots created by HYCU by using the **Delete snapshots** switch.

Dell PowerScale OneFS user permissions

Depending on the topology of the file shares that you plan to protect, do one of the following:

- *When managing each access zone as a separate source:*
 - Create a user in the System zone and assign this user the following permissions:
 - Job Engine – Write
 - Platform API – Read
 - Snapshot – Write
 - Network – Read

- Create users in the specific zones and assign these users the following permissions:
 - NFS – Read
 - SMB – Read
 - Auth – Read
 - Platform – Read
- *When managing multiple access zones as a single source:* Create a user in the System zone and assign this user the following permissions:
 - Job Engine – Write
 - NFS – Read
 - Platform API – Read
 - SMB – Read
 - Snapshot – Write
 - Network – Read
 - Auth – Read

Adding a NetApp ONTAP file server

Prerequisites

- HYCU must have access to the added file server. See the requirements in the section [“NetApp ONTAP user permissions” on page 73](#).
- *Only if you plan to protect SMB file shares.* The SMB user must have full read and write access to all the file shares that you plan to back up. If the user belongs to an Active Directory domain, the user must be a member of the Backup Operators group.
- *For NetApp ONTAP version 9.10.1 or later:* Client access to the Snapshot copy directory on volumes and shares must be enabled. For details, see NetApp ONTAP documentation.
- *Only if you plan to protect NFS file shares using the NFSv4 protocol.*
 - *For generic NFS file shares:* To be able to list exports for the NFSv4 servers, generic file shares must also support the NFSv3 protocol.

 **Note** If the exports cannot be listed, first-level folders of the global NFSv4 file system will be added as individual shares in HYCU.

- Root squashing must be disabled on the file server.
- The export policy for the file shares that are going to be backed up must

include a rule that allows super user access from the HYCU backup controller, and HYCU instance subnets or IP addresses.

Considerations

- When adding a NetApp ONTAP file server, you can configure HYCU to manage the source in the following ways:
 - Manage all storage virtual machines (SVM) on your NetApp ONTAP file server as a single source. This is the basic option for environments where all the SVMs use the same network and authentication configuration.
 - Manage individual SVMs on your NetApp ONTAP file server as separate sources. This is the advanced option for environments where each SVM uses a customized network and authentication configuration.
- Only if you plan to protect SMB file shares.* As part of adding a file server, HYCU tests the SMB credentials using a random SMB file share on the server. If a user does not have permissions for all file shares, the test may fail. The failed test causes the file server adding process to end with a reported error. You can skip the test by changing the `afs.skip.smb.test` configuration setting in the HYCU `config.properties` file to `false`. For details on how to customize the HYCU configuration settings, see [“Customizing HYCU configuration settings” on page 523](#).

Accessing the Sources dialog box

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

Procedure

- In the Sources dialog box, click the **File Servers** tab, and then click  **New**.
- Select **NetApp ONTAP Server**, and then click **Next**.
- Depending on how you want to manage your NetApp ONTAP file server, do the following:

I want to...	Instructions
Use the same network and authentication configuration for all	<ol style="list-style-type: none"> Make sure that the Manage all SVMs switch is enabled. <i>Optional.</i> Enter a custom display name for the file server.

I want to...	Instructions
SVMs.	<p>c. In the URL field, enter the host name or IP address of the NetApp ONTAP file server in the following format:</p> <p><code>https://<Hostname/IP>:<Port></code></p> <p>Entering the port is optional if the default value is used (8080).</p> <p>! Important If you enter the host name, make sure the name is unique.</p> <p>d. Specify the user name and password of a user with administrative rights for REST API access on the file server, and then click Next.</p> <p>e. Enable the Use SMB protocol for accessing shares switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually.</p> <p>f. Enable the Use NFSv4/NFSv3 protocol for accessing shares switch if you plan to protect NFS file shares.</p>
Use customized network and authentication configurations for each SVM.	<p>a. Disable the Manage all SVMs switch.</p> <p>b. <i>Optional.</i> Enter a custom display name for the file server.</p> <p>c. In the URL field, enter the host name or IP address of the NetApp ONTAP file server in the following format:</p> <p><code>https://<Hostname/IP>:<Port></code></p> <p>Entering the port is optional if the default value is used (8080).</p> <p>! Important If you enter the host name, make sure the name is unique.</p> <p>d. Specify the user name and password of a user</p>

I want to...	Instructions
	<p>with administrative rights for REST API access on the file server, and then click Next.</p> <p>e. From the Managed SVM drop-down menu, select the SVM that hosts the file shares that you want to protect.</p> <p>f. <i>Optional.</i> If you want to filter out certain SMB/NFS IP addresses due to your network environment needs, add the trusted IP addresses to the SVM IP whitelist. Enter a space after each IP address to add multiple IP addresses.</p> <p>g. Enable the Use SMB protocol for accessing shares switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually.</p> <p>h. Enable the Use NFSv4/NFSv3 protocol for accessing shares switch if you plan to protect NFS file shares.</p>

4. Click **Save**.

You can later do the following:

- Edit any of the existing file servers (click  **Edit** and make the required modifications).
- Delete the file servers that you do not need anymore (click  **Delete**). When deleting a file server, you can choose to delete or keep snapshots created by HYCU by using the **Delete snapshots** switch.

NetApp ONTAP user permissions

The following is a list of the required endpoint permissions for the user with administrative rights for REST API access on the file server:

Endpoint	Permissions
/api/cluster	Read-only

Endpoint	Permissions
/api/network/ip/interfaces	Read-only
/api/protocols/cifs/shares	Read-only
/api/storage/volumes	Read and write
/api/svm/svms	Read-only

Adding NetApp Cloud Volumes Service for Google Cloud

Prerequisites

- HYCU must be deployed in the Google Cloud project with the file shares created by NetApp Cloud Volumes Service for Google Cloud that you plan to protect. For details, see [“Deploying HYCU to Google Cloud” on page 46](#).
- HYCU must have access to the added service. For details, see [“Enabling HYCU to access NetApp Cloud Volumes Service for Google Cloud” on the next page](#).
- Only if you plan to protect SMB file shares.* When setting up the Active Directory connection in the Cloud Volumes configuration, make sure that the SMB user is a member of your company's Active Directory Backup Operators group.
- Only if you plan to protect NFS file shares using the NFSv4 protocol.* The export policy for the file shares that you plan to protect must include a rule that gives full access permission to the HYCU backup controller subnet or IP addresses.
- The file shares that you plan to protect must have the Make snapshot directory (.snapshot) visible setting in the Cloud Volumes configuration set to true.
- The HYCU instance must run in the same Google Cloud region as NetApp Cloud Volumes Service for Google Cloud.

Consideration

Only if you plan to protect SMB file shares. As part of adding a file server, HYCU tests the SMB credentials using a random SMB file share on the server. If a user does not have permissions for all file shares, the test may fail. The failed test causes the file server adding process to end with a reported error. You can skip the test by changing the afs.skip.smb.test configuration setting in the HYCU config.properties file to false. For details on how to customize the HYCU

configuration settings, see “[Customizing HYCU configuration settings](#)” on [page 523](#).

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, click the **File Servers** tab, and then click  **New**.
2. Select **NetApp Cloud Volumes Service for Google Cloud**, and then click **Next**.
3. *Optional.* Enter a custom display name for the file server.
4. From the Cloud Account drop-down menu, select the Google Cloud service account. For details on how to add a Google Cloud service account to HYCU, see “[Adding a Google Cloud service account](#)” on [page 367](#).
5. Enable the **Use SMB protocol for accessing shares** switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually.
6. Enable the **Use NFSv4/NFSv3 protocol for accessing shares** switch if you plan to protect NFS file shares.
7. Click **Save**.

You can later do the following:

- Edit any of the existing file servers (click  **Edit** and make the required modifications).
- Delete the file servers that you do not need anymore (click  **Delete**). When deleting a file server, you can choose to delete or keep snapshots created by HYCU by using the **Delete snapshots** switch.

Enabling HYCU to access NetApp Cloud Volumes Service for Google Cloud

To enable HYCU to access NetApp Cloud Volumes Service for Google Cloud, your Google Cloud service account must be granted the NetApp Cloud Volumes Admin and Service Account Token Creator roles or a custom role with the following permissions:

- `cloudvolumesgcp-api.netapp.com/snapshots.create`
- `cloudvolumesgcp-api.netapp.com/snapshots.delete`
- `cloudvolumesgcp-api.netapp.com/snapshots.get`
- `cloudvolumesgcp-api.netapp.com/snapshots.list`
- `cloudvolumesgcp-api.netapp.com/volumes.get`
- `cloudvolumesgcp-api.netapp.com/volumes.list`
- `iam.serviceAccounts.signJwt`
- `resourcemanager.projects.get`

Adding an S3 compatible file server

Prerequisite

The file shares that you want to protect must be owned and managed by the authentication account whose access key you enter when adding the file server.

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, click the **File Servers** tab, and then click  **New**.
2. Select **S3 Compatible Server**, and then click **Next**.
3. *Optional.* Enter a custom display name for the file server.
4. Enter the service endpoint, the access key ID, and the secret access key.
5. Use the **Path style access** switch if you want HYCU to use a path-style URL (`http://<ServerName>:<PortNumber>/<FileShareName>`) to access the protected S3 file shares.

 **Important** Using the Path style access option to access the Cloudian S3 file shares is mandatory.

6. Click **Save**.

You can later do the following:

- Edit any of the existing file servers (click  **Edit** and make the required modifications).
- Delete the file servers that you do not need anymore (click  **Delete**).

When deleting a file server, you can choose to delete or keep snapshots created by HYCU by using the **Delete snapshots** switch.

Adding a generic file server

Prerequisites

- HYCU must have access to the added file server.
- Only if you plan to protect SMB file shares.* The SMB user must have full read and write access to all the file shares that you plan to back up. If the user belongs to an Active Directory domain, the user must be a member of the Backup Operators group.
- Only if you plan to protect NFS file shares using the NFSv4 protocol.* To be able to list exports for the NFSv4 servers, generic file shares must also support the NFSv3 protocol.

 **Note** If the exports cannot be listed, first-level folders of the global NFSv4 file system will be added as individual file shares in HYCU.

Accessing the Sources dialog box

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

Procedure

- In the Sources dialog box, click the **File Servers** tab, and then click  **New**.
- Select **Generic File Server**, and then click **Next**.
- Optional.* Enter a custom display name for the file server.
- Enter the host name of the file server. Do not include `https://` or a port number in the host name.
- Enable the **Use SMB protocol for accessing shares** switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually.
- Enable the **SMB advanced settings** switch to specify a custom SMB port number. By default, the SMB port number is 445. If the destination port of the SMB server differs from the default one, you can specify an alternate port. Valid port numbers are from 0 through 65535.
- Enable the **Use NFSv4/NFSv3 protocol for accessing shares** switch if you

plan to protect NFS file shares.

8. Click **Save**.

You can later do the following:

- Edit any of the existing file servers (click  **Edit** and make the required modifications).
- Delete the file servers that you do not need anymore (click  **Delete**).

Adding a server

Adding one or more servers to HYCU is the first step to protecting your server data.

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, click the **Servers** tab, and then click  **New**.
2. Enter the name of the server.
3. Enter the host name or the IP address of the server.
4. Click **Save**.

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

 **Note** If you delete a server from HYCU and then add it again (with the same name and IP address), keep in mind that this server will be treated as a new one and therefore no old restore points will be available.

Setting up targets

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

Backup target type	Description
Target	<p>Backup data can be stored on the following types of targets: NFS, SMB, Nutanix, Nutanix Objects, iSCSI, Azure (including Azure Government and Azure China), Amazon S3 / S3 Compatible (including AWS Government East and AWS Government West), Google Cloud, tape, and Data Domain.</p> <p>Note A file server can be used as an NFS or SMB target. If you plan to use a file server only as a target and not as a source, there is no need to add it to HYCU.</p> <p>The approach to setting up targets is common for different target types. However, there are specific prerequisites and steps that are required for each target type. Depending on which target you want to set up, see one of the following sections:</p> <ul style="list-style-type: none"> “Setting up an NFS target” on the next page “Setting up an SMB target” on page 82 “Setting up a Nutanix target” on page 85 “Setting up a Nutanix Objects target” on page 88 “Setting up an iSCSI target” on page 90 “Setting up an Azure target” on page 92 “Setting up an Amazon S3 / S3 Compatible target” on page 95 “Setting up a Google Cloud target” on page 99 “Setting up a tape target” on page 101 “Setting up a Data Domain target” on page 105
Snapshot	<p><i>Not available for vSphere virtual machines residing on VMFS or NFS datastores.</i> Backup data is stored as a snapshot on the original location.</p> <p>Important If snapshots created by HYCU are corrupted or unavailable due to a disaster occurring in your data protection environment, you will not be able to restore backup data from this location. However,</p>

Backup target type	Description
	you can still restore your data from targets if data archives exist.

Setting up an NFS target

Prerequisites

- The service must be configured and accessible for the HYCU backup controller and the HYCU instances.
- There must be enough free space on the target for storing the data.
- If deduplication is enabled on the target, the target must be dedicated exclusively to HYCU backups. By dedicating a target exclusively to HYCU backups, you ensure that accurate storage utilization reports are provided.
- If the target resides on Windows, local permissions (security) must be set to **Full Control for Everyone**. If you want to limit access to this system only for HYCU, use the HYCU backup controller IP address for this purpose.
- *For protecting server data:* The target must be accessible from the server.

Limitations

For protecting server data:

- You can store only Linux server backups to this type of target.
- Target encryption and compression are not supported.

Recommendation

It is highly recommended that public access is disabled for a target on which backup data is stored. HYCU automatically detects if public access is enabled for the target and issues a warning message to notify you to adjust the security settings to restrict access to data.

Accessing the Targets panel

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

Procedure

1. In the Targets panel, click **+** **Add**. The Add Target dialog box opens.
2. Select **NFS**, and then click **Next**. The Target Options dialog box opens.
3. Enter a name for the target and, optionally, its description.
4. In the Concurrent Backups field, specify the maximum number of concurrent backups.

If the backup throughput allows, you can specify that more backup jobs run concurrently to reduce the duration of backups and the amount of queued backup jobs.

5. Enable the **Use for archiving** switch if you want this target to be reserved for data archives.

ⓘ Important The target that you use for archiving data cannot be used for backing up data or storing copies of backup data.

6. Use the **Enable compression** switch if you want HYCU to compress backup data before storing it on this target. Compression can be used for backup data, copies of backup data, and data archives.

ⓘ Important Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, archiving of compressed data to targets with enabled compression may increase system requirements for the HYCU backup controller.

7. Click **Next**.
8. *Optional.* In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB). If your target is not dedicated exclusively to HYCU backups, you must leave this field empty.

When this field is left empty, HYCU retrieves the available amount of storage space from the target itself.

ⓘ Note If the target has deduplication enabled, HYCU's estimation of required storage space on the target may be higher than the actual amount of space required on the storage media. Therefore, it is recommended to leave this field empty in such cases.

9. Enter the NFS server name or IP address.

10. Enter the path to the NFS shared folder from the root of the server (for example, /backups/HYCU).
11. Use the **Target encryption** switch if you want the data stored on this target to be encrypted.

 **Note** If you enable target encryption, keep in mind the following:

- The deduplication ratio may be affected by it (in cases where the target has deduplication enabled).
- To be able to import the encrypted target for restoring virtual machines, applications, file shares, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see “[Exporting an encryption key](#)” on page 373.

12. *Only if charges for reading data from the target may apply.* Enable the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges.

With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.

 **Note** If you plan to archive file share data, it is recommended that you enable this option because archiving of file share data is by default performed from the target.

13. Click **Save**.

The target is added to the list of targets. For details on managing targets, see “[Managing targets](#)” on page 325.

Setting up an SMB target

Prerequisites

- The service must be configured and accessible for the HYCU backup controller and the HYCU instances.
- There must be enough free space on the target for storing the data.
- If deduplication is enabled on the target, the target must be dedicated exclusively to HYCU backups. By dedicating a target exclusively to HYCU backups, you ensure that accurate storage utilization reports are provided.

- The supported SMB version must be used. For a list of supported SMB versions, see the *HYCU Compatibility Matrix*.
- *For protecting server data:* The target must be accessible from the server.

Limitations

For protecting server data:

- You can store only Windows server backups to this type of target.
- Target encryption and compression are not supported.

Recommendation

It is highly recommended that public access is disabled for a target on which backup data is stored. HYCU automatically detects if public access is enabled for the target and issues a warning message to notify you to adjust the security settings to restrict access to data.

Accessing the Targets panel

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

Procedure

1. In the Targets panel, click  **Add**. The Add Target dialog box opens.
2. Select **SMB**, and then click **Next**. The Target Options dialog box opens.
3. Enter a name for the target and, optionally, its description.
4. In the Concurrent Backups field, specify the maximum number of concurrent backups.
If the backup throughput allows, you can specify that more backup jobs run concurrently to reduce the duration of backups and the amount of queued backup jobs.
5. Enable the **Use for archiving** switch if you want this target to be reserved for data archives.
Important The target that you use for archiving data cannot be used for backing up data or storing copies of backup data.
6. Use the **Enable compression** switch if you want HYCU to compress backup data before storing it on this target. Compression can be used for backup data, copies of backup data, and data archives.

ⓘ Important Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, archiving of compressed data to targets with enabled compression may increase system requirements for the HYCU backup controller.

7. Click **Next**.
8. *Optional.* In the **Size** field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB). If your target is not dedicated exclusively to HYCU backups, you must leave this field empty.

When this field is left empty, HYCU retrieves the available amount of storage space from the target itself.

ⓘ Note If the target has deduplication enabled, HYCU's estimation of required storage space on the target may be higher than the actual amount of space required on the storage media. Therefore, it is recommended to leave this field empty in such cases.

9. *Optional.* Enter the domain and user credentials.
10. Enter the SMB server name or IP address.
11. Enter the path to the SMB shared folder from the root of the server (for example, /backups/HYCU).
12. Use the **Target encryption** switch if you want the data stored on this target to be encrypted.

ⓘ Note If you enable target encryption, keep in mind the following:

- The deduplication ratio may be affected by it (in cases where the target has deduplication enabled).
- To be able to import the encrypted target for restoring virtual machines, applications, file shares, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see “[Exporting an encryption key](#)” on page 373.

13. *Only if charges for reading data from the target may apply.* Enable the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges.

With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains

this data and for which no additional charges apply. If this is not possible, the data will be read from the target.

 **Note** If you plan to archive file share data, it is recommended that you enable this option because archiving of file share data is by default performed from the target.

14. Click **Save**.

The target is added to the list of targets. For details on managing targets, see “Managing targets” on page 325.

Setting up a Nutanix target

Prerequisite

The Nutanix cluster on which a Nutanix target will be created must be accessible to the HYCU backup controller.

Limitations

- A Nutanix target cannot be used for storing file share data.
- Storing server backup data on this type of target is not supported.

Considerations

- The storage container on a Nutanix cluster that HYCU creates automatically and uses as a Nutanix target must be dedicated exclusively to storing backup data. Because the names of such storage containers start with the HYCU-prefix, make sure not to create your own storage containers with the same prefix. Keep in mind that these storage containers are not available as destinations when restoring data, cloning data, and creating HYCU instances.
- *Only if you plan to employ Nutanix Mine with HYCU.* While adding a Nutanix target, you can also decide to add the related Nutanix cluster as a source to HYCU, if not already added.
- *For Nutanix Mine with HYCU:* In the Nutanix Mine with HYCU dashboard, the Nutanix targets are listed as Mine Storage.

Recommendation

For better performance, it is recommended that an iSCSI Data Service IP address is specified on the Nutanix cluster on which a Nutanix target will be

created. This automatically enables the Nutanix load balancing feature during data protection operations, which eliminates heavy I/O load on the Nutanix cluster and storage containers. For details on how to specify an iSCSI Data Service IP address, see Nutanix documentation.

Accessing the Targets panel

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

Procedure

1. In the Targets panel, click  **Add**. The Add Target dialog box opens.
2. Select **Nutanix**, and then click **Next**. The Target Options dialog box opens.
3. Enter a name for the target and, optionally, its description.
4. In the Concurrent Backups field, specify the maximum number of concurrent backups.

If the backup throughput allows, you can specify that more backup jobs run concurrently to reduce the duration of backups and the amount of queued backup jobs.

5. Enable the **Use for archiving** switch if you want this target to be reserved for data archives.

 **Important** The target that you use for archiving data cannot be used for backing up data or storing copies of backup data.

6. Use the **Enable compression** switch if you want HYCU to compress backup data before storing it on this target. Compression can be used for backup data, copies of backup data, and data archives.

 **Important** Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, archiving of compressed data to targets with enabled compression may increase system requirements for the HYCU backup controller.

7. Click **Next**.
8. *Optional.* In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).

If you leave this field empty, HYCU retrieves the available amount of storage space from the target itself.

9. Enter the name of the Nutanix cluster in the following URL format:

`https://<ServerName>:<Port>`

10. Enter the user name and password of a user with cluster administration rights.

! Important When adding a Nutanix cluster that has client authentication enabled, make sure that you specify credentials of a local user.

11. Use one or more of the following switches if you want to enable the respective Nutanix options on the storage container to increase your Nutanix cluster's effective storage capacity:

- **Deduplication**
- **Erasure coding**
- **Hardware compression**

For more information on these options, see Nutanix documentation.

12. Use the **Target encryption** switch if you want the data stored on this target to be encrypted.

! Note If you enable target encryption, keep in mind the following:

- Enabling target encryption in combination with options intended to increase your cluster's effective storage capacity will prevent such options from taking effect.
- To be able to import the encrypted target for restoring virtual machines, applications, and volume groups, export the encryption key to a file and keep this file on safe. For instructions, see “[Exporting an encryption key](#)” on page 373.

13. *Only if charges for reading data from the target may apply.* Enable the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges.

With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.

14. Click **Save**.

The target is added to the list of targets. For details on managing targets, see “[Managing targets](#)” on page 325.

Setting up a Nutanix Objects target

Prerequisites

- The service must be configured and accessible.
- If you want to provide secure HTTPS access, import the CA certificate/chain to HYCU. For details, see “[Importing a custom certificate](#)” on page 435.

Limitations

- Storing backup data to targets on which expiration for HYCU objects and versions is enabled in a lifecycle policy is not supported.
- Storing server backup data on this type of target is not supported.

Consideration

A Nutanix Objects target that has WORM enabled is represented by the  icon in the list of targets.

Accessing the Targets panel

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

Procedure

1. In the Targets panel, click  **Add**. The Add Target dialog box opens.
2. Select **Nutanix Objects**, and then click **Next**. The Target Options dialog box opens.
3. Enter a name for the target and, optionally, its description.
4. In the Concurrent Backups field, specify the maximum number of concurrent backups.

If the backup throughput allows, you can specify that more backup jobs run concurrently to reduce the duration of backups and the amount of queued backup jobs.

5. Enable the **Use for archiving** switch if you want this target to be reserved for data archives.

Important The target that you use for archiving data cannot be used for backing up data or storing copies of backup data.

6. Use the **Enable compression** switch if you want HYCU to compress backup data before storing it on this target. Compression can be used for backup data, copies of backup data, and data archives.

ⓘ Important Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, backing up, copying backup data, and archiving to targets with enabled compression and archiving of compressed data may increase system requirements for the HYCU backup controller.

7. Click **Next**.
8. In the **Size** field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
If you leave this field empty, HYCU retrieves the available amount of storage space from the target itself.
9. Provide the following information:

Required information	Notes
Service endpoint	You must enter the full service endpoint URL, including the HTTP or HTTPS protocol.
Bucket name	Specify the name of the bucket. If the bucket does not exist, HYCU will create it automatically.
Access key ID	The access key ID and the secret access key are used to authenticate S3 REST API service calls.
Secret access key	

10. Use the **Path style access** switch if you want HYCU to use a path-style URL (`https://<ServiceEndpointURL>/<BucketName>`) to access the bucket. HYCU by default uses a virtual-hosted-style URL (`https://<BucketName>.<ServiceEndpointURL>`).
11. Use the **Target encryption** switch if you want the data stored on this target to be encrypted.

ⓘ Note To be able to import the encrypted target for restoring virtual machines, applications, file shares, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see “[Exporting an encryption key](#)” on page 373.

12. *Only if charges for reading data from the target may apply.* Enable the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges.

With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.

 **Note** If you plan to archive file share data, it is recommended that you enable this option because archiving of file share data is by default performed from the target.

13. Click **Save**.

The target is added to the list of targets. For details on managing targets, see “[Managing targets](#)” on page 325.

Setting up an iSCSI target

Prerequisites

- The service must be configured and accessible.
- The target has not been initialized yet.
- The HYCU iSCSI Initiator secret must be added on the iSCSI server if you want to enable mutual authentication between HYCU and the iSCSI server.

Limitations

- An iSCSI target cannot be used for storing file share data.
- Storing server backup data on this type of target is not supported.

Considerations

- If you have more than one volume created on the selected iSCSI target, HYCU uses the disks from all the volumes that it can access for storing data.
- Nutanix volume groups used as iSCSI targets automatically discard unused blocks. For other types of iSCSI targets, this option can be added manually. For details, contact [HYCU Support](#).

Accessing the Targets panel

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

Procedure

1. In the Targets panel, click **+** **Add**. The Add Target dialog box opens.
2. Select **iSCSI**, and then click **Next**. The Target Options dialog box opens.
3. Enter a name for the target and, optionally, its description.
4. In the Concurrent Backups field, specify the maximum number of concurrent backups.

If the backup throughput allows, you can specify that more backup jobs run concurrently to reduce the duration of backups and the amount of queued backup jobs.

5. Enable the **Use for archiving** switch if you want this target to be reserved for data archives.

Important The target that you use for archiving data cannot be used for backing up data or storing copies of backup data.

6. Use the **Enable compression** switch if you want HYCU to compress backup data before storing it on this target. Compression can be used for backup data, copies of backup data, and data archives.

Important Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, archiving of compressed data to targets with enabled compression may increase system requirements for the HYCU backup controller.

7. Click **Next**.
8. *Optional.* In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
If you leave this field empty, HYCU retrieves the available amount of storage space from the target itself.
9. Enter the target portal IP address and the target name.
Note If data from sources other than HYCU resides on the storage device, such a target cannot be set for HYCU backups.
10. If the iSCSI server requires CHAP authentication, in the CHAP section, do the following:
 - a. Use the **CHAP** switch to enable CHAP authentication, and then provide a user name and the target secret (the security key) for the user's account

to access the iSCSI server.

- b. Use the **Perform mutual authentication** switch if you want the iSCSI target to be authenticated by HYCU. In this case, the HYCU iSCSI Initiator secret must be specified on the iSCSI server. For details about setting the iSCSI Initiator secret, see “[Setting the iSCSI Initiator secret](#)” on [page 384](#).
11. Use the **Target encryption** switch if you want the data stored on this target to be encrypted.

! Important To be able to import the encrypted target for restoring virtual machines, applications, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see “[Exporting an encryption key](#)” on [page 373](#).

12. *Only if charges for reading data from the target may apply.* Enable the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges.

With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.

13. Click **Save**.

The target is added to the list of targets. For details on managing targets, see “[Managing targets](#)” on [page 325](#).

Setting up an Azure target

Prerequisites

- The service must be configured and accessible.
- *For Azure targets that have the immutability policy (WORM) set:* Blob versioning must be enabled for the storage account, and version-level immutability support must be enabled on the container. For details, see Azure documentation.

Limitations

- Storing server backup data on this type of target is not supported.
- Backing up data to a target that has a hierarchical namespace enabled is not supported.

Considerations

- *For virtual machines and volume groups:* Your data on the Azure target can be stored in the hot, cool, and archive storage tiers. When restoring data archives, HYCU performs data rehydration during which the tier of the Blob object storage is changed from the archive storage tier to the hot storage tier. Keep in mind that this can take a few hours to complete. HYCU moves data back to the archive storage tier afterward.
- *For virtual machines, applications, and volume groups:* HYCU automatically moves a data archive that has a retention period set to at least 180 days from the Azure cool or hot storage tier to the archive storage tier during the next archive synchronization. By moving data archives to the archive storage tier, HYCU ensures your data is stored most cost-efficiently because the archive storage tier is optimized for storing data that is not accessed frequently and is stored for at least 180 days.
- *For file shares:* HYCU does not automatically move file share data archives to the archive storage tier. If after backing up file share data you move the data yourself to the archive storage tier, you must also make sure to move it back to the cool or hot storage tier before you can restore data.
- An Azure target that has the immutability policy (WORM) set is represented by the  icon in the list of targets.

Recommendations

- It is highly recommended that public access is disabled for a target on which backup data is stored. HYCU automatically detects if public access is enabled for the target and issues a warning message to notify you to adjust the security settings to restrict access to data.
- To enable HYCU to delete the temporary containers that are created during backup and restore operations, it is recommended that you do not set up Azure targets that belong to storage accounts on which version-level immutability support is enabled.

Accessing the Targets panel

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

Procedure

1. In the Targets panel, click  **Add**. The Add Target dialog box opens.
2. *Only if you want to set up an Azure target.* Select **AZURE**, and then click **Next**.

3. Only if you want to set up an Azure Government or Azure China target.
 - a. Click **Other**, and then select the relevant target type.
 - b. Click **Next**.
4. Enter a name for the target and, optionally, its description.
5. In the Concurrent Backups field, specify the maximum number of concurrent backups.

If the backup throughput allows, you can specify that more backup jobs run concurrently to reduce the duration of backups and the amount of queued backup jobs.
6. Enable the **Use for archiving** switch if you want this target to be reserved for data archives.

Important The target that you use for archiving data cannot be used for backing up data or storing copies of backup data.
7. Use the **Enable compression** switch if you want HYCU to compress backup data before storing it on this target. Compression can be used for backup data, copies of backup data, and data archives.

Important Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, backing up, copying backup data, and archiving to targets with enabled compression and archiving of compressed data may increase system requirements for the HYCU backup controller.
8. Click **Next**.
9. In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
10. Enter the storage account name, the secret access key, and the container name.

Note If the container does not exist, it is created automatically.
11. Use the **Target encryption** switch if you want the data stored on this target to be encrypted.

Note To be able to import the encrypted target for restoring virtual machines, applications, file shares, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see “Exporting an encryption key” on page 373.

12. Only if charges for reading data from the target may apply. Use the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges.

With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.

 **Note** If you plan to archive file share data, it is recommended that you keep this option enabled because archiving of file share data is by default performed from the target.

13. Click **Save**.

The target is added to the list of targets. For details on managing targets, see “Managing targets” on page 325.

Setting up an Amazon S3 / S3 Compatible target

Prerequisites

- The service must be configured and accessible.
- The S3 bucket must be created and configured in AWS or any other supported S3 compatible environment. For a list of supported S3 compatible cloud storage solutions, see the *HYCU Compatibility Matrix*.
- The following minimum required Amazon S3 permissions must be specified:
 - General permissions:
`s3:GetObject, s3:GetObjectRetention, s3:DeleteObject,`
`s3:PutObject, s3>ListBucket, s3:GetBucketAcl,`
`s3>ListBucketMultipartUploads, s3:GetBucketLocation,`
`s3:GetBucketObjectLockConfiguration, s3>DeleteObjectVersion,`
`s3>ListBucketVersions, and s3:GetBucketVersioning.`
 - Additional permissions:
 - For Amazon S3 targets: `s3:GetBucketPublicAccessBlock`.
 - For S3 compatible targets: `s3>ListMultipartUploadParts` and `s3:AbortMultipartUpload`.
 - For targets that have Object Lock (WORM) enabled:
`s3:PutObjectRetention, s3:PutObjectTagging`, and

`s3:GetObjectTagging`.

- *For Wasabi S3 compatible targets that have Object Lock (WORM) enabled:* `s3>ListMultipartUploadParts`, `s3:AbortMultipartUpload`, and `s3:GetObjectVersion`.
- *Only if you plan to store data to an Amazon S3 target in Amazon Virtual Private Cloud (VPC).* An interface VPC endpoint must be set up.
- *For S3 compatible targets:* If you want to provide secure HTTPS access, import the CA certificate/chain to HYCU. For details, see [“Importing a custom certificate” on page 435](#).
- *For setting up a Tencent Cloud target:* Make sure the service endpoint URL does not contain the bucket name. For example, if the Tencent Cloud access domain is `https://testbucket-1234567890.cos.ap-chengdu.myqcloud.com`, in the HYCU Service endpoint field, enter the URL without the bucket name:

`https://cos.ap-chengdu.myqcloud.com`

Limitations

- HYCU does not support Amazon S3 targets that use the Glacier Flexible Retrieval and Glacier Deep Archive storage classes.
- HYCU currently supports only AWS Signature Version 4.
- Storing backup data to targets on which expiration for HYCU objects and versions is enabled in a lifecycle policy is not supported.
- Storing server backup data on this type of target is not supported.
- *For Wasabi S3 compatible targets that have Object Lock (WORM) disabled:* Backing up data to such targets is supported only if compliance mode is not enabled.

Considerations

- Amazon S3, Cloudian S3 compatible, Scality S3 compatible, and Wasabi S3 compatible targets that have Object Lock (WORM) enabled are represented by the  icon in the list of targets.
- If the service endpoint that you plan to specify when adding the target is not an Amazon S3 endpoint, check with your data storage vendor if setting a storage class is supported.

Recommendations

It is highly recommended that public access is disabled for a target on which backup data is stored. HYCU automatically detects if public access is enabled for the target and issues a warning message to notify you to adjust the security settings to restrict access to data.

Accessing the Targets panel

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

Procedure

1. In the Targets panel, click  **Add**. The Add Target dialog box opens.
2. *Only if you want to set up an Amazon S3 or S3 compatible target.* Select **Amazon S3 / S3 Compatible**, and then click **Next**.
3. *Only if you want to set up an AWS Government East or AWS Government West target.*
 - a. Click **Other**, and then select the relevant target type.
 - b. Click **Next**.
4. Enter a name for the target and, optionally, its description.
5. In the Concurrent Backups field, specify the maximum number of concurrent backups.

If the backup throughput allows, you can specify that more backup jobs run concurrently to reduce the duration of backups and the amount of queued backup jobs.

6. Enable the **Use for archiving** switch if you want this target to be reserved for data archives.

 **Important** The target that you use for archiving data cannot be used for backing up data or storing copies of backup data.

7. Use the **Enable compression** switch if you want HYCU to compress backup data before storing it on this target. Compression can be used for backup data, copies of backup data, and data archives.

 **Important** Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, backing up, copying backup data, and archiving to targets with enabled compression and archiving of compressed data may

increase system requirements for the HYCU backup controller.

8. Click **Next**.
9. In the **Size** field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
10. Enter the service endpoint URL.
11. From the **Storage Class** drop-down menu, select the storage class that you want to use for storing the data. If you leave the **Default** option selected, Amazon S3 selects the storage class autonomously.

Important Different storage classes incur different data storage and retrieval costs. For details about Amazon S3 storage classes, see AWS documentation.

12. Enter the bucket name, the access key ID, and the secret access key. The access key and the secret access key are used to authenticate the Amazon API service calls.
13. Use the **Path style access** switch if you want HYCU to use a path-style URL (`https://s3.amazonaws.com/<BucketName>`) to access the bucket. HYCU by default uses a virtual-hosted-style URL (`https://<BucketName>.s3.amazonaws.com`).
14. Use the **Target encryption** switch if you want the data stored on this target to be encrypted.

Note To be able to import the encrypted target for restoring virtual machines, applications, file shares, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see “[Exporting an encryption key](#)” on page 373.

15. *Only if charges for reading data from the target may apply.* Use the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges.

With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.

Note If you plan to archive file share data, it is recommended that you keep this option enabled because archiving of file share data is by default performed from the target.

16. Click **Save**.

The target is added to the list of targets. For details on managing targets, see [“Managing targets” on page 325](#).

Setting up a Google Cloud target

Prerequisites

- A Google Cloud service account must be created and then added to HYCU. For instructions on how to add a cloud account to HYCU, see [“Adding a Google Cloud service account” on page 367](#).
- A Google Cloud Storage bucket must be created in the project that is linked to the created Google Cloud service account you added to HYCU.
- The service must be configured and accessible.
- *Only if Bucket Lock and Object Retention Lock, or Bucket Lock alone are enabled on the target (WORM)*. The Google Cloud service account that you add to HYCU must have the following permissions granted in Google Cloud:
 - *For Bucket Lock:* `storage.buckets.create`, `storage.buckets.delete`, `storage.buckets.get`, `storage.buckets.update`, `storage.objects.get`, and `storage.objects.update`.
 - *For Object Retention Lock:* The same permissions as for Bucket Lock with the additional `storage.objects.setRetention` permission.

Limitation

Storing server backup data on this type of target is not supported.

Considerations

- To ensure your data is stored most cost-efficiently, HYCU stores data in the Google Cloud storage class that is optimal for the retention period set in your policy. Therefore, data can be stored in a different storage class than the one set as the bucket's default storage class. However, if the bucket's default storage class is set to standard, backup data and copies of backup data are always stored in the standard storage class.
- Each data archive that has a retention period set to at least 365 days is automatically moved to the Google Cloud archive storage class during the next archive synchronization.
- Google Cloud targets that have Bucket Lock and Object Retention Lock, or Bucket Lock alone enabled are WORM-compliant and represented by the  icon in the list of targets.

Recommendations

- It is highly recommended that public access is disabled for a target on which backup data is stored. HYCU automatically detects if public access is enabled for the target and issues a warning message to notify you to adjust the security settings to restrict access to data.
- *Only if you plan to enable WORM on the target.* Enabling retention for both the bucket (Bucket Lock) and the objects (Object Retention Lock) in Google Cloud is recommended.

Accessing the Targets panel

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

Procedure

1. In the Targets panel, click  **Add**. The Add Target dialog box opens.
2. Select **Google Cloud**, and then click **Next**. The Target Options dialog box opens.
3. Enter a name for the target and, optionally, its description.
4. In the Concurrent Backups field, specify the maximum number of concurrent backups.
If the backup throughput allows, you can specify that more backup jobs run concurrently to reduce the duration of backups and the amount of queued backup jobs.
5. Enable the **Use for archiving** switch if you want this target to be reserved for data archives.

 **Important** The target that you use for archiving data cannot be used for backing up data or storing copies of backup data.
6. Use the **Enable compression** switch if you want HYCU to compress backup data before storing it on this target. Compression can be used for backup data, copies of backup data, and data archives.

 **Important** Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, backing up, copying backup data, and archiving to targets with enabled compression and archiving of compressed data may increase system requirements for the HYCU backup controller.

7. Click **Next**.
8. In the **Size** field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
9. In the **Bucket Name** field, enter the bucket name.

 **Note** The specified bucket should be created in a project that is linked to the Google Cloud service account you added to HYCU.

10. From the **Cloud Account** drop-down menu, select the Google Cloud service account you added to HYCU.
11. Use the **Target encryption** switch if you want the data stored on this target to be encrypted.

 **Note** To be able to import the encrypted target for restoring virtual machines, applications, file shares, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see “[Exporting an encryption key](#)” on page 373.

12. *Only if charges for reading data from the target may apply.* Use the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges.

With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.

 **Note** If you plan to archive file share data, it is recommended that you keep this option enabled because archiving of file share data is by default performed from the target.

13. Click **Save**.

The target is added to the list of targets. For details on managing targets, see “[Managing targets](#)” on page 325.

Setting up a tape target

HYCU supports using tape to archive data that you intend to keep for a longer period of time through Integral Volume sets provided by QStar Archive Storage Manager (ASM).

Prerequisites

- The licensed capacity must be sufficient for storing archive data.
- The QStar cache must be large enough.
- There must be enough free space for storing archive data on QStar.

For details, see QStar documentation.

Limitations

- Target compression is not supported—archive data cannot be compressed before it is stored on the target.
- A tape target cannot be used for storing the following kind of data: individual files, whole applications, SQL Server databases, Exchange Server databases, mailboxes, and public folders, Oracle database instances and tablespaces, and file shares.

Considerations

- Make sure to use a tape target only for storing archive data.
- Each Integral Volume set is treated as a separate target in HYCU.

Procedure

1. In the Targets panel, click  **Add**. The Add Target dialog box opens.
2. Select **QStar NFS** or **QStar SMB**, and then click **Next**. The Target Options dialog box opens.
3. Enter a name for the target and, optionally, its description.
4. In the Concurrent Backups field, specify the maximum number of concurrent archive jobs. You can specify several archive jobs to run concurrently to reduce the duration of archiving data and the amount of queued archive jobs.

ⓘ Important You must ensure that the QStar cache is large enough to support concurrent archive operations. Keep in mind that specifying several archive jobs to run concurrently may also increase system requirements for the HYCU backup controller.
5. Make sure the **Use for archiving** option is enabled.
6. Make sure the **Enable compression** option is disabled.
7. Click **Next**.
8. Follow the instructions that are relevant for your tape target type:

Target type	Instructions
QStar NFS	<p>a. <i>Optional.</i> In the Size field, enter the maximum space that should be reserved for archive data (in MiB, GiB, or TiB).</p> <p>b. Provide user credentials that HYCU will use to access the shared folder and make web service calls.</p> <p>c. Enter the name of the Integral Volume set where you want to archive data.</p> <p>d. Provide the web service information. If the default port is used and HTTPS access to the QStar server is configured, enter the host name of the QStar server. Otherwise, specify the URL that will be used to access the QStar server in the following format:</p> <p><code>https://<QStarServer>:<Port></code></p> <p>e. <i>Optional.</i> Enter the path to the shared folder of the mounted Integral Volume set. If you leave this field empty, HYCU tries to retrieve the path to the shared folder.</p> <p>f. Use the Target encryption switch if you want the data stored on this target to be encrypted.</p> <p>Note If you enable target encryption, keep in mind the following:</p> <ul style="list-style-type: none"> • The compression ratio may be affected by it (in cases where tape compression is enabled). • To be able to import the encrypted target for restoring virtual machines, applications, file shares, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see “Exporting an encryption key” on page 373. <p>g. <i>Only if charges for reading data from the target may apply.</i> Use the Metered target switch if you want HYCU to try to read the data from other locations first to avoid additional charges.</p> <p>With the Metered target switch enabled, HYCU will try</p>

Target type	Instructions
QStar SMB	<p>to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.</p> <p>Note If you plan to archive file share data, it is recommended that you keep this option enabled because archiving of file share data is by default performed from the target.</p> <ol style="list-style-type: none"> a. <i>Optional.</i> In the Size field, enter the maximum space that should be reserved for archive data (in MiB, GiB, or TiB). b. <i>Optional.</i> Specify the domain in which the account that has access permissions on the shared folder is registered. c. Provide user credentials that HYCU will use to access the shared folder and make web service calls. d. Enter the name of the Integral Volume set where you want to archive data. e. Provide the web service endpoint information. If the default port is used and HTTPS access to the QStar server is configured, enter the host name of the QStar server. Otherwise, specify the URL that will be used to access the QStar server in the following format: <code>https://<QStarServer>:<Port></code> f. <i>Optional.</i> Enter the path to the shared folder of the mounted Integral Volume set. If you leave this field empty, HYCU tries to retrieve the path to the shared folder. g. Use the Target encryption switch if you want the data stored on this target to be encrypted. <p>Note If you enable target encryption, keep in mind the following:</p> <ul style="list-style-type: none"> • The compression ratio may be affected by it (in cases where tape compression is enabled).

Target type	Instructions
	<ul style="list-style-type: none"> • To be able to import the encrypted target for restoring virtual machines, applications, file shares, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see “Exporting an encryption key” on page 373. <p data-bbox="514 579 1308 736">h. <i>Only if charges for reading data from the target may apply.</i> Use the Metered target switch if you want HYCU to try to read the data from other locations first to avoid additional charges.</p> <p data-bbox="562 750 1308 961">With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.</p> <p data-bbox="593 974 1276 1147">Note If you plan to archive file share data, it is recommended that you keep this option enabled because archiving of file share data is by default performed from the target.</p>

9. Click **Save**.

After you create a tape target, it is added to the list of targets and represented by the  icon.

Setting up a Data Domain target

Prerequisites

- DD Boost must be enabled on the Data Domain system.
- A storage unit must be created on the Data Domain system.
- *Only if retention lock is applied to data on your Data Domain system.* Automatic retention lock must be enabled. For details, see Dell documentation.

Limitation

Storing server backup data on this type of target is not supported.

Consideration

A Data Domain target that has retention lock (WORM) enabled is represented by the  icon in the list of targets.

Accessing the Targets panel

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

Procedure

1. In the Targets panel, click  **Add**. The Add Target dialog box opens.
2. Select **Data Domain**, and then click **Next**.
3. Enter a name for the target and, optionally, its description.
4. In the Concurrent Backups field, specify the maximum number of concurrent backups.

If the backup throughput allows, you can specify that more backup jobs run concurrently to reduce the duration of backups and the amount of queued backup jobs.

5. Enable the **Use for archiving** switch if you want this target to be reserved for data archives.

 **Important** The target that you use for archiving data cannot be used for backing up data or storing copies of backup data.

6. Click **Next**.
7. In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
8. Enter the host name or IP address of the Data Domain server that runs on the system.
9. Provide user credentials that HYCU will use to access the Data Domain server.
10. In the Storage Unit Name field, enter the name of the storage unit that will be used for storing data.
11. From the Authentication Mode drop-down menu, select the required level of authentication:
 - **None**
 - **Two-way password** (*the default mode*)

12. *Only if you selected the two-way password authentication mode.* From the Encryption Strength drop-down menu, select the required level of encryption that will be used when transferring data to the target:
 - **None** (the default value)
 - **Medium**
 - **High**

 **Important** Encryption settings are negotiated between the client (HYCU) and the Data Domain server. The highest configured encryption setting is used. For details, see Dell documentation.

13. Use the **Target encryption** switch if you want the data stored on this target to be encrypted.

 **Note** To be able to import the encrypted target for restoring virtual machines, applications, file shares, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see “[Exporting an encryption key](#)” on page 373.

Keep in mind that if you enable the Target encryption option, HYCU will by default compress backup data before storing it on this target.

14. *Only if charges for reading data from the target may apply.* Use the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges.

With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.

 **Note** If you plan to archive file share data, it is recommended that you keep this option enabled because archiving of file share data is by default performed from the target.

15. Click **Save**.

The target is added to the list of targets. For details on managing targets, see “[Managing targets](#)” on page 325.

Defining your backup strategy

HYCU enables you to schedule automatic backups to achieve the optimum level of data protection based on your recovery point and time objectives, and

backup retention requirements. Backups can be scheduled to start each time the specific number of minutes, hours, days, weeks, or months has passed.

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

- Recovery Point Objective (RPO)

RPO is the maximum period of time for which data loss is considered acceptable (in months, weeks, days, hours, or minutes). For example, setting the RPO to 24 hours means that your business can tolerate losing only data from the last 24 hours.

- Recovery Time Objective (RTO)

RTO is the maximum amount of time (in months, weeks, days, hours, or minutes) that can be spent on restoring data after a disaster occurs.

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

- Taking advantage of predefined policies

You can use any of the predefined policies (Gold, Silver, or Bronze) to simplify the data protection implementation. For details, see “[Taking advantage of predefined policies](#)” below.

- Creating a custom policy

If none of the predefined policies meets the needs of your environment, you can create a new policy and tailor it to your needs. For details, see “[Creating a custom policy](#)” on the next page.

After you decide for a policy approach, consider the following:

- If one of the predefined or custom policies satisfies all data protection goals of your environment, you can set such a policy as default. For details, see “[Setting a default policy](#)” on page 126.
- You can set up the automatic assignment of policies to virtual machines. For details, see “[Setting up automatic policy assignment](#)” on page 124.

Taking advantage of predefined policies

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

HYCU comes with the following predefined policies:

Type of predefined policy	Description
Gold	Data is backed up every 4 hours and restored within 4 hours.
Silver	Data is backed up every 12 hours and restored within 12 hours.
Bronze	Data is backed up every 24 hours and restored within 24 hours.

If you want to exclude entities from being backed up, you can use the Exclude policy.

Creating a custom policy

If the needs of your environment are not covered with any of the predefined policies, you can create a new policy and tailor it to your needs. While tailoring a policy to your needs and setting the preferred RPO, RTO, and targets, you can also enable one or more policy options for optimal policy implementation.

These policy options are the following:

Policy option	Description
Copy	<i>Available only if Target is selected as the backup target type.</i> Allows you to create a copy of backup data.
Archiving	Allows you to preserve your data for future reference.
Fast restore	<i>Not available for vSphere virtual machines residing on VMFS or NFS datastores, or if Snapshot is selected as the backup target type.</i> Allows you to restore virtual machine, application, and volume group data to the original storage container in a fast way by keeping local snapshots for the specified retention time. With this option enabled, HYCU will keep more than one snapshot on the original location, depending on your retention settings. This will allow you to restore virtual machine, application, and volume group data in a fast way, reducing downtime.

Policy option	Description
Backup from replica	<p><i>Available only for Nutanix clusters.</i> Allows you to back up your virtual machines and volume groups from their replicas in remote office/branch office (ROBO) environments.</p> <p>① Important Make sure that the schedule interval you set for the Nutanix protection domains that include the virtual machines and volume groups you want to protect is less than or equal to the RPO set in the HYCU policy.</p>
Auto-assignment	<p>Keep in mind that the replication retention for the respective snapshot on the Nutanix cluster is automatically adjusted to the RPO set in the HYCU policy. This allows HYCU to use the Changed Block Tracking (CBT) feature to get a list of changed data since the last snapshot and perform an incremental backup.</p> <p>For details on protecting virtual machines and volume groups through the Nutanix Prism web console, see Nutanix documentation.</p> <p>Allows you to set up the automatic assignment of policies to virtual machines. You do this by first assigning tags to virtual machines in Nutanix Prism, VMware vSphere, AWS GovCloud (US), or Azure Government, and then specifying the corresponding keys and values in HYCU policies.</p>

Creating a policy

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

Prerequisites

- If you plan to specify time windows for backup and backup copy jobs, make sure you have created them. By specifying time windows, you define time frames when your backup and backup copy jobs are allowed to start. For details on time windows, see “[Creating a time window](#)” on page 116.

- If you plan to enable the Archiving policy option, make sure you have created a data archive. For details on how to do this, see “[Creating a data archive](#)” on page 121.
- If you plan to back up data from replicas in ROBO environments:
 - A protection domain that includes the virtual machines and volume groups that you want to protect must be created and the specified schedule interval must be less than or equal to the RPO set in the HYCU policy. For details on protecting virtual machines and volume groups through the Nutanix Prism web console, see Nutanix documentation.
 - Both the central site Nutanix cluster and the branch office site cluster must be added to HYCU. For details, see “[Adding a Nutanix cluster](#)” on page 54.
- If you plan to enable the Auto-assignment policy option, make sure you are familiar with the information in “[Setting up automatic policy assignment](#)” on page 124.

Limitations

- The Snapshot backup target type is not available for vSphere virtual machines residing on VMFS or NFS datastores.
- The Copy and Fast restore options are not available if you select Snapshot as the backup target type.
- The Backup from replica option is not available for vSphere virtual machines and applications, as well as for the HYCU backup controller.
- The Fast restore option is not available for vSphere virtual machines residing on VMFS or NFS datastores.
- *Only if you plan to enable the Backup from replica option.* Depending on whether Nutanix Disaster Recovery (Nutanix DR) is enabled in Prism Central, the following limitations apply:
 - *Nutanix DR is not enabled:* Backing up data from replicas is supported for virtual machines and volume groups.
 - *Nutanix DR is enabled:*
 - Backing up data from replicas is supported only for virtual machines.
 - The Snapshot backup target type is not available.
 - The Fast restore option is not available.

Considerations

- *Only if you plan to select Snapshot as the backup target type.* When setting the RPO and the retention period, keep in mind that the number of snapshots that will be created by HYCU must not exceed source maximums and snapshot limitations.
- *Only if you plan to protect vSphere virtual machines and enable the Fast restore option.* Keeping snapshots on the source is possible only if all virtual machine disks are located on vVols or vSAN datastores. If any of the disks are located on a VMFS or NFS datastore, such a policy cannot be assigned to the virtual machine.
- *Only if you plan to enable the Backup from replica option.* If Nutanix DR is enabled in Prism Central, make sure that you configure the protection policy in Prism Central in such a way that at least two snapshots are always available and snapshot retention is longer than the RPO set in the HYCU policy. For instructions, see Nutanix documentation.

Accessing the Policies panel

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

Procedure

1. In the Policies panel, click  **New**. The New Policy dialog box opens.
2. Enter a name and, optionally, a description of your policy.
3. Add any of the following policy options to the list of the enabled options by clicking it:
 - **Backup** (*mandatory*)
 - **Copy**
 - **Archiving**
 - **Fast restore**
 - **Backup from replica**
 - **Auto-assignment**
4. In the Backup section, do the following:
 - a. In the Backup every field, set the RPO (in months, weeks, days, hours, or minutes).
 - b. *Only if Target is selected as the backup target type.* In the Recover within field, set the RTO (in months, weeks, days, hours, or minutes).

- c. In the Retention field, set a retention period (in months, weeks, days, or hours) for the data. The retention period defines when a restore point will be expired. For details on data retention, see “[Managing data retention](#)” on page 396.

 **Note** Only if you use Object Lock on Amazon S3 or Nutanix Objects targets. It is recommended that the retention period is approximately the same as the object retention period specified on the cloud target.

- d. Under Backup target type, select the location for storing the protected data:

- **Snapshot**
- **Target**

- e. Only if Target is selected as the backup target type. Under start new backup chain, select when you want a new backup chain to be started:

- **Backup threshold**

A new backup chain is started when the percentage of data changes since the last full backup exceeds the value you specify for this option. The default value is 25.

- **Backup chain length**

A new backup chain is started when the number of the full and subsequent incremental backups in a backup chain exceeds the value you specify for this option. The default value is 7.

For details about the impact of the backup chain length on the backup time window, see “[Creating a backup window](#)” on page 117.

 **Note** If you select both options, the new backup chain is started when either of the specified values has been exceeded.

- f. Only if Target is selected as the backup target type. From the Targets drop-down menu, select one or more targets that you want to use for storing protected data.

If you want your target to be selected automatically, make sure the **Automatically selected** option is selected. In this case, the HYCU advanced scheduler automatically selects only the targets that can guarantee compliance with the RPO and RTO policy settings. Targets that have their estimated backup time lower than the RPO and estimated recovery time lower than the RTO are added to the pool of targets. Based on each entity size, as well as target backup and restore throughput and

queue, the HYCU advanced scheduler calculates the backup and recovery end time and selects the target where the backup will complete the fastest.

Note The target for incremental backups can be any target in the selected pool of targets. To have a single target for all backups in a backup chain, make sure to select a single target per policy.

- g. *Only if you want to specify a backup window.* Enable the **Use backup window** switch, and then, from the Backup Window drop-down menu, select a backup window for backup jobs. If no backup window is available and you want to create one, see “[Creating a backup window](#)” on [page 117](#).

By clicking **Manage**, you are automatically directed to the Time Windows dialog box from where you can manage backup windows.

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

Enabled option	Procedure
Copy	<p><i>Available only if Target is selected as the backup target type.</i> To create a copy of backup data, in the Copy section, do the following:</p> <ol style="list-style-type: none"> a. Set a retention period (in months, weeks, or days) for the copy of backup data. b. From the Targets drop-down menu, select one or more targets that you want to use for storing the copy of backup data. <p>If you want your target to be selected automatically, make sure the Automatically selected option is selected. The copy target will be different from the target for data safety reasons.</p> <p>Note When there are several targets available for storing the copy of backup data and multiple copies of backup data are being created in parallel, HYCU distributes these copies accordingly among targets based on the estimated size of queued and running backups on them.</p>

Enabled option	Procedure
	<p>c. <i>Only if you want to specify a copy window.</i> Enable the Use copy window switch, and then, from the Copy window drop-down menu, select a copy window for backup copy jobs. If no copy window is available and you want to create one, see “Creating a copy window” on page 119. By clicking Manage, you are automatically directed to the Time Windows dialog box from where you can manage copy windows.</p>
Archiving	<p>To archive data, in the Archiving section, from the Data archive drop-down menu, select a data archive. If no data archive is available and you want to create one, see “Creating a data archive” on page 121.</p>
Fast restore	<p><i>Not available for vSphere virtual machines residing on VMFS or NFS datastores, or if Snapshot is selected as the backup target type.</i> To keep more than one snapshot on the source, which allows a fast restore, in the Fast restore section, set a retention period (in months, weeks, days, hours, or minutes) for snapshots. For example, if you set the RPO to two days and the snapshot retention period to four days, you will have two snapshots available on the source.</p> <p>Note The snapshot retention period cannot be shorter than the RPO or longer than the backup retention period.</p>
Backup from replica	<p><i>Available only for Nutanix clusters.</i> To back up data from replicas, in the Backup from replica section, from the Central site cluster drop-down menu, select the cluster on which the replicas of your entities reside.</p>
Auto-assignment	<p>To set up automatic policy assignment, in the Auto-assignment section, enter a key and a value, and then click Add. If required, repeat this step for all the keys and values that you want to add. For details on automatic policy assignment, see “Setting up automatic policy assignment”</p>

Enabled option	Procedure
	<p>on page 124.</p> <p>① Important If the category in Nutanix Prism includes more than one value and you want to add the same key with different values to HYCU, you must repeat this step for each value that you want to add.</p>

6. Click **Save**.

The custom policy is created and added to the list of policies. For details on managing policies, see “[Managing policies](#)” on page 330.

Creating a time window

HYCU enables you to define time frames when your backup and backup copy jobs are allowed to start. If you use a time window, the backup or backup copy jobs are started only within the specified hours, therefore improving effectiveness and avoiding an overloaded environment. For example, you can schedule your backup or backup copy jobs to run on non-production hours to reduce loads during peak hours.

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

① Important When defining a time window, make sure that the RPO specified in the affected policy can be achieved within this time window. If the RPO is shorter than any time frame during which backup or backup copy jobs are not allowed to start, this will result in your entity not being compliant with backup requirements.

Depending on whether you want to create a backup window or a copy window, see one of the following sections:

- “[Creating a backup window](#)” on the next page
- “[Creating a copy window](#)” on page 119

Accessing the Time Window dialog box

To access the Time Window dialog box, in the Policies panel, click  **Time Windows**.

Creating a backup window

Procedure

1. In the Time Windows dialog box, click  **New**. The Select Window dialog box opens.
2. Select **Backup Window**, and then click **Next**.
3. Enter a name for your backup window and, optionally, its description.
4. From the Time zone drop-down menu, specify the time zone for your backup window. You can click one of the displayed time zones (your local time zone or your HYCU backup controller time zone) or select one from the drop-down menu.
5. Click **Full/Incremental** or **Incremental Only** to schedule backups depending on the backup type.

 **Note** Keep in mind the following:

- *Only if Target is selected as the backup target type in your policy.* During the Full/Incremental time frame, backups of any backup type are started, whereas during the Incremental Only time frame, only incremental backups are started. However, if for some reason (for example, due to the Copy policy option being enabled, a snapshot missing, a disk being added to the virtual machine, and so on) an incremental backup cannot be started, a full backup is started instead, also during the Incremental Only time frame.
- *Only if Target is selected as the backup target type in your policy and the backup chain length has been reached:*
 - If the Full/Incremental schedule option is selected, the first backup is a full backup that starts a new backup chain. The future backups will be incremental until the chain length reaches its limit again.
 - If the Incremental Only schedule option is selected, all backups are incremental.
- *Only if Target is selected as the backup target type in your policy and the backup chain length has been reached.* The backup is incremental. All future backups in the Full/Incremental time frame will be full.
- *Only if Snapshot is selected as the backup target type in your policy.* Because the backups performed by HYCU have a minimal effect on your data protection environment, they are started in both the Full/Incremental and the Incremental Only time frame.

6. Select the week days and hours during which you want backups of the selected backup type to start running. To specify time frames for backups of a different backup type, select another backup type, and then repeat this step.

 **Tip** You can click and drag to quickly select a time frame that includes the days and hours you want to add.

The selected time frames are displayed in the Time Frames field. If you want to delete any of the selected time frames, click  next to it.

7. Click **Save**.

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

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

- Specify a backup window when creating a new policy. For details, see “[Creating a policy](#)” on page 110.
- Assign a backup window to the existing policy. To do so, select the policy, click  **Edit**, and then make the required modifications.

Example

You have created the windows_bronze time window and specified the time frames for backups of any type to start on Saturday and Sunday and for incremental only backups to start 6 PM to 6 AM on weekdays.

Backup Window > New

Name

Description - *Optional*

Time Zone

Your local Timezone is  Europe/Ljubljana (UTC+01:00) Controller Timezone is  Europe/Ljubljana (UTC+01:00)

Full/Incremental Incremental Only

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

MON	00:00 - 06:00					18:00 - 24:00					00:00 - 06:00					18:00 - 24:00								
TUE	00:00 - 06:00					18:00 - 24:00					00:00 - 06:00					18:00 - 24:00								
WEN	00:00 - 06:00					18:00 - 24:00					00:00 - 06:00					18:00 - 24:00								
THU	00:00 - 06:00					18:00 - 24:00					00:00 - 06:00					18:00 - 24:00								
FRI	00:00 - 06:00					18:00 - 24:00					00:00 - 06:00					18:00 - 24:00								
SAT	00:00 - 24:00																							
SUN	00:00 - 24:00																							

Time Frames

In this case, the backup jobs will be started every 24 hours (full backups will be started only during the weekends) at any point of time within the specified backup windows.

Creating a copy window

Procedure

1. In the Time Windows dialog box, click  **New**. The Select Window dialog box opens.
2. Select **Copy Window**, and then click **Next**.
3. Enter a name for your copy window and, optionally, its description.
4. From the Time zone drop-down menu, specify the time zone for your copy window. You can click one of the displayed time zones (your local time zone

or your HYCU backup controller time zone) or select one from the drop-down menu.

5. Select the weekdays and hours during which you want backup copy jobs to start running.

 **Tip** You can click and drag to quickly select a time frame that includes the days and hours you want to add.

The selected time frames are displayed in the Time Frames field. If you want to delete any of the selected time frames, pause on the time frame, and then click  that appears next to it.

6. Click **Save**.

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

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

- Specify a copy window when creating a new policy. For details, see “[Creating a policy](#)” on page 110.
- Assign a copy window to the existing policy. To do so, select the policy, click  **Edit**, and then make the required modifications.

Example

You have created the `copy_window_bronze` time window and specified the time frames that allow backup copy jobs to be started from Monday to Friday from 6 PM to 6 AM, and from Saturday to Sunday all day long.

Copy Window > New

Name
copy_window_bronze

Description - Optional
copy_window_bronze_desc

Time Zone
Europe/Ljubljana (UTC+01:00)

Your local Timezone is Europe/Ljubljana (UTC+01:00) Controller Timezone is Europe/Ljubljana (UTC+01:00)

Time Frames

Clear All

MON 00:00 - 06:00 X MON 18:00 - 24:00 X TUE 00:00 - 06:00 X TUE 18:00 - 24:00 X WEN 00:00 - 06:00 X WEN 18:00 - 24:00 X
 THU 00:00 - 06:00 X THU 18:00 - 24:00 X FRI 00:00 - 06:00 X FRI 18:00 - 24:00 X SAT 00:00 - 24:00 X SUN 00:00 - 24:00 X

Close **Back** **Save**

In this case, the backup copy jobs will be started every 24 hours at any point of time within the specified time frames.

Creating a data archive

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

Prerequisites

- The archive target is reserved only for data archives (no backup data is stored on the archive target).

- *For archiving data to the Azure archive storage tier:* Data archives are stored in Azure with the Blob Storage or General Purpose v2 (GPv2) accounts.

Limitations

- *For archiving data to the Azure archive storage tier:* General Purpose v1 (GPv1) accounts do not support moving data archives to the archive storage tier.
- *For archiving data to the Azure archive storage tier and the Google Cloud archive storage class:* Data archives created with any of the previous versions of HYCU are not moved to the archive storage tier.

Consideration

Only if you selected Snapshot as the backup target type in your policy. The configuration settings that HYCU uses for archiving are the ones that the virtual machine has at the time when archiving starts.

Accessing the Policies panel

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

Procedure

1. In the Policies panel, click  **Archiving**.
2. Click  **New**.
3. Enter a name for your data archive and, optionally, its description.
4. Depending on whether you want to create a daily, weekly, monthly, and/or yearly archive of data, add any of the preferred archiving options to the list of the enabled options by clicking it:
 - **Daily**
 - **Weekly**
 - **Monthly**
 - **Yearly**
5. Specify the hour and the minute when the archive job should begin running.
6. From the Time zone drop-down menu, select the appropriate time zone for the archive job.

 **Note** All scheduled archive jobs are by default started based on the HYCU backup controller time zone and are not affected by the time windows specified for the same policy.

7. Depending on the selected archiving options, specify at what intervals you want your data to be archived:

Archiving option	Instructions
Daily	<ul style="list-style-type: none"> a. In the Recur every field, specify whether you want the data to be archived every day or every few days. b. Use the Apply only on weekdays switch if you want the data to be archived only on weekdays.
Weekly	<ul style="list-style-type: none"> a. In the Recur every field, specify whether you want the data to be archived every week or every few weeks. b. Select one or more days of the week on which you want the data to be archived. <p>Note If you select more than one day, archive compliance is calculated by taking into account data archives of all the selected days, not only the latest data archive.</p>
Monthly	<ul style="list-style-type: none"> a. In the Recur every field, specify whether you want the data to be archived every month or every few months. b. Select whether you want the data to be archived on the same day of the month (for example, on the fifth day of the month), or on a specific day of the month (for example, on the second Friday of the month).
Yearly	<ul style="list-style-type: none"> a. In the Recur every field, specify whether you want the data to be archived every year or every few years. b. Select whether you want the data to be archived on the same day of the preferred month (for example, on the fifth day of January), or on a specific day of the preferred month (for example, on the second Friday of April).

8. In the Retention field, set the retention period to be used.

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

9. From the Target drop-down menu, select one or more archive targets.
10. Click **Save**.

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

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

- Specify a data archive when creating a new policy. For details, see “[Creating a policy](#)” on page 110.
- Assign a data archive to the existing policy. To do so, select the policy, click  **Edit**, and then make the required modifications.
- Archive data manually. For details, see “[Archiving data manually](#)” on page 339.

Setting up automatic policy assignment

By setting up automatic policy assignment, you ensure that policies are automatically assigned to all virtual machines to which tags are assigned. This is especially useful in complex data protection environments where the data protection approach often requires the use of various policies.

After you assign tags to virtual machines and specify the matching keys and values, and the comparison of these values shows that the specified values match, the corresponding policies are automatically assigned to the virtual machines during the next virtual machine synchronization.

HYCU performs the automatic synchronization of virtual machines every five minutes. However, you can at any time update the list of virtual machines also manually by clicking  **Refresh** in the Virtual Machines panel.

 **Note** HYCU uses the term tags to refer also to categories and custom attributes that are assigned to virtual machines on Nutanix clusters and in vSphere environments.

Considerations

- If you want a predefined policy to be automatically assigned to a virtual machine, when specifying the tag and the matching key and value, you can use the name of the policy (Gold, Silver, Bronze, or Exclude). Keep in mind that if you use the Exclude value, the virtual machine will be excluded from the backup.
- Assigning policies automatically does not affect virtual machines that already have a policy assigned.
- If the default policy is set, it is never assigned to newly discovered virtual machines that have tags applied, but only to the ones for which no automatic assignment of policies is set up. For details on setting the default policy, see “[Setting a default policy](#)” on the next page.
- If the comparison of tags and keys and values returns multiple match results, the policy with the lowest RPO is assigned to the virtual machine.
- *For Nutanix ESXi clusters and vSphere environments:* After you restore a virtual machine for which you have set up automatic policy assignment, the tag value is kept on the restored virtual machine only if the original tag still exists in VMware vSphere.

Procedure

1. Sign in to the Nutanix Prism web console, the vSphere (Web) Client, the AWS GovCloud (US) console, or the Azure Government portal.
2. Assign tags to virtual machines for which you want to set up automatic assignment. For instructions, see the relevant documentation.
3. Sign in to the HYCU web user interface.
4. Specify the matching keys and values in HYCU policies as described in “[Creating a policy](#)” on page 110.

! Important Depending on your data protection environment, the key and the value that you should enter represent the following:

- *For Nutanix AHV clusters:* The name and the value of the category.
- *For Nutanix ESXi clusters or vSphere environments:* The tag name and the category of the tag, or the attribute and the value of the custom attribute.
- *For AWS GovCloud (US) or Azure Government environments:* The name and the value of the tag.

Setting a default policy

You can select one of the predefined or custom policies to be the default policy for your data protection environment. When you set the default policy, depending on your choice, the default policy will be assigned to the following entities (applications, virtual machines, volume groups, and/or file shares):

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

Accessing the Policies panel

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

Procedure

1. In the Policies panel, select the policy that you want to set as the default one, and then click  **Set Default**. The Set Default Policy dialog box opens.
2. Select the entities to which you want the default policy to be assigned:
 - **Virtual Machines**
 - **Applications**
 -  **Important** Setting the default policy for applications is possible only if the default policy is set also for virtual machines.
 - **Volume Groups**
 - **Shares**
3. Depending on whether you want the default policy to be assigned to only newly discovered entities, or both newly discovered entities and already existing entities without an assigned policy, do one of the following:

I want the default policy to be assigned to...	Instructions
Only newly discovered entities.	Click Save .
Both newly discovered entities and all existing entities that do not have an assigned policy.	<ol style="list-style-type: none"> a. Enable the Assign to entities without policy switch. b. Click Save.

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

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

Chapter 4

Protecting virtual machines

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

Depending on your source, you can also protect the following:

Source	Item available for protection
Nutanix cluster	<p>Volume groups (collections of logically related virtual disks) in storage containers</p> <p>! Important If one or more volume groups are attached to a virtual machine at backup time, they are also backed up during the virtual machine backup. You can view such volume groups and their details in the Volume Groups panel together with all existing volume groups residing on Nutanix clusters that have been added to HYCU. For instructions on how to enable data protection for volume groups independent of virtual machine protection, see “Protecting volume groups” on page 274.</p>
vSphere environment	Virtual machine templates (virtual machines that are used as templates to create other virtual machines)

The preparation steps and instructions for protecting virtual machines (including the HYCU backup controller) and servers may differ.

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

- [“Planning virtual machine protection” on the next page](#)
- [“Backing up virtual machines” on page 153](#)
- [“Restoring virtual machines” on page 155](#)
- [“Restoring individual files” on page 202](#)

Planning virtual machine protection

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

- “Preparing your data protection environment” below
- “Preparing for disaster recovery” on page 132
- “Preparing for the restore to a different source” on page 135
- “Server specifics” on page 140
- “SpinUp specifics” on page 142
- “Enabling access to data” on page 146
- “Setting up virtual machine backup configuration options” on page 149

Preparing your data protection environment

Prerequisites

- *For vSphere environments and Nutanix ESXi clusters:* VMware Tools of the latest version is installed on the virtual machines. For detailed information about installing VMware Tools, see VMware documentation.
- *For ROBO environments:* If volume groups are attached to virtual machines that you plan to back up and you want these volume groups to be backed up during the virtual machine backup, make sure they are in the same Nutanix protection domain as the virtual machines.
- *For archiving data to a QStar tape target:* 1 GiB of additional free memory is available on the HYCU backup controller for each concurrent archive job.
- *Only if you plan to validate the virtual machine backup and specify a custom script.*
 - The script must be available on the virtual machine in the accessible folder and must have one of the following extensions:
 - Windows: bat, ps1, cmd
 - Linux: sh
 - *For Linux:* You must have permissions to run the script on the virtual machine with the assigned credentials.

- *For Linux virtual machines that you plan to clone to an AWS GovCloud (US) environment:* If you have SSH password authentication and the SSH root login enabled in the SSH service, and you want to use them on the restored virtual machine, the `cloud-init` service, if present, must be properly configured on the original virtual machine to not override the SSH service configuration on the restored virtual machine.

Limitations

- Only the backup of local fixed disks and Nutanix volume groups is supported. When backing up a virtual machine with remote volumes (for example, iSCSI, disk arrays, mapped network disks), such volumes are not included in the snapshot and are consequently not backed up.
- *For Linux virtual machines:* Restoring files is possible only from file systems that are permanently mounted. Therefore, make sure the required file systems are specified in the `/etc/fstab` file before the backup is performed.
- *For Nutanix clusters:* Protecting the following types of virtual machines is not supported: Nutanix Controller VMs, Prism Central VMs, Nutanix Files file server VMs, and Nutanix Objects nodes. Therefore, such virtual machines are not shown in the Virtual Machines panel. If you want to protect these types of virtual machines, contact your Nutanix Sales representative.
- *For Nutanix ESXi clusters:*
 - Protecting virtual machines that have NVMe controllers added is not supported.
 - If you enabled the Backing up from replica policy option, backing up virtual machines that have disks on different containers is not supported.
- *For AWS GovCloud (US) virtual machines with instance store volumes:* Protecting instance store volumes is not supported.
- *For vSphere environments:* Protecting virtual machines that have vSphere Fault Tolerance enabled is not supported.

Considerations

- In large or medium size data protection environments with virtual machines of larger size (2–4 TiB), keep in mind, that the first backup of such virtual machines takes more time and resources. Consider protecting these virtual machines in such a way that they are not backed up simultaneously. You can assign a policy to a large virtual machine, wait until it gets protected,

and then continue with protecting other virtual machines.

- *For Nutanix clusters, AWS GovCloud (US) environments, and Azure Government environments:* Archiving is performed from a snapshot if the snapshot is available. Otherwise, archiving is performed from the target (if Target is defined as the backup target type in your policy).
- *For vSphere environments:*
 - The number of snapshots that can be created per virtual machine may differ due to snapshot limitations. For details, see VMware documentation.
 - If something unexpected occurs during the backup of a virtual machine template (for example, a network problem), the virtual machine template that is converted to a virtual machine as part of the backup process will remain converted. In this case, make sure to convert the virtual machine back to the virtual machine template. For details on how to do this, see VMware documentation.
 - *Only if you use HotAdd and plan to restore individual files.*
 - When restoring from a snapshot, using HotAdd is supported if all the prerequisites are met. For details about HotAdd prerequisites, see VMware documentation.
 - Using HotAdd is supported only for snapshots that were created automatically by HYCU, not for snapshots that were created by using the Recreate Snapshot option.
- *For Nutanix ESXi clusters, AWS GovCloud (US) environments, and Azure Government environments:* If the snapshot that HYCU used to perform a full backup is missing on the source, the next virtual machine backup will be a full backup.
- *For protection domains configured with NearSync:* Although snapshots in a protection domain are created in a 1–15 minute interval, HYCU uses only the snapshots that are created on an hourly basis for backing up and restoring from snapshots. This applies to the following environments:
 - Nutanix ESXi clusters
 - Nutanix clusters when using the Backup from replica option
- *For Nutanix ESXi clusters:* If a storage container of the Nutanix ESXi cluster is presented as an NFS datastore to the VMware infrastructure, a full backup of a virtual disk on such a storage container performed using a corresponding vSphere source will copy the entire allocated disk, not only the used blocks.

- If you want the virtual machine details section in the Nutanix Prism web console and vSphere (Web) Client to contain the information on which HYCU policy is assigned to a virtual machine, in the `HYCU config.properties` file, set the `hycu.policy.description` configuration setting to true. For details on how to customize the HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on page 523.
- *For ROBO environments:* The number of snapshots in the protection domain may be higher than configured if HYCU uses these snapshots for backing up virtual machines and volume groups.

Recommendations

- *For ROBO environments:* If a volume group is attached to several virtual machines that you plan to back up and you want this volume group to be backed up as well, it is recommended that it is attached only to the virtual machines inside the same Nutanix protection domain.
- *For virtual machines in a ROBO environment:* To ensure that applications on virtual machines are up and running after restoring the virtual machines, it is recommended that application-consistent snapshots are created for them. For details on how to do this, see Nutanix documentation.
- *For Nutanix AHV clusters:* If your virtual machine is protected with the synchronous replication schedule, the virtual machine snapshots created by HYCU are replicated also to the remote sites that HYCU cannot access. This results in HYCU not being able to perform the snapshot cleanup. Therefore, to avoid excessive storage consumption and having to manually delete snapshots, it is recommended to protect your virtual machines with the asynchronous or NearSync replication schedules.

Preparing for disaster recovery

To achieve high reliability and resilience of your data protection environment, you must also protect the HYCU backup controller itself. By doing so, you ensure integrity and safety of the protected data, and avoid data loss in case of a disaster, for example, when your HYCU backup controller is deleted by accident or the cluster node on which it is running stops operating. In addition, if your data protection environment also includes HYCU instances, you must protect these as well.

Make sure to take a note of the configuration parameters of the target on which you plan to store the HYCU backup controller backups. You can also take a

note of the configuration parameters of any target on which you plan to store the backups of virtual machines, applications, file shares, and volume groups if you decide to recover them without recovering the HYCU backup controller. You will need to provide the correct configuration data when importing the target for disaster recovery.

Target type	Required information for importing
NFS	<ul style="list-style-type: none"> • NFS server name or IP address • Shared folder
SMB	<ul style="list-style-type: none"> • Domain (if used) • User name (if used) • Password (if used) • SMB server name or IP address • Shared folder
Nutanix	<ul style="list-style-type: none"> • URL • User name • Password
Nutanix Objects	<ul style="list-style-type: none"> • Service endpoint • Bucket name • Access key ID • Secret access key • Path style access
iSCSI	<ul style="list-style-type: none"> • Target portal • Target name • User (if CHAP authentication is enabled) • Target secret (if CHAP authentication is enabled) • Perform mutual authentication (if CHAP authentication is enabled)
Amazon S3 / S3 Compatible	<ul style="list-style-type: none"> • Service endpoint • Bucket name • Access key ID • Secret access key • Path style access
Azure	<ul style="list-style-type: none"> • Storage account name • Secret access key • Storage container name
Google Cloud	<ul style="list-style-type: none"> • Bucket name • Google Cloud service account
QStar NFS	<ul style="list-style-type: none"> • User name • Password (if used) • Integral volume set name • Web service endpoint

	<ul style="list-style-type: none"> • Shared folder (if used)
QStar SMB	<ul style="list-style-type: none"> • Domain (if used) • User name • Password (if used) • Integral volume set name • Web service endpoint • Shared folder (if used)
Data Domain	<ul style="list-style-type: none"> • Data Domain server host name or IP address • User name • Password • Storage unit name • Authentication mode • Encryption strength

Considerations

- The RPO in the policy that is assigned to the HYCU backup controller should always be lower than any RPO already set for other protected entities in the data protection environment.
- Assigning a policy that has the Backup from replica policy option enabled to the HYCU backup controller is not supported.
- *Only if you use more than one HYCU backup controller for data protection.* Each HYCU backup controller must be protected from within its own web user interface.

Recommendation

To further increase safety, it is recommended that you combine protection of the HYCU backup controller with protection of the source that hosts the HYCU backup controller. You can use, for example, Nutanix protection domains or VMware vSphere Data Protection. For more information, see Nutanix or VMware documentation.

Preparing for the restore to a different source

If you plan to restore your virtual machines to a different source, keep in mind the prerequisites, limitations, considerations, and recommendations described in this section.

 **Note** During the backup, HYCU performs the platform readiness check to ensure that the virtual machine can be successfully restored to a different source. You can view the platform readiness check status in the backup job report.

Prerequisites

- For Linux virtual machines that you plan to restore to a Nutanix ESXi cluster or a vSphere environment: The VMware Paravirtual SCSI driver (`vmw_pvscsi`) must be included in initramfs.

 **Important** Adding the `vmw_pvscsi` driver is required only if the driver is built as a module and not included in the kernel. To check if the `vmw_pvscsi` driver is already a part of the kernel, see the value of the `CONFIG_VMWWARE_PVSCSI` setting in file `/boot/config-`/usr/bin/uname -r``. If the value is "m", the driver must be added as a kernel module. If the value is "y", the driver is already present in the kernel.

To add the driver, on the virtual machine, run the following command as the root user:

```
dracut -f add-drivers "vmw_pvscsi"
```

To check if the driver is present after adding, on the virtual machine, run the following command as the root user:

```
lsinitrd | grep "vmw_pvscsi"
```

- For Linux virtual machines that you plan to restore to a Nutanix AHV cluster: Virtio drivers (`virtio_pci`, `virtio_net`, and `virtio_scsi`) must be included in initramfs.

 **Important** Adding the listed drivers is required only if the drivers are built as a module and not included in the kernel. To check if the listed drivers are already included in the kernel, see the value of these settings in file `/boot/config-`/usr/bin/uname -r``: `CONFIG_VMWWARE_PVSCSI`, `CONFIG_SCSI_VIRTIO`, and `CONFIG_VIRTIO_PCI`. If the values are "m", the drivers must be added as a kernel module. If the value is "y", the drivers are already present in the kernel.

To add the drivers, on the virtual machine, run the following command as the root user:

```
dracut -f --add-drivers "virtio_pci virtio_net virtio_scsi"
```

- *For Linux virtual machines that you plan to clone to an AWS GovCloud (US) environment:* If you have SSH password authentication and the SSH root login enabled in the SSH service, and you want to use them on the restored virtual machine, the `cloud-init` service, if present, must be properly configured on the original virtual machine to not override the SSH service configuration on the restored virtual machine.
- *For Windows virtual machines that you plan to restore to a Nutanix AHV cluster:* The Nutanix VirtIO driver package must be installed.
- *For restoring a virtual machine to an AWS GovCloud (US) or Azure Government environment:*
 - Access to the virtual machines through SSH or remote desktop connection must be enabled and a firewall must be configured to allow a remote desktop or SSH connection using a public network.
 - Appropriate credentials must be assigned to the virtual machines that you plan to restore or to the virtual machines on which the applications that you plan to restore are running. For instructions on how to assign credentials to a virtual machine, see [“Enabling access to data” on page 146](#).
 - *For Linux virtual machines:*
 - DHCP must be enabled on the virtual machines that you want to migrate to cloud.
 - Privileged access to the Linux system as root or by using the `sudo` command without a password is required.
 - The use of persistent network device names based on MAC addresses must be disabled. For details on how to do this, see your Linux distribution documentation. Legacy network naming is used instead.
 - Hyper-V drivers (`hv_vmbus`, `hv_storvsc`, and `hv_netvsc`) must be included in `initramfs`. To add the drivers, on the virtual machine, run the following command as the root user:

```
dracut -f --add-drivers "hv_vmbus hv_storvsc hv_netvsc"
```

- *For Linux virtual machines:* In the `/etc/fstab` system configuration file of the virtual machine, LABEL or UUID must be used instead of device names for file system device identification (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`). If not, platform readiness check will issue a warning in the backup job report.

- *For disaster recovery:* Configure your environment to provide the DR-ready status for the virtual machines. A virtual machine has the DR-ready status if the following is true:
 - A successful platform readiness check is performed during the virtual machine backup.
 - All backups in the current backup chain are stored on one of the cloud targets (Google Cloud, Azure, or Azure Government).

You can check the DR-ready status of a virtual machine in the Virtual Machines panel.

Limitation

For virtual machines that you plan to restore to AWS GovCloud (US): Only virtual machines with no more than 11 disks can be restored.

Considerations

- If during a restore of the selected virtual machine you receive a warning message indicating that there is a guest operating system mismatch detected (between the guest operating system that is running on the virtual machine and the one specified during the configuration of the virtual machine) or a memory size mismatch detected while creating a new virtual machine, make sure to modify the virtual machine configuration after the restore by specifying the appropriate guest operating system or memory. By doing so, you make sure that the restored virtual machine has the same configuration as it had before the restore. For details on how to do this, see Nutanix or VMware documentation.
- *For virtual machines with attached volume groups:* You must reattach the volume groups to the virtual machine after the restore. For details on how to do this, see Nutanix and guest operating system documentation.
- Depending on your virtual machine original environment and target environment, some additional steps may be required after the restore. For details, see “[After restoring a virtual machine to a different source](#)” on [page 530](#).

Recommendations

- It is recommended that all virtual machine disks are online. If the disks are offline, a warning is issued in the platform readiness check job report.
- *For restoring a virtual machine to a Nutanix AHV cluster:* Follow these recommendations before backing up your virtual machine to ensure that the

virtual machine will start after the restore (otherwise, you will need to perform additional manual steps as described in [“After restoring a virtual machine to a Nutanix AHV cluster” on page 531](#)):

- *For Windows virtual machines:* The Nutanix VirtIO package is installed on the virtual machine.
- *For Linux virtual machines on Nutanix ESXi clusters:* Nutanix Guest Tools (NGT) is installed on your virtual machine.
- *For Linux virtual machines in vSphere, AWS GovCloud (US), or Azure Government environments:* The VirtIO drivers are available as a kernel module which is added to initramfs.

How to determine the availability of the VirtIO drivers and add them if necessary

To check if the VirtIO drivers are available in the installed kernel, as the root user, run the following command:

```
grep -i virtio /boot/config-`uname -r`
```

The following output confirms that the VirtIO drivers are available:

```
CONFIG_VIRTIO_BLK=m
CONFIG_SCSI_VIRTIO=m
CONFIG_VIRTIO_NET=m
CONFIG_VIRTIO_CONSOLE=m
CONFIG_HW_RANDOM_VIRTIO=m
CONFIG_DRM_VIRTIO_GPU=m
CONFIG_VIRTIO=m
# Virtio drivers
CONFIG_VIRTIO_PCI=m
CONFIG_VIRTIO_PCI_LEGACY=y
CONFIG_VIRTIO_BALLOON=m
CONFIG_VIRTIO_INPUT=m
# CONFIG_VIRTIO_MMIO is not set
```

To check if the VirtIO drivers are added to initramfs, as the root user, run the following commands:

```
cp /boot/initramfs-`uname -r`.img /tmp/initramfs-`uname -r`.img.gz
```

```
zcat /tmp/initramfs-`uname -r`.img | cpio -it | grep virtio
```

An output similar to the following one appears if the VirtIO drivers are added to initramfs:

```
97084 blocks
```

If the output is blank, the VirtIO drivers are not added to initramfs. To add the VirtIO drivers to initramfs, as the root user, run the following command:

```
dracut --add-drivers "virtio_pci virtio_blk virtio_scsi
virtio_net" -f -v
```

To check if the VirtIO drivers are added to initramfs, as the root user, run the following commands:

```
cp /boot/initramfs-`uname -r`.img /tmp/initramfs-`uname -
r`.img.gz
```

```
zcat /tmp/initramfs-`uname -r`.img | cpio -it | grep virtio
```

An output similar to the following one should appear:

```
usr/lib/modules/`uname -r`/kernel/drivers/scsi/virtio_scsi.ko
usr/lib/modules/`uname -r`/.x86_
64/kernel/drivers/block/virtio_blk.ko
usr/lib/modules/`uname -r`/kernel/drivers/char/virtio_
console.ko
usr/lib/modules/`uname -r`/kernel/drivers/net/virtio_net.ko
usr/lib/modules/`uname -r`/kernel/drivers/virtio
usr/lib/modules/`uname -r`/kernel/drivers/virtio/virtio.ko
usr/lib/modules/`uname -r`/kernel/drivers/virtio/virtio_pci.ko
usr/lib/modules/`uname -r`/kernel/drivers/virtio/virtio_
ring.ko
97084 blocks
```

For details, see Nutanix documentation.

Server specifics

The instructions for protecting virtual machine data apply also to servers except where specifically stated otherwise.

Prerequisites

- Access to the file system data is enabled. For instructions, see “[Enabling access to data](#)” on page 146.
- Sufficient disk space—estimated at up to 1.8 percent of the space of all volumes that you plan to back up—is available for the index created by HYCU for data protection purposes at the following location:
 - Linux: `/var/opt/hycu/hycuraw`
 - Windows: `%programdata%\HYCU\hycuraw`
- *For Windows servers:*
 - The VSS service is enabled and running, and the VSS writer status is stable.
 - WinRM is enabled and configured by using the `winrm quickconfig` command.
 - *For cloning a Windows server to a Nutanix AHV cluster:* Make sure the Nutanix VirtIO package is installed on the server before you back it up. For detailed information about installing Nutanix VirtIO, see Nutanix documentation.
- *For Linux servers:*
 - Access to the server through SSH is enabled.
 - *Only if you plan to back up data by using LVM snapshots (the recommended approach):* Sufficient space in the volume group is available for LVM snapshots. It is recommended that at least 10 percent of free space is available in each volume. However, the percent should be higher if a large number of writes to volumes is expected during the backup. For more information, see LVM documentation.
 - Privileged access to the Linux system as root or by using the `sudo` command without a password is required.
 - The dm-snapshot kernel module must be included in initramfs. To add the module, on the server, run the following command as the root user:


```
dracut -f --add-drivers "dm-snapshot"
```
 - *For cloning a Linux server:* The following drivers must be added to the guest OS kernel:
 - *For cloning to a Nutanix AHV cluster:* Nutanix VirtIO drivers (`virtio_pci`, `virtio_blk`, `virtio_scsi`, `virtio_net`)
 To add the drivers, run the following command as the root user:

```
dracut -f --add-drivers "virtio_pci virtio_blk virtio_
scsi virtio_net"
```

- For cloning to a Nutanix ESXi cluster or a vSphere environment: VMware driver `vmw_pvscsi`

To add the driver, run the following command as the root user:

```
dracut -f --add-drivers "vmw_pvscsi"
```

Limitations

- Protecting servers that use Virtual Data Optimizer (VDO) is not supported.
- For Linux servers that use UEFI firmware:
 - Only the default boot loaders of the supported operating systems are supported. For a list of supported operating systems, see the *HYCU Compatibility Matrix*.
 - The EFI system partition must be mounted on the default location used by the operating system (`/boot/efi`).

Consideration

For Linux servers: By default, HYCU uses LVM snapshots for data protection. However, if you cannot provide the required space for LVM snapshot storage in each volume, you can configure HYCU to use device mapper (DM) snapshots as an alternative. For details, see “[Enabling DM snapshots](#)” on page 152.

SpinUp specifics

If you plan to use the SpinUp functionality to migrate your virtual machines across on-premises and cloud (AWS, Google Cloud, global Azure, or Azure Government) environments, make sure that the following prerequisites are met:

Prerequisites

- Configure your environment to provide a successful platform readiness check during the virtual machine backup:
 - Access to the virtual machines through SSH or remote desktop connection is enabled and a firewall is configured to allow a remote desktop or SSH connection using a public network.

- Appropriate credentials are assigned to the virtual machines that you plan to migrate or to the virtual machines on which the applications that you plan to migrate are running. For instructions on how to assign credentials to a virtual machine, see “[Enabling access to data](#)” on [page 146](#).
- *For migration of Linux virtual machines:*
 - DHCP is enabled on the virtual machines that you want to migrate to cloud.
 - Privileged access to the Linux system as root or by using the `sudo` command without a password is required.
 - The use of persistent network device names based on MAC addresses is disabled. For details on how to do this, see your Linux distribution documentation. Legacy network device naming is used instead.
 - In the `/etc/fstab` system configuration file of the virtual machine, `LABEL` or `UUID` (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`) must be used instead of device names for file system device identification.
 - The following drivers must be included in initramfs:
 - *For migration to AWS:* `ixgbevf`, `ena`, `nvme`, `nvme-core`, `xen_netfront`, and `xen_blkfront`
To add the drivers, on the virtual machine, run the following command as the root user:

```
dracut -f --add-drivers "ixgbevf ena nvme nvme-core
xen_netfront xen_blkfront"
```
 - *For migration to Google Cloud:* `virtio_pci`, `virtio_net`, and `virtio_scsi`
To add the drivers, on the virtual machine, run the following command as the root user:

```
dracut -f --add-drivers "virtio_pci virtio_net virtio_scsi"
```

ⓘ Important Adding the `virtio_pci` driver is required only if it is built as a module and not included in the kernel.

To add the drivers, on the virtual machine, run the following command as the root user:

```
dracut -f --add-drivers "hv_vmbus hv_storvsc hv_
netvsc"
```

- For migration from AWS or Azure to a Nutanix AHV cluster: `virtio_net.ko`, `virtio_scsi.ko`, and `virtio_pci.ko`

To add the drivers, on the virtual machine, run the following command as the root user:

```
dracut -f --add-drivers "virtio_net.ko virtio_scsi.ko
virtio_pci.ko"
```

- For migration of Windows virtual machines:

- For migration to Google Cloud:
 - The Nutanix VirtIO package is installed on the virtual machines that you plan to migrate.
 - DHCP is enabled on the virtual machines that you want to migrate to cloud.
- For migration to Azure: DHCP is enabled on the virtual machines that you want to migrate to cloud.
- For migration from AWS or Azure to a Nutanix AHV cluster: The Nutanix VirtIO package is installed on the virtual machines that you plan to migrate.

You can view the platform readiness check status in the backup job report.

- For disaster recovery to cloud: Configure your environment to provide the DR-ready status for the virtual machines. A virtual machine has the DR-ready status if the following is true:
 - All backups in the current backup chain are stored on the respective cloud target.
 - A successful platform readiness check is performed during the virtual machine backup.

You can check the DR-ready status of a virtual machine in the Virtual Machines panel.

Limitations

- Migrating protected data across the on-premises and cloud environments is not supported for multi-boot systems.
- *For migration of virtual machines from cloud:* You can migrate virtual machines that use UEFI firmware only to a Nutanix AHV cluster or a vSphere environment. Migrating such virtual machines to a Nutanix ESXi cluster is not supported.

Considerations

- If you are migrating data from a Nutanix cluster, the data is migrated from the snapshot if the snapshot is available. Otherwise, the data is migrated from the target (if Target is defined as the backup target type in your policy).

⚠️ Important If a restore point contains only a Snapshot tier, you cannot use it for migrating data.

- *For Windows virtual machines:* If the virtual machine has more than one disk, additional disks are put offline during the migration by default. You can put the disks back online manually after the migration or you can change the default setting before performing a backup by running the following command in PowerShell:

```
Set-StorageSetting -NewDiskPolicy OnlineAll
```

Recommendations

- It is recommended that all virtual machine disks are online. If the disks are offline, a warning is issued in the platform readiness check job report.
- *For Windows virtual machines:* It is recommended to enable EMS console redirection for troubleshooting purposes. Having it enabled allows you to gather more information in the case a virtual machine does not boot after being migrated to cloud.
- *For Linux virtual machines:* It is recommended to enable serial console redirection for troubleshooting purposes. Having it enabled allows you to configure the virtual machine network in the case this is required after migration to cloud. A virtual machine with serial console redirection enabled has the successful platform readiness check status even if the network is not working.

- *For migration of Linux servers that use UEFI firmware to cloud:* If the virtual machine does not boot after the migration, reboot the machine.

Enabling access to data

When the recovery goals of your environment require backing up data inside the file systems of your virtual machine or server, you must enable HYCU to access it.

Enabling access to data is a prerequisite in the following data protection scenarios:

- You plan to protect servers.
- You plan to validate virtual machine backups.
- You plan to restore individual files to the virtual machine.
- You plan to protect applications.
- You plan to protect volume groups that are attached to a virtual machine by using iSCSI as part of protecting the virtual machine.
- You plan to use pre- and post-scripts.
- You plan to use the SpinUp functionality to migrate your virtual machines and applications to cloud.

Prerequisites

- A firewall must be configured to allow inbound network traffic through the required TCP port.
- *Only if the WinRM protocol over HTTPS will be used.* HYCU must be configured to use HTTPS for WinRM connections to virtual machines. For instructions, see “[Enabling HTTPS for WinRM connections](#)” on page 446.

Limitation

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

Considerations

- *For Windows virtual machines:* When specifying a user name, make sure to use one of the following formats:
 - If the virtual machine is added to an Active Directory domain:
`<Domain>\<Username> or <Username>@<Domain>`

- If the virtual machine is not added to an Active Directory domain: <Username>, .\<Username>, or <Hostname>\<Username> (in this case, <Hostname> is the value of the COMPUTERNAME variable).
- *For virtual machines that you plan to back up from their replicas in ROBO environments:* Make sure that the most recent replica reflects the state of the virtual machine.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, select the virtual machine to which you want to enable access.
2. Click  **Credentials**. The Credential Groups dialog box opens.
3. Click  **New**.
4. Enter a name for the credential group.
5. From the Protocol drop-down menu, select one of the following protocol options:

Protocol options	Instructions
Automatic	<p>Select this option if you want HYCU to automatically select a protocol for accessing the virtual machine: the SSH protocol (TCP port 22) or the WinRM protocol (HTTPS transport and TCP port 5986, or HTTP transport and TCP port 5985), and then enter the user name and password of a user account that has required permissions to access the virtual machine.</p> <p> Note <i>For Linux virtual machines:</i> Password authentication is used by default. If you want to use public key authentication, select the SSH protocol option and make the required modifications.</p>
SSH	<p>Select this option if you want to use the SSH protocol, and then do the following:</p> <ol style="list-style-type: none"> a. In the Port field, enter the SSH server port number.

Protocol options	Instructions
	<p>b. From the Authentication type drop-down menu, select the type of authentication you want to be used and provide the required information:</p> <ul style="list-style-type: none"> • Password authentication Enter the user name and password of a user account that has required permissions to access the virtual machine. • Public key authentication <ul style="list-style-type: none"> • In the Username field, enter the user name of a user account that has required permissions to access the virtual machine. • Choose a private key. <p>Note Only if you are signed in to HYCU as a self-service group administrator. If you use Conjur for managing your HYCU secrets, you can enable the Retrieve values from secrets manager switch if you want to provide the secret instead of browsing for the file. For details on managing secrets, see “Managing secrets” on page 404.</p> • Only if the private key is encrypted. Enter the private key passphrase.
WinRM	<p>Select this option if you want to use the WinRM protocol, and then do the following:</p> <ol style="list-style-type: none"> a. From the Transport drop-down menu, select the type of transport you want to be used. b. In the Port field, enter the WinRM server port number. c. Enter the user name and password of a user account that has required permissions to access the virtual machine.

6. Click **Save**.

7. Click **Assign**.

The name of the assigned credential group appears in the Credential group column of the Virtual Machines panel. HYCU performs virtual machine and application discovery after you assign the credentials to the virtual machines and the Discovery status in the Virtual Machines and Applications panels is updated accordingly.

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

To unassign a credential group from a virtual machine, in the Virtual Machines panel, select the virtual machine, click  **Credentials**, and then click **Unassign**.

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

Setting up virtual machine backup configuration options

For each virtual machine, you can set up configuration options to better adjust the scope and flow of a specific virtual machine backup to the needs of your data protection environment.

You can set the backup configuration options on the selected virtual machine for the following purposes:

I want to...	Instructions
Specify the pre/post-backup and pre/post-snapshot scripts.	“Specifying pre/post-backup and pre/post-snapshot scripts” on the next page
Specify any disks or volume groups to exclude or include when backing up a virtual machine.	“Excluding or including disks in the backup” on page 151
<i>Applicable only for Linux servers.</i> Configure HYCU to use DM snapshots instead of LVM snapshots for backing up data.	“Enabling DM snapshots” on page 152

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Specifying pre/post-backup and pre/post-snapshot scripts

You can use the pre/post-backup and pre/post-snapshot scripts to perform necessary actions before the backup is performed or the snapshot is created (for example, to suspend application I/O), and after the backup is performed or the snapshot is created (for example, to resume application I/O). For details on how to specify the scripts, follow the procedure described in this section. For details on exit codes and exported environment variables, see “[Using the pre and post scripts](#)” on page 461.

Prerequisites

- Access to the virtual machine file system is enabled. For instructions, see “[Enabling access to data](#)” on page 146.
- A script is available in the accessible folder and has one of the following extensions:
 - Windows: bat, ps1, cmd
 - Linux: sh
- *For Linux:* You have permissions to run a script on the virtual machine with the assigned credentials.

Procedure

1. In the Virtual Machines panel, select the virtual machine on which you want to specify pre/post scripts, and then select  **Configuration**. The Configuration dialog box opens.
2. In the Pre/post scripts tab, use the fields of your choice to specify which pre/post-snapshot and pre/post-backup scripts should be run. Enter the script path names to one or more fields:
 - **Run pre-backup script**
 - **Run pre-snapshot script**
 - **Run post-snapshot script**
 - **Run post-backup script**

 **Note** In the script path name field, a sample path name is displayed. Make sure to enter the valid script path name.

3. Click **Save**.

Excluding or including disks in the backup

By default, all disks and volume groups that are attached to a virtual machine are backed up during the virtual machine backup. However, if you want specific disks to be excluded from or included in the backup, HYCU enables you to select these disks before the virtual machine backup is performed:

- By excluding disks, you make sure that only the selected disks are not backed up.
- By including disks, you make sure that only the selected disks are backed up. In this case, any temporary disks are automatically excluded from the backup.

Prerequisite

You are an owner of the virtual machine whose disks you want to exclude from or include in the backup. For instructions on how to set ownership of a virtual machine, see “[Setting ownership of virtual machines](#)” on page 354.

Limitations

- *Only if you plan to restore individual files.* If you exclude all virtual machine disks from the backup and leave only the volume groups attached to the virtual machine, you will not be able to restore individual files.
- *For SQL Server:* Excluding or including disks in the backup is not supported if the Optimized SQL Server HADR protection option is enabled.
- *For Exchange Server:* Excluding or including disks in the backup is not supported if the Optimized Exchange Server DAG protection option is enabled.

Considerations

- The next backup after changing the virtual machine backup scope will be a full backup.
- Excluding disks with protected applications may affect application protection.
- If any disks are excluded from the backup (manually or automatically), the virtual machine will be restored or migrated to cloud without such disks or

with blank disks if you select the option to create excluded disks as blank. The corresponding restore point labels are marked with a red circle. For details, see “[Viewing entity details](#)” on page 313.

- *For vSphere virtual machines:* If independent and/or RDM disks are attached to the virtual machine, they are excluded from the backup automatically. Keep in mind that the option to create excluded disks as blank when restoring data or migrating data to cloud is available only for independent disks and not for RDM disks.
- *For servers with dynamic disks:* Dynamic disks are automatically excluded from the backup.

Procedure

1. In the Virtual Machines panel, select the virtual machine whose disks and volume groups you want to exclude from or include in the backup, and then select  **Configuration**. The Configuration dialog box opens.
2. In the Exclude/Include vDisks tab, depending on whether you want to exclude or include disks and volume groups in the backup, do one of the following:

I want to...	Instructions
Exclude disks and volume groups from the backup.	<ol style="list-style-type: none"> a. Click Exclude selected vDisks, and then select the disks or volume groups that you want to exclude from the backup. b. Click Save. <p>Important For vSphere environments: If you plan to restore individual files, make sure not to exclude the operating system disk from the backup.</p>
Include disks and volume groups in the backup.	<ol style="list-style-type: none"> a. Click Include selected vDisks, and then select the disks or volume groups that you want to include in the backup. b. Click Save.

You can later make changes to the selection of the excluded or included disks.

Enabling DM snapshots

By default, HYCU uses LVM snapshots for Linux server data protection. However, you can also configure a Linux server to be backed up by using DM

snapshots.

Considerations

- Although you can configure HYCU to use DM snapshots for data protection, DM volumes are not supported. For details on supported volumes, see the *HYCU Compatibility Matrix*.
- For snapshot storage, you can specify a directory that is hosted on any volume that is excluded from the backup, or on an NFS share.
- *Only if you plan to use an NFS share for snapshot storage.* Make sure the connection to the NFS server has low latency and high throughput (10 GiBps or higher) to avoid system performance issues.

Procedure

1. In the Virtual Machines panel, select the virtual machine that you want to back up by using DM snapshots, and then select  **Configuration**. The Configuration dialog box opens.
2. In the Snapshots tab, use the **Enable DM snapshots** switch, and then specify the path to the directory that you want to use for snapshot storage (for example, `/mnt/nfs/snapshotdir`).
3. Click **Save**.

Backing up virtual machines

With HYCU, you can back up your virtual machines in a fast and efficient way.

 **Note** The procedure for backing up virtual machine templates is the same as for virtual machines. Therefore, you can follow the same instructions as for backing up virtual machines.

Prerequisites

- *For Nutanix ESXi clusters and vSphere environments:* You have the required backup privileges assigned. For details, see “[Assigning privileges to a vSphere user](#)” on page 456.
- *Only if you plan to protect servers or volume groups that are attached to a virtual machine by using iSCSI.* Credentials are assigned to servers that you want to protect or to virtual machines whose volume groups you want to protect. For instructions, see “[Enabling access to data](#)” on page 146.

Limitations

- Assigning a policy that has the Backup from replica policy option enabled to the HYCU backup controller is not supported.
- *For vSphere virtual machines residing on VMFS or NFS datastores:* If you select Snapshot as the backup target type in your policy, such a policy cannot be assigned to the virtual machine.
- *For Nutanix clusters:* Backing up virtual machines from their replicas created with Nutanix DR is supported only for Nutanix AOS version 6.5 or later.
- *For virtual machines that have Azure Disk Encryption enabled:* The key vault is not protected by HYCU.

Considerations

- If during virtual machine synchronization, a virtual machine cannot be found in a source environment, the status of this virtual machine and any discovered applications running on it is set to PENDING_REMOVAL. The policy is still assigned to the virtual machine and the applications, but you cannot perform any data protection actions (they are grayed out in HYCU). Depending on whether this virtual machine is found in the source environment during the time interval of two automatic virtual machine synchronization processes, the following happens:
 - *The virtual machine is found in the source environment:* Its status and the status of the applications running on it is changed to Protected.
 - *The virtual machine is not found in the source environment:* If the virtual machine still has at least one valid restore point available, its status and the status of the applications running on it is changed to Protected deleted. This means that the virtual machine that is deleted from the source is still considered protected and is not removed from HYCU.
- *For Nutanix clusters:* If you plan to migrate a protection domain with protected virtual machines from one cluster to another through Nutanix Prism and you want these virtual machines to remain protected, make sure that both these clusters are added to HYCU. The next virtual machine synchronization after migration will add the corresponding virtual machines to the list of the virtual machines on the cluster to which you migrated the protection domain. The migrated virtual machines have the same UUIDs as before the migration and also keep the assigned policies. Keep in mind that the next backup of such virtual machines will be a full backup.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, select the virtual machines that you want to back up.

 **Tip** You can update the list of virtual machines by clicking  **Refresh**. To narrow down the list of displayed virtual machines, you can use the filtering options described in “[Filtering and sorting data](#)” on [page 317](#).

2. Click  **Set Policy**.
3. From the list of available policies, select the preferred policy.
4. Click **Assign** to assign the policy to the selected virtual machines.

 **Note** When you assign the policy to the selected virtual machines, the same policy is also assigned to the applications running on them if these applications already have an assigned policy. In this case, the policy assigned to the virtual machines takes precedence over the policy assigned to the applications and is automatically assigned to the applications.

The backup is scheduled according to the values that you defined for your policy. If required, you can also perform a manual backup at any time. For details, see “[Performing a manual backup](#)” on [page 333](#).

Restoring virtual machines

HYCU enables you to restore either an entire virtual machine or only virtual disks (virtual machine disks and/or Nutanix volume groups attached to virtual machines) that became corrupted. You can also validate the virtual machine backup by creating a virtual machine clone.

 **Note** *For vSphere environments:* The procedure for restoring virtual machine templates is the same as for virtual machines. Therefore, you can follow the same instructions as for restoring virtual machines.

Prerequisites

- If you are restoring a virtual machine to the same source and you want the existing ISO image to be attached to the restored virtual machine, make sure the ISO image that was attached to the virtual machine at backup time still exists on the source at virtual machine restore time and its name and location are the same.
- *For restoring data from tape:* If the tape target is being actively used for archiving data, its mode should be set to Read Only. For details on how to edit a target, see [“Managing targets” on page 325](#).
- *For Nutanix ESXi clusters and vSphere environments:* You must have the required restore privileges assigned. For details, see [“Assigning privileges to a vSphere user” on page 456](#).
- *For servers:* At least one Nutanix cluster, vCenter Server, AWS GovCloud (US) region, or Azure Government subscription must be added to HYCU to provide a storage container for storing the restore data. For details on how to add sources to HYCU, see [“Adding sources” on page 53](#).
- *For AWS GovCloud (US) environments:* If you are a self-service group user, you must add an AWS GovCloud (US) account to HYCU. For details, see [“Adding an AWS GovCloud \(US\) account” on page 366](#).
- *For Azure Government environments:* If you are a self-service group user, you must add an Azure Government service principal to HYCU. For details, see [“Adding an Azure Government service principal” on page 371](#).

Limitations

- If you are restoring a virtual machine from one source to another, the ISO image that was attached to the virtual machine at backup time will not be attached to the restored virtual machine.
- You can restore a virtual machine for which UEFI boot mode is enabled only to a source that supports UEFI boot configuration.
- *For vSphere environments:* After you restore an encrypted virtual machine, the virtual machine data is not encrypted anymore. You must manually apply the encryption storage policy to the restored virtual machine to re-enable encryption.
- *For Azure Government environments:* The OS profile of a virtual machine cannot be restored.

Considerations

- A restore is performed from the snapshot only if you are restoring to the same source (the source where the original virtual machine was running). If you are restoring to a different source, depending on the tier that you select for the restore, the following will happen:
 - If you select Snapshot, the restore will fail.
 - If you select Automatic, the restore will be performed from the target if there is an available target. Otherwise, it will fail.
- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data or validating the virtual machine backup.
- You cannot perform a restore of a virtual machine whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up` configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see [“Customizing HYCU configuration settings” on page 523](#).
- *For AWS GovCloud (US) virtual machines with encrypted volumes:* Depending on whether you are restoring such a virtual machine to the same or a different region, the following applies:
 - To the same region: The restored volumes will be encrypted with the same KMS key as the original ones.
 - To a different region: The restored volumes will be encrypted with the default KMS managed key for EBS encryption.

Restore options

Virtual machine restore options

You can select among the following virtual machine restore options:

I want to...	VM restore option	Description
Restore a virtual	Restore VM	Enables you to restore a virtual machine to the same source. Select this option if you want to

I want to...	VM restore option	Description
machine to the original location.		<p>replace the original virtual machine with the restored one. For instructions, see “Restoring a virtual machine” on the next page.</p> <p>Important You cannot restore servers by using this option.</p>
Restore a virtual machine by cloning it to a new location.	Clone VM	<p>Enables you to restore a virtual machine by creating a virtual machine clone on the same or a different source. Select this option if you want to keep the original virtual machine. For instructions, see “Cloning a virtual machine” on page 171.</p> <p>Important If you plan to restore a virtual machine to a different source, keep in mind that depending on your virtual machine original environment and target environment, you might have to perform some additional steps after the restore. For details, see “After restoring a virtual machine to a different source” on page 530.</p>
Validate a virtual machine backup.	Validate VM backup	<p>Enables you to validate the virtual machine backup by creating a virtual machine clone. Select this option if you want to verify that the virtual machine has no corrupted backups. For instructions, see “Validating the virtual machine backup” on page 188.</p>

Disk restore options

You can select among the following disk restore options:

Disk restore option	Description
Restore vDisks	Enables you to restore virtual disks. Select this option if you want to replace the original virtual disks with the

Disk restore option	Description
	<p>restored ones. For instructions, see “Restoring virtual disks” on page 193.</p> <p>ⓘ Important You cannot restore server disks by using this option.</p>
Clone vDisks	<p>Enables you to restore virtual disks by creating their clones. Select this option if you want to keep the original virtual disks. For instructions, see “Cloning virtual disks” on page 195.</p>
Export vDisks	<p>Enables you to restore virtual disks to an NFS or SMB share. Select this option if you want to make the virtual disks available to users with specific access permissions, or if you want to use the virtual disks later to restore data to a server or to a source not supported by HYCU or not added to HYCU. For instructions, see “Exporting virtual disks” on page 197.</p>

Restoring a virtual machine

You can restore a virtual machine to its original or a new location on the same source. In this case, the original virtual machine will be overwritten.

For details on how to restore a virtual machine, depending on your data protection environment, see one of the following sections:

- “Restoring a virtual machine to a Nutanix cluster or a vSphere environment” on the next page
- “Restoring a virtual machine to an AWS GovCloud (US) environment” on page 164
- “Restoring a virtual machine to an Azure Government environment” on page 168

Restoring a virtual machine to a Nutanix cluster or a vSphere environment

Considerations

- *Only if volume groups are attached to the virtual machine that you are restoring.*
You can choose to restore the volume groups together with the virtual machine if they were attached to it at backup time. In this case, the original volume groups are deleted and the restored ones are automatically attached to the restored virtual machine as well as all other virtual machines to which they were attached at backup time.
- The restored virtual machine retains the original MAC address.
- *Only if you plan to restore a vSphere virtual machine.* Depending on how you plan to restore data, consider the following:
 - *From a target:* The original virtual machine and all its snapshots will be deleted as part of the restore process.
 - *From a snapshot:* The entire virtual machine will be reverted to the selected snapshot and any excluded or included disk configuration will be ignored.
- *Only if you plan to restore vSphere virtual machine data to the original storage container.* If the storage container is mounted to several hosts and the original host is powered off or in maintenance mode at restore time, data will be restored to the same storage container on a different host.
- *Only if you plan to restore a vSphere virtual machine to a datacenter that was not added to HYCU.* The protection status of such a virtual machine will be Protected deleted after the restore.
- *Only if you plan to restore a vSphere virtual machine from a datacenter that was removed from HYCU.* After you remove the datacenter, the protection status of the virtual machine changes from Protected to Protected deleted. When restoring such a virtual machine, consider the following:
 - If restoring to a datacenter that is added to HYCU, the protection status of the virtual machine changes back to Protected.
 - If restoring to any datacenter that is not added to HYCU, the protection status of the virtual machine stays Protected deleted.
- *Only if you plan to restore a virtual machine running on a Nutanix ESXi cluster.* If Snapshot is selected as the backup target type in your policy, the NVRAM file will not be restored.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click the virtual machine that you want to restore. The Detail view appears at the bottom of the screen.

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

2. In the Detail view, select the preferred restore point.
3. Click  **Restore**.
4. Select **Virtual machine options**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. In the General section, do the following:

- a. From the Storage container drop-down menu, select where you want to restore the virtual machine. By default, the original storage container is selected.

 **Note** If you decide to restore the virtual machine to another storage container, keep in mind the following:

- Restore from the Snapshot tier cannot be performed to another storage container.
- If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.

- b. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
- c. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original

virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
- In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

 **Note** The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- d. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore. The original virtual machine will be deleted automatically.

 **Important** Only if you are restoring a vSphere virtual machine to a vSphere environment and you have disabled the Power virtual machine on switch. When you try to power on the virtual machine and you are prompted to answer whether the virtual machine has been moved or copied, make sure to answer **I Moved It**.

- e. Only if virtual disks have been excluded from the backup (manually or automatically): Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
- f. For volume groups attached to the virtual machine: Use the **Restore volume groups** switch if you want to restore also the volume groups that are attached to the virtual machine.

7. In the Network section, review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

Depending on whether the original networks are available, proceed as follows:

- If the original networks are available, you can leave the default values and restore the virtual machine with the original network settings, or you

can modify the network settings.

- If the original networks are not available, you must modify the network settings.

Modifying network settings

Original networks are...	Instructions
Available	<p>You can do the following:</p> <ul style="list-style-type: none"> • Add a new network adapter by clicking  New, and then selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting it, and then clicking  Edit and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking  Delete.
Unavailable	<p>You can do the following:</p> <ul style="list-style-type: none"> • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking  Edit and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking  Delete. • Add a new network adapter by clicking  New, and then selecting the preferred network.

 **Note** You can restore the virtual machine without a network adapter. Make sure to configure the network settings on the virtual machine afterward.

8. Click **Restore**.

 **Note** For Nutanix ESXi clusters: Because the minimum RAM required for restoring a virtual machine is 256 MiB, any virtual machine with less RAM is automatically set to 256 MiB during the restore.

Restoring a virtual machine to an AWS GovCloud (US) environment

Considerations

- Make sure that the virtual machine you are restoring is not deleted from AWS GovCloud (US). If you delete a virtual machine from AWS GovCloud (US), you cannot restore it even if it still has a valid restore point available in HYCU (that is, even if its status is Protected deleted).
- When restoring a virtual machine, the original virtual machine disks are deleted and replaced with the restored ones.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click the virtual machine that you want to restore. The Detail view appears at the bottom of the screen.
 **Note** The Detail view appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Detail view.
2. In the Detail view, select the preferred restore point.
3. Click  **Restore**.
4. Select **Virtual machine options**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. The following information is displayed and preselected:
 - The AWS GovCloud (US) account to which the virtual machine will be restored.
 - The account ID of the AWS GovCloud (US) account to which the virtual machine will be restored.
 - The region to which the virtual machine will be restored.
 - *Only if found by HYCU*. The key pair name for connection to the restored virtual machine. **Important** If HYCU does not find the name of your key pair and you want to use it, you can select it from the Key pair name drop-

down menu.

- The availability zone to which the virtual machine will be restored.

7. Click **Next**.

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

- **Automatic**: Ensures the fastest restore to the latest state.
- **Backup**
- **Copy**
- **Archive**
- **Snapshot**

9. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU threads field, enter the number of CPUs for the restored virtual machine multiplied by the number of cores per CPU and the number of threads per core.
- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- From the Virtual machine type drop-down menu, select the virtual machine type.

 **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

10. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.

11. Under Network interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the Virtual Private Cloud (VPC) to which the original virtual machine belongs. If required, you can also modify network settings.

Modifying network settings

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

 **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same Virtual Private Cloud (VPC). The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
 - a. The Virtual Private Cloud (VPC) to which the network interface will be added is displayed and preselected.
 - b. From the Subnets drop-down menu, select the subnet to which the network interface should be assigned.
 - c. From the Security groups drop-down menu, select one or more security groups that will be associated with the network interface. If you want to select all the available security groups, select **Select all**.
 - d. In the Public address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Auto-assign	An automatically allocated public IP address will be assigned to the network interface on the restored virtual machine.
Elastic IP (Reserved)	An elastic public IP address that you reserved in AWS GovCloud (US) will be assigned to the network interface on the restored virtual machine.
Elastic IP (New)	An elastic public IP address will be assigned to the network interface on the restored virtual machine.

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

Option	Description
Auto-assign	An automatically allocated private IP address will be assigned to the network interface on the restored virtual machine.
Custom	A private IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

12. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:

- Linux**
- Windows**

13. Under Operating system license, select one of the following options:

OS license option	Select this option if you want to...
Keep existing license	Keep the existing OS license on the restored virtual machine. Important Make sure that the existing license is applicable also in AWS.
<i>Available only for the Windows Server OS.</i> Replace existing license with AWS license	Replace the existing OS license with an AWS license on the restored virtual machine.

14. Click **Restore**.

Restoring a virtual machine to an Azure Government environment

Consideration

If you want the restored virtual machine to have the same static IP address as the original virtual machine, do one of the following:

- Before the restore, in Azure Government, disassociate the IP address from the original virtual machine, and then select this IP address for the network interface during the restore in HYCU.
- During the restore, select a different IP address for the network interface. After the restore, in Azure Government, assign the preferred IP address to the restored virtual machine.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click the virtual machine that you want to restore. The Detail view appears at the bottom of the screen.
 **Note** The Detail view appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Detail view.
2. In the Detail view, select the preferred restore point.
3. Click  **Restore**.
4. Select **Virtual machine options**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. From the Location drop-down menu, select the geographic region for the restored virtual machine.
7. From the Availability zone drop-down menu, select the zone for the restored virtual machine.
 **Note** The selected geographic region and the size of the virtual machine determine to which zones you can restore data. If you do not want to restore data to any zone, select None.
8. Click **Next**.

9. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic:** Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**

10. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU cores field, enter the number of virtual CPUs for the restored virtual machine.
- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- From the Virtual machine type drop-down menu, select the virtual machine type.

 **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

11. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
12. Under Network interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the subscription to which the original virtual machine belongs. If required, you can also modify network settings.

Modifying network settings

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

 **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
 - Only if you are adding a network interface.* From the Network drop-down menu, select the network for the network interface.
 -  **Note** The list of available networks includes only the ones within the region you selected for the restored virtual machine.
 - Select the subnet to which the network interface should be assigned.
 - In the Public IP address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	A static IP address will be assigned to the network interface on the restored virtual machine.
Existing	A preferred public IP address resource that you have created in Azure Government will be assigned to the network interface on the restored virtual machine.

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

options:

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the restored virtual machine.

- e. Click **Add** or **Save**.
 - Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot restore the virtual machine without a network interface.
13. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
 - **Linux**
 - **Windows**
14. Click **Restore**.

Cloning a virtual machine

You can create a clone of the original virtual machine by restoring the virtual machine to its original or a new location on the same or a different source. In this case, the original virtual machine will not be overwritten.

For details on how to clone a virtual machine, depending on your data protection environment, see one of the following sections:

- “[Cloning a virtual machine to a Nutanix cluster or a vSphere environment](#)” on the next page
- “[Cloning a virtual machine to an AWS GovCloud \(US\) environment](#)” on page 177
- “[Cloning a virtual machine to an Azure Government environment](#)” on page 181

 **Important** After you clone the virtual machine, make sure that everything works as expected by going through the considerations and the recommendations listed in “[After cloning a virtual machine](#)” on page 185.

Cloning a virtual machine to a Nutanix cluster or a vSphere environment

Prerequisites

- *For virtual machines that you plan to clone to a new location:* A Nutanix cluster or a vCenter Server for a vSphere environment to which you plan to clone the virtual machine must be added to HYCU. For details on how to do this, see “[Adding a Nutanix cluster](#)” on page 54 or “[Adding a vCenter Server](#)” on page 57.
- *For Linux servers:* In the `/etc/fstab` system configuration file of the server, UUIDs (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`) must be used instead of device names for file system device identification.

Limitation

For vSphere environments: Attaching the ISO image to the restored virtual machine is not supported.

Considerations

- *Only if volume groups are attached to the virtual machine that you are cloning.* You can choose to restore the volume groups together with the virtual machine if they were attached to it at backup time. In this case, the original volume groups are kept alongside of the restored ones. If the volume groups are also attached to other virtual machines, the following applies (depending on how they are attached to the virtual machines):
 - **Directly:** Volume groups are automatically attached only to the cloned virtual machine.
 - **By using iSCSI:** Volume groups are automatically attached to all virtual machines to which they were attached at backup time.
- *For restoring a virtual machine running on a Nutanix AHV cluster to a Nutanix ESXi cluster:* If virtual machine disks are attached to the PCI bus, the bus type will be automatically changed to SCSI after the restore. Because of this configuration change, the restore finishes with a warning.
- *For Linux virtual machines running on a Nutanix ESXi cluster:* If after restoring a virtual machine that was created through the vSphere (Web) Client, the virtual machine does not boot, follow the steps described in “[After restoring a virtual machine to a Nutanix ESXi cluster](#)” on page 532.

- After you restore a virtual machine, it might happen that the order of virtual disks differs from the one on the original virtual machine if you performed the restore:
 - From a Nutanix AHV cluster to a Nutanix ESXi cluster or a vSphere environment
 - From a Nutanix ESXi to another Nutanix ESXi cluster
 - From a vSphere environment to a Nutanix ESXi cluster
 In this case, make the necessary adjustments, including the selection of the correct boot disk.
- *Only if you plan to restore vSphere virtual machine data to the original storage container.* If the storage container is mounted to several hosts and the original host is powered off or in maintenance mode at restore time, data will be restored to the same storage container on a different host.
- *Only if ownership is set for the virtual machine.* The same owner is automatically assigned to the restored virtual machine.
- *Only if you plan to restore a virtual machine running on a Nutanix ESXi cluster.* If Snapshot is selected as the backup target type in your policy, the NVRAM file will not be restored.
- *Only if the original virtual machine resides on a source other than a vSphere environment.* Make sure to modify the virtual machine configuration by specifying the appropriate guest operating system.

Recommendation

For Linux virtual machines: It is recommended that the use of persistent network device names based on MAC addresses is disabled. Otherwise, you will have to configure the network manually. For details on how to disable the use of persistent network device names, see your Linux distribution documentation.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click the virtual machine that you want to restore. The Detail view appears at the bottom of the screen.

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

2. In the Detail view, select the preferred restore point.
3. Click  **Restore**.
4. Select **Virtual machine options**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. In the General section, do the following:
 - a. From the Storage container drop-down menu, select the storage container where you want to restore the virtual machine.

 **Note** By default, the original storage container is selected. If you decide to restore the virtual machine to another storage container, keep in mind the following:

- Restore from the Snapshot tier cannot be performed to another storage container.
- If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.
- If the selected storage container is on a different source, additional prerequisites apply. For details, see “[After restoring a virtual machine to a different source](#)” on page 530.

- b. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
- c. In the New VM name field, specify a new name for the virtual machine.
- d. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
- In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

 **Note** The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- e. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore.

 **Important** Make sure to consider the following:

- This option is disabled for virtual machines that have volume groups attached by using iSCSI. For details on what needs to be done before turning on the restored virtual machine, see “[After cloning a virtual machine](#)” on page 185.
- *Only if you are cloning a virtual machine from a Nutanix cluster or a vSphere environment.* If you turn the restored virtual machine on, the original virtual machine will be turned off automatically.
- *Only if you are cloning a vSphere virtual machine to a vSphere environment and you have disabled the Power virtual machine on switch.* When you try to power on the virtual machine and you are prompted to answer whether the virtual machine has been moved or copied, make sure to answer **I Copied It**.
- f. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
- g. *For volume groups attached to the virtual machine:* Use the **Clone volume groups** switch if you want to restore also the volume groups that are attached to the virtual machine.
- 8. In the Network section, do the following:
 - a. Review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual

machine was connected). If any of the original networks is no longer available, N/A is shown.

Depending on whether the original networks are available, proceed as follows:

- If the original networks are available, you can leave the default values and clone the virtual machine with the original network settings, or you can modify the network settings.
- If the original networks are not available, you must modify the network settings.

Modifying network settings

Original networks are...	Instructions
Available	<p>You can do the following:</p> <ul style="list-style-type: none"> • Add a new network adapter by clicking  New and selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting the virtual adapter, and then clicking  Edit and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking  Delete.
Unavailable	<p>You can do the following:</p> <ul style="list-style-type: none"> • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking  Edit and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking  Delete. • Add a new network adapter by clicking  New, and then selecting the preferred network.

 **Note** You can clone the virtual machine without a network adapter. Make sure to configure the network settings on the virtual machine afterward.

- b. *Only if you are restoring the virtual machine to a different Nutanix cluster or vSphere environment.* Use the **Keep original MAC address** switch if you want the restored virtual machine to keep the original MAC address. Keep in mind that this is applicable only if at least one network adapter has a MAC address assigned.

9. Click **Restore**.

Cloning a virtual machine to an AWS GovCloud (US) environment

Prerequisite

For virtual machines that you plan to restore to a new location: The AWS GovCloud (US) region to which you plan to restore the virtual machine must be added to HYCU. For instructions, see “[Adding an AWS GovCloud \(US\) region](#)” on page 59.

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for restoring data to a new location.
- *For virtual machines that have BitLocker volumes encrypted with TPM-based keys:* Restoring such volumes is not supported.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click the virtual machine that you want to restore. The Detail view appears at the bottom of the screen.

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

2. In the Detail view, select the preferred restore point.
3. Click  **Restore**.
4. Select **Virtual machine options**, and then click **Next**.

5. Select **Clone VM**, and then click **Next**.
6. From the Destination source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. From the AWS GovCloud (US) account, select the account to which the virtual machine will be restored.

The following information is displayed and preselected:

- The account ID of the AWS GovCloud (US) account to which the virtual machine will be restored.
- The region to which the virtual machine will be restored.

8. *Optional.* From the Key pair name drop-down menu, select the key pair name that you want to use for connection to the restored virtual machine.

! Important *For Windows virtual machines:* The key pair name that you select can be used only if the EC2Config or EC2Launch service was configured on the original virtual machine or if you configure it later on the restored virtual machine.

9. From the Availability zone drop-down menu, select the availability zone to which the virtual machine will be restored.
10. Click **Next**.
11. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic:** Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
12. In the New VM name field, specify a name for the restored virtual machine.
13. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.
If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:
 - In the vCPU threads field, enter the number of CPUs for the restored virtual machine multiplied by the number of cores per CPU and the number of threads per core.

- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- From the Virtual machine type drop-down menu, select the virtual machine type.

 **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

14. *Only if virtual disks have been excluded from the backup (manually or automatically)*: Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
15. Under Network interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the Virtual Private Cloud (VPC) to which the original virtual machine belongs. If required, you can also modify network settings.

Modifying network settings

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

 **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same Virtual Private Cloud (VPC). The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
 - a. The Virtual Private Cloud (VPC) to which the network interface will be added is displayed and preselected.
 - b. From the Subnets drop-down menu, select the subnet to which the network interface should be assigned.
 - c. From the Security groups drop-down menu, select one or more security groups that will be associated with the network interface. If you want to select all the available security groups, select **Select all**.

d. In the Public address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Auto-assign	An automatically allocated public IP address will be assigned to the network interface on the restored virtual machine.
Elastic IP (Reserved)	An elastic public IP address that you reserved in AWS GovCloud (US) will be assigned to the network interface on the restored virtual machine.
Elastic IP (New)	An elastic public IP address will be assigned to the network interface on the restored virtual machine.

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

Option	Description
Auto-assign	An automatically allocated private IP address will be assigned to the network interface on the restored virtual machine.
Custom	A private IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

16. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:

- **Linux**
- **Windows**

17. Under Operating system license, select one of the following options:

OS license option	Select this option if you want to...
Keep existing license	Keep the existing OS license on the restored virtual machine. ① Important Make sure that the existing license is applicable also in AWS.
<i>Available only for the Windows Server OS.</i> Replace existing license with AWS license	Replace the existing OS license with an AWS license on the restored virtual machine.

18. Click **Restore**.

Cloning a virtual machine to an Azure Government environment

Prerequisites

- *For virtual machines that you plan to restore to a new location:* The Azure Government subscription to which you plan to restore the virtual machine must be added to HYCU. For details on how to do this, see “[Adding an Azure Government subscription](#)” on page 61.
- *For virtual machines that have Azure Disk Encryption enabled:* The key vault must be available on the location to which you are restoring the virtual machine.

Limitation

If a restore point contains only a Snapshot tier, you cannot use it for restoring data to a new location.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click the virtual machine that you want to restore. The Detail view appears at the bottom of the screen.

Note The Detail view appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Detail view.

2. In the Detail view, select the preferred restore point.
3. Click  **Restore**.
4. Select **Virtual machine options**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. From the Service principal drop-down menu, select the service principal that has access to the required resources (the source from which and to which you are restoring the virtual machine).
8. From the Subscription drop-down menu, select the subscription for the restored virtual machine.
9. From the Resource group drop-down menu, select the resource group for the restored virtual machine.
10. From the Location drop-down menu, select the geographic region for the restored virtual machine.
11. From the Availability zone drop-down menu, select the zone for the restored virtual machine.

Note The selected geographic region and the size of the virtual machine determine to which zones you can restore data. If you do not want to restore data to any zone, select **None**.

12. Click **Next**.
13. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
14. In the New VM name field, specify a name for the restored virtual machine.
15. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original

virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- From the Virtual machine type drop-down menu, select the virtual machine type.

 **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

16. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
17. Under Network interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the subscription to which the original virtual machine belongs. If required, you can also modify network settings.

Modifying network settings

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

 **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:

- a. *Only if you are adding a network interface.* From the Network drop-down menu, select the network for the network interface.

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

- b. Select the subnet to which the network interface should be assigned.
- c. In the Public IP address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	A static IP address will be assigned to the network interface on the restored virtual machine.
Existing	A preferred public IP address resource that you have created in Azure Government will be assigned to the network interface on the restored virtual machine.

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

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the restored virtual machine.

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

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

18. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
 - **Linux**
 - **Windows**
19. Click **Restore**.

After cloning a virtual machine

After cloning a virtual machine, go through considerations and recommendations listed in this section to make sure that everything works as expected.

Considerations

- If a new MAC address is assigned to a network adapter on the cloned virtual machine, make sure that the guest operating system is configured appropriately to connect the cloned virtual machine to the selected network.
- If after cloning a virtual machine from a Nutanix AHV cluster to a Nutanix ESXi cluster or a vSphere environment, the virtual machine does not turn on due to an IDE device not being configured properly, you must edit the IDE device configuration manually. For details on how to do this, see VMware documentation.
- *For vSphere environments:* Some operating systems (for example, Red Hat Enterprise Linux 7) might require network configuration. For details, see VMware documentation.
- *For virtual machines to which volume groups are attached by using iSCSI:* Because the original virtual machine and the restored one have the same network and iSCSI configuration settings after the restore, make sure both the virtual machines are not turned on at the same time to avoid any potential issues. As one way of preventing issues, you can disconnect the restored virtual machine from the network before turning it on and make the required changes such as replacing the network adapter and updating the iSCSI settings on it.
- *For servers:*
 - *Only if you cloned a Windows server to a Nutanix ESXi cluster.* Make sure to modify the machine configuration after the restore by specifying the appropriate guest OS and to install the latest version of VMware Tools on the machine. For detailed information, see VMware documentation.

- Only if you cloned a Linux server that uses UEFI firmware to a Nutanix AHV cluster. If the virtual machine does not boot after the restore, reboot the machine.

Recommendation

For Linux servers: Because the original boot loader of the server is replaced with a temporary one during the backup, it is recommended to update the boot configuration after the restore. Depending on what firmware the server uses, see one of the following sections for instructions on how to do this:

- “[Updating the boot configuration of Linux servers that use BIOS firmware](#)” below
- “[Updating the boot configuration of Linux servers that use UEFI firmware](#)” on the next page

Updating the boot configuration of Linux servers that use BIOS firmware

Procedure

- In the `/etc/default/grub` system configuration file, do the following:
 - Edit the `GRUB_CMDLINE_LINUX` option and remove the following kernel parameters (if present):
 - `rd.lvm`. (except `rd.lvm=0`)
 - `rd.md`. (except `rd.md=0`)
 - `rd.dm`. (except `rd.dm=0`)
 - `rd.luks`.
 - Set the resume device on the virtual machine to match the resume device UUID on the original server. For example, if the resume device on the original server is `resume=/dev/mapper/cl-swap`, the resume device on the virtual machine should be `resume=UUID=4044243b-612b-42bc-ba22-4736c4eadde6`.
- Optional.* If you want to speed up the boot process and skip mounting non-existent volumes, in the `/etc/fstab` system configuration file, comment all the lines for volumes for which a warning was triggered at backup time.

Example

The following warning message was triggered:

```
Non LVM volumes detected: Following volumes are not
backupable: /dev/sdf3:/test_mount.
```

In the `/etc/fstab` system configuration file, comment the line that contains the `/test_mount` mountpoint.

3. Update the GRUB configuration by running the following command:

```
grub2-mkconfig -o /boot/grub2/grub.cfg
```

4. Install the boot loader on the boot disk by running the following command:

```
grub2-install /dev/sdc
```

Tip The boot disk is the one that contains the boot partition. To identify the boot partition, run the following command:

```
findmnt -nT /boot -o SOURCE
```

5. Reboot the virtual machine.

Updating the boot configuration of Linux servers that use UEFI firmware

Procedure

1. *Only if you cloned the server to a Nutanix ESXi cluster or a vSphere environment.* When the virtual machine enters the firmware setup mode, select the **Boot from file** option, and then specify the `<EFIPartition>/EFI/hycu/shimx64.efi` file. For details, see Nutanix or VMware documentation.
2. In the `/etc/default/grub` system configuration file, do the following:
 - a. Edit the `GRUB_CMDLINE_LINUX` option and remove the following kernel parameters (if present):
 - `rd.lvm`. (except `rd.lvm=0`)
 - `rd.md`. (except `rd.md=0`)
 - `rd.dm`. (except `rd.dm=0`)
 - `rd.luks`.
 - b. Set the resume device on the virtual machine to match the resume device UUID on the original server. For example, if the resume device on the original server is `resume=/dev/mapper/cl-swap`, the resume device on the virtual machine should be `resume=UUID=4044243b-612b-42bc-ba22-4736c4eadde6`.
3. *Optional.* If you want to speed up the boot process and skip mounting non-existent volumes, in the `/etc/fstab` system configuration file, comment all

the lines for volumes for which a warning message was triggered at backup time.

Example

The following warning message was triggered:

```
Non LVM volumes detected: Following volumes are not
backupable: /dev/sdf3:/test_mount.
```

In the `/etc/fstab` system configuration file, comment the line that contains the `/test_mount` mountpoint.

4. Update the GRUB configuration by running the following command:

- For Red Hat Enterprise Linux and Oracle Linux:

```
grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

- For CentOS:

```
grub2-mkconfig -o /boot/efi/EFI/centos/grub.cfg
```

5. Reboot the virtual machine.

6. *Only if secure boot was enabled on the original server and you use third-party kernel modules.* Enroll the Machine Owner Key (MOK) used to sign third-party kernel modules. For details on how to do this, see the respective operating system documentation.

7. Create the default boot entry in the UEFI firmware setup. The boot entry should point to the following system file:

- For Red Hat Enterprise Linux and Oracle Linux:

```
<EFIPartition>/EFI/redhat/shimx64.efi
```

- For CentOS:

```
<EFIPartition>/EFI/centos/shimx64.efi
```

Validating the virtual machine backup

You can validate the virtual machine backup by creating a virtual machine clone. In this case, the original virtual machine will not be overwritten and turned off. You can also specify whether you want to keep the virtual machine clone after the backup validation is performed.

 **Note** You can also set up a validation policy and schedule the backup validation according to the values that you define in your validation policy.

For details on how to do this, see “[Setting up a validation policy](#)” on [page 334](#).

Prerequisites

- If you are cloning the virtual machine to a vSphere environment, the latest version of VMware Tools must be installed on the virtual machine.
- *Only if you plan to specify the Advanced validation type.*
 - Credentials must be assigned to the virtual machine. For prerequisites, limitations, considerations, and instructions, see “[Enabling access to application data](#)” on [page 209](#).
 - A network card must be added to the virtual machine.

Limitation

Performing the backup validation is not supported for the HYCU backup controller, AWS GovCloud (US) environments, and Azure Government environments.

Considerations

- Network conflicts may occur during the backup validation if the virtual machine is configured with a static IP address, resulting in unreliable backup validation data.
- *Only if you plan to specify the Advanced validation type when performing the backup validation for a Windows virtual machine.* Checking for disk errors may fail in some cases, which does not mean that your virtual machine is corrupted. However, it is highly recommended that you check the status of such a virtual machine manually.
- After you perform the backup validation, consider the following:
 - You can view the backup validation status of a virtual machine in the Validation column in the Virtual Machines panel (represented by an icon). By pausing on the icon, you can also see which validation policy is assigned to the virtual machine, if you have set it up. For details on setting up a validation policy, see “[Setting up a validation policy](#)” on [page 334](#).
 - The Exclude policy is automatically assigned to the cloned virtual machine.

Procedure

1. In the Virtual Machines panel, click the virtual machine for which you want to perform the backup validation. The Detail view appears at the bottom of the screen.

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

2. In the Detail view, select the preferred restore point.
3. Click  **Restore**.
4. Select **Virtual machine options**, and then click **Next**.
5. Select **Validate VM backup**, and then click **Next**.
6. From the Storage container drop-down menu, select where you want to clone the virtual machine for which you are performing the backup validation.
7. From the Restore from drop-down menu, select which tier you want to use for the backup validation. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**

 **Note** If you select Automatic, the tier for the backup validation is by default selected in the following priority order: Backup > Copy > Archive > Snapshot. This means that HYCU will always use the first available tier in the specified order for the backup validation. However, you can at any time change this default behavior by customizing the `backup.validation.restore.source.priority.order` configuration setting in the HYCU `config.properties` file and adjusting the tier order to your data protection needs. For details on how to customize HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on [page 523](#).

8. In the New VM name field, specify a name for the cloned virtual machine.
9. Use the **Use original VM configuration** switch if you want the cloned virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU(s) field, enter the number of virtual CPUs for the cloned virtual machine.
- In the Cores per vCPU field, enter the number of cores per virtual CPU for the cloned virtual machine.

 **Note** The total number of cores of the cloned virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

- In the Memory field, set the amount of memory (in GiB or MiB) for the cloned virtual machine.

10. From the Keep VM after validation drop-down menu, depending on whether you want to keep the virtual machine after the backup validation is performed, select one of the following options:

Option	Description
Always	The virtual machine will be kept after the backup validation is performed.
On validation error	The virtual machine will be kept after the backup validation is performed only if a validation error occurs during the validation.
Never	The virtual machine will be automatically deleted after the backup validation is performed.

11. From the Validation type drop-down menu, select one of the following types:

Validation type	Description
Basic	During the backup validation, the following tasks will be performed: <ul style="list-style-type: none"> • The virtual machine will be cloned and turned on. • The guest OS will be shut down.
Advanced	During the backup validation, the following tasks will be performed:

Validation type	Description
	<ul style="list-style-type: none"> • The virtual machine will be cloned and turned on. • Any applications running on the virtual machine will be discovered. • Virtual disks will be validated, which includes checking the virtual machine file system and existing disks on the virtual machine. For Windows virtual machines, checking for disk errors is also performed. • The custom scripts will be run, if specified. • The guest OS will be shut down.

12. *Only if you selected the Advanced validation type.* Do the following:

- Enable the **Run custom script** switch if you want the custom script to be run on the virtual machine as part of the backup validation process, and then make sure that the proper path to the script is specified.

 **Note** The script returns an exit code of 0 for success and any other value for failure.

- In the Network section, review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

Depending on whether the original networks are available, proceed as follows:

- If the original networks are available, you can leave the default values and clone the virtual machine with the original network settings, or you can modify the network settings.
- If the original networks are not available, you must modify the network settings.

Modifying network settings

Original networks are...	Instructions
Available	<p>You can do the following:</p> <ul style="list-style-type: none"> • Add a new network adapter by clicking  New,

Original networks are...	Instructions
	<p>and then selecting the preferred network.</p> <ul style="list-style-type: none"> • Edit the existing network adapter to connect the virtual machine to a different network by selecting it, and then clicking  Edit and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking  Delete.
Unavailable	<p>You can do the following:</p> <ul style="list-style-type: none"> • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking  Edit and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking  Delete. • Add a new network adapter by clicking  New, and then selecting the preferred network.

13. Click **Validate**.

Restoring virtual disks

You can restore virtual disks to their original or a new location. In this case, the original virtual disks will be overwritten.

Limitations

- Restoring server disks by using the Restore vDisks option is not supported.
- Restoring virtual disks is not supported for AWS GovCloud (US) and Azure Government environments.
- *Only if you plan to restore vSphere virtual disks.* The virtual disks that you plan to restore must be available on the original virtual machine. If the virtual disks are not available, the restore will fail.

Considerations

- If any virtual disks were excluded from the backup, you cannot select them for the restore. The corresponding restore point labels are marked with a red circle. For details, see “[Viewing entity details](#)” on page 313.
- The original virtual disks are deleted and the restored ones are automatically attached to all virtual machines to which they were attached at backup time.
- *Only if restoring volume groups attached to the virtual machine.* The virtual machines to which the volume groups are attached must be turned off.
- If you plan to restore vSphere virtual disks, keep in mind that the virtual machine will be powered off during restore.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click the virtual machine whose virtual disks you want to restore.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

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

3. Click  **Restore**.
4. Select **Disk options**, and then click **Next**.
5. Select **Restore vDisks**, and then click **Next**.
6. From the list of virtual disks that are available for the restore, select the ones that you want to restore, and then click **Next**.

 **Important** *Only if restoring volume groups attached to the virtual machine.* You cannot select individual disks, but only the whole volume group.

7. From the Storage container drop-down menu, select where you want to restore the virtual disks.

 **Note** By default, the original storage container is selected. If you decide to restore the virtual disks to another storage container, they will

not be restored from the snapshot, but from the target. Therefore, no fast restore will be performed.

8. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic:** This type of restore ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot** (*Nutanix and vSphere only*)
9. Click **Restore**.

Cloning virtual disks

You can create clones of virtual disks by restoring them to their original or a new location. In this case, the original virtual disks will not be overwritten.

Prerequisite

For vSphere environments: Make sure the number of disks on the virtual machine to which you are restoring data does not exceed 15 on SCSI controller 0.

Limitations

- Restoring virtual disks is not supported for AWS GovCloud (US) and Azure Government environments.
- Restoring virtual disks to the original virtual machine is not possible if the original virtual machine is powered off.
- Restoring virtual disks to a virtual machine running on a different source is not supported.

Considerations

- If any virtual disks are excluded from backup, you cannot select them for restore. The corresponding restore point labels are marked with a red circle. For details, see “[Viewing entity details](#)” on page 313.
- *Only if restoring volume groups attached to the virtual machine.* The original volume groups are kept alongside of the restored ones and the following applies regarding their attachment:

- If you are restoring the volume groups to the original virtual machine, they are attached to all the virtual machines to which they were attached at backup time.
- If you are restoring the volume groups to a virtual machine other than original running on a Nutanix AHV cluster, they are attached only to the selected virtual machine. If you are restoring the volume groups to a virtual machine other than original running on a Nutanix ESXi cluster, you must attach them manually after the restore.

The name format of the cloned volume groups is as follows:

`<OriginalVGName>-<Timestamp>`

- For virtual machine disks:

- The original virtual machine disks are kept alongside the restored ones that are automatically attached to the virtual machine as the first available interface index (per interface type). For example, if you have the `scsi.0`, `scsi.1`, and `scsi.4` virtual disks already attached to your virtual machine, the restored one will be `scsi.2`.
- If the bus type of the original virtual disks is IDE, it is automatically changed to SCSI during the restore.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click the virtual machine whose virtual disks you want to restore.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

 **Note** The Detail view appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Detail view.
3. Click  **Restore**.
4. Select **Disk options**, and then click **Next**.
5. Select **Clone vDisks**, and then click **Next**.
6. From the list of virtual disks that are available for the restore, select the ones that you want to restore, and then click **Next**.

ⓘ Important Only if restoring volume groups attached to the virtual machine. You cannot select individual disks, but only the whole volume group.

7. From the Select VM drop-down menu, select the virtual machine to which you want to attach the restored virtual disks. The restored virtual disks can be attached to the original virtual machine (the default selection) or any other virtual machine. Consider the following:
 - If you are attaching the virtual disks to the original virtual machine, make sure it is turned on.
 - You cannot attach the restored disks to a server.
8. From the Storage container drop-down menu, select where you want to restore the virtual disks.

ⓘ Note For virtual machines: You can select only among the storage containers that are created on the Nutanix cluster on which the selected virtual machine resides.

9. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: This type of restore ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot** (Nutanix clusters only)
10. Click **Restore**.

Exporting virtual disks

You can restore virtual disks to an NFS or SMB share. You can use exported virtual disks to restore data to a server. For details, see “Restoring data to a server” on page 200.

Prerequisites

- For restoring virtual disks to an SMB share: The SMB server must be configured to stop creating sparse files (the `strict_allocate` parameter must be set to `yes` in the `smb.conf` file).

- Only if you are restoring data that is stored in the archive storage tier on an Azure target. You must recreate a snapshot and use this snapshot for restoring data, or manually rehydrate data. For instructions on how to recreate a snapshot, see “[Recreating snapshots](#)” on page 340. For instructions on how to manually rehydrate data, see Azure documentation.

Limitation

For AWS GovCloud (US) and Azure Government environments: Restoring virtual disks from a snapshot is not supported.

Consideration

If any virtual disks were excluded from the backup, you cannot select them for the restore. The corresponding restore point labels are marked with a red circle. For details, see “[Viewing entity details](#)” on page 313.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click the virtual machine whose virtual disks you want to restore.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

 **Note** The Detail view appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Detail view.
3. Click  **Restore**.
4. Select **Disk options**, and then click **Next**.
5. Select **Export vDisks**, and then click **Next**.

 **Important** During the restore of virtual disks, you cannot perform additional restores or expire backups for this virtual machine.
6. From the list of virtual disks that are available for the restore, select the ones that you want to restore, and then click **Next**.
7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among

which you can select:

- **Automatic:** This type of restore ensures the fastest restore to the latest state.
- **Backup**
- **Copy**
- **Archive**
- **Snapshot** (*Nutanix clusters only*)

8. From the Type drop-down menu, select where you want to restore the virtual disks, and then provide the required information:

- **SMB**

- a. *Optional.* Enter the domain and user credentials.
- b. Enter the SMB server name or IP address and the path to the SMB shared folder from the root of the server (for example, /backups/HYCU).

- **NFS**

Enter the NFS server name or IP address and the path to the NFS shared folder from the root of the server (for example, /backups/HYCU).

9. Click **Restore**.

After exporting virtual disks

After the restore of the virtual disks is complete, you can use them to restore data to a server or to an environment with a source not supported by HYCU or not added to HYCU.

Data is restored to the following location:

`/<SharedPath>/<VMName>/<Timestamp>/<Filename>`

In this instance, `<SharedPath>` is the path to the shared folder, `<VMName>` is the virtual machine name, `<Timestamp>` is the time of the restore, and `<Filename>` is the virtual machine disk UUID.

What kind of files are created by the restore depends on the environment in which the virtual machine whose virtual disks you restored was backed up. Depending on the type of source in your environment, the following files are created for each selected disk:

Source	Files
Nutanix AHV	<p><i><DiskName></i> (without extensions) A raw image of the disk, including unallocated space as zeroes</p>
Nutanix ESXi	
vSphere	<ul style="list-style-type: none"> • <i><DiskName>-flat.vmdk</i> A raw image of the disk • <i><DiskName>.vmdk</i> A VMDK descriptor file, referencing <i><DiskName>-flat.vmdk</i>

Restoring data to a server

The procedure described in this section is an example of how to restore data to a Windows server.

Prerequisites

- The server to which you want to restore data must have the same number of disks as the original machine and the disk size must be equal to or greater than the original size.
- You have downloaded a Linux live CD (for example, Ubuntu) and booted it on the server where you want to restore your data.

Considerations

- Make sure you run all the commands as root.
- You can safely ignore the following error message:

The backup GPT table is corrupt, but the primary appears OK, so that will be used.

Procedure

1. Identify your destination disk.

Because HYCU performs the backup at the disk level, you must identify the path of each disk to which you will restore data. To list all the disks on your system, run the following command:

```
fdisk -l
```

The following is an example of the output:

```
Disk /dev/sda: 32 GiB, 34359738368 bytes, 67108864 sectors
```

Disk /dev/sdb: 5 GiB, 5368709120 bytes, 10485760 sectors

2. Mount the share to which you exported the disks.
3. Identify the path to the exported disks on the mounted share by running the following commands:

```
cd /<SharedPath>/<VMName>/<Timestamp>
```

```
ls
```

The following is an example of the output:

```
ServerDisk0 ServerDisk1
```

4. Verify each exported disk by running the following command:

```
fdisk -l <ExportedDiskName>
```

For example:

```
fdisk -l ServerDisk0
```

The information about the exported disk (for example, disk size and a list of partitions) is displayed. Use this information to identify a suitable destination disk for restoring the data. For example, the size of exported disk ServerDisk0 matches the size of disk /dev/sda. Therefore, disk ServerDisk0 can be restored to disk /dev/sda.

The following is an example of the output:

```
Disk ServerDisk0: 32 GiB, 34359738368 bytes, 67108864 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x36bab260

Device Boot Start End Sectors Size Id Type
ServerDisk0p1 * 2048 718847 716800 350M 7
HPFS/NTFS/exFAT
ServerDisk0p2 718848 67106815 66387968 31.7G 7
HPFS/NTFS/exFAT
```

5. Restore data by running the following command for each disk:

```
dd if=<ExportedDiskName> of=<DestinationDiskPath> bs=1024
status=progress
```

For example:

```
dd if=ServerDisk0 of=/dev/sda bs=1024k status=progress
```

The following is an example of the output:

```
33540483072 bytes (34 GB, 31 GiB) copied, 229 s, 146 MB/s
33554432+0 records in
33554432+0 records out
34359738368 bytes (34 GB, 32 GiB) copied, 229.78 s, 150 MB/s
```

- Eject the Linux live CD and reboot the server.

Restoring individual files

You can restore individual files to the same or a different virtual machine, to an SMB or NFS share, or to the local machine. This alternative to restoring an entire virtual machine allows you to restore only one or more files that have become corrupted or have been deleted for some reason and are now missing on the virtual machine.

Individual files can be restored from a target or a snapshot. A restore is always performed from the snapshot if the snapshot is available for the selected restore point (this speeds up the restore process). Otherwise, the restore is performed from the target if Target is selected as the backup target type in your policy (this saves space in your environment). If you want to restore individual files from a snapshot and no snapshot is available for the selected virtual machine restore point, HYCU enables you to manually recreate it. For details on how to do this, see “[Recreating snapshots](#)” on page 340.

You can use the pre-restore and post-restore scripts to perform necessary actions before and after the restore of individual files is performed. For details on how to specify the scripts, follow the procedure described in this section. For details on exit codes and exported environment variables, see “[Using the pre and post scripts](#)” on page 461.

Prerequisites

Windows	<ul style="list-style-type: none"> The NTFS, FAT, or FAT32 file system must be used.
---------	---

virtual machines	<ul style="list-style-type: none"> For improved restore performance, the startup type of the Microsoft iSCSI Initiator Service may not be set to Disabled. <i>For restoring files to a virtual machine:</i> <ul style="list-style-type: none"> <i>For Windows 8 and 10 virtual machines:</i> WinRM must be enabled and configured by using the <code>winrm quickconfig</code> command. A Windows operating system user account must exist. This account must have WinRM permissions granted and must be a member of the virtual machine's local Administrators group. Access to the virtual machine file system must be enabled. For instructions, see “Enabling access to data” on page 146. <i>For pre/post-restore scripts:</i> A script must be available in the accessible folder and must have one of the following extensions: <code>bat</code>, <code>ps1</code>, <code>cmd</code>. <i>For AWS GovCloud (US) and Azure Government environments:</i> The virtual machine to which you plan to restore files must be in the same virtual network as the HYCU backup controller.
Linux virtual machines	<ul style="list-style-type: none"> The FAT32, xfs, ext4/ext3/ext2, reiserfs, or btrfs file system must be used. <i>For restoring individual system files with the non-root sudo user and better overall restore performance:</i> <ul style="list-style-type: none"> The sudo user must have the NOPASSWD option set. For example, to set the NOPASSWD option for the user on a Red Hat Enterprise Linux 8.x system, add the following line to the <code>/etc/sudoers</code> file: <pre>restoreuser ALL=(ALL) NOPASSWD: ALL</pre> The <code>cifs-utils</code> package must be installed on virtual machines whose files you plan to restore. <i>For AWS GovCloud (US) and Azure Government environments:</i> The virtual machine to which you plan to restore files must be in the same virtual network as

	<p>the HYCU backup controller.</p> <ul style="list-style-type: none"> • References in the <code>/etc/fstab</code> system configuration file entries must use universally unique identifiers (for example, <code>UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5</code>) rather than device names (for example, <code>/dev/sda1</code>) unless the entries refer to logical volumes (for example, <code>/dev/mapper/ol-root</code>). • <i>For restoring files to a virtual machine:</i> <ul style="list-style-type: none"> ◦ Access to the virtual machines through ssh must be enabled. ◦ Access to the virtual machine file system must be enabled. For instructions, see “Enabling access to data” on page 146. ◦ <i>For pre/post-restore scripts:</i> A script must be available in the accessible folder and has the <code>sh</code> extension. You must have permissions to run a script on the virtual machine with the assigned credentials.
Nutanix ESXi clusters	<ul style="list-style-type: none"> • <i>For restoring files to a virtual machine:</i> The latest versions of VMware Tools and NGT must be installed on the client virtual machine. <p>For detailed information about installing VMware Tools, see VMware documentation. For detailed information about installing NGT, see Nutanix documentation.</p> <ul style="list-style-type: none"> • You must have the required restore privileges assigned. For details, see “Assigning privileges to a vSphere user” on page 456.
vSphere environments	<ul style="list-style-type: none"> • You must have the required restore privileges assigned. For details, see “Assigning privileges to a vSphere user” on page 456.
All environments	<ul style="list-style-type: none"> • <i>Only if you are restoring data that is stored in the archive storage tier on an Azure target.</i> You must recreate a snapshot and use this snapshot for restoring data, or manually rehydrate data. For instructions on how to recreate a snapshot, see “Recreating snapshots” on page 340. For instructions on how to manually rehydrate data, see Azure documentation.

Limitations

- Restoring individual files on multi-boot systems is not supported.
- Restoring individual files is not supported for virtual machines with encrypted disks, folders, or files.
- Restoring individual file data from tape is not supported.
- On Linux, you can restore symbolic links and soft links only to the original location.
- Restoring files from the same snapshot simultaneously by two different users is not possible.
- You cannot restore individual files if you excluded all virtual machine disks from the backup and left only the attached volume groups.
- *For restoring files to a different virtual machine:* You can restore files only to a virtual machine that belongs to the same operating system family as the original one.
- *For restoring files to a local machine:* You can download only a data archive whose size is less than or equal to 2 GiB.
- *For Windows virtual machines running on a Nutanix cluster that have Storage Replica enabled:* Restoring individual files to a virtual machine is supported only if the restore is performed from the target.
- *For Azure Government environments:* You cannot restore individual files if Azure Disk Encryption is enabled on the virtual machine.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- You cannot perform a restore of a virtual machine whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up` configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see [“Customizing HYCU configuration settings” on page 523](#).
- *For restoring files to a virtual machine:* To be able to restore some types of files (for example, system files), the account you specify to access a virtual machine must be a member of the virtual machine's local Administrators

group on Windows or have root permissions on Linux.

- If any virtual disks are excluded from backup, you cannot select them for restore. The corresponding restore point labels are marked with a red circle. For details, see “[Viewing entity details](#)” on page 313.
- *For using the Backup from replica option:* When restoring to the central or remote site (the original location), the restore is always performed from the snapshot on the central site.
- *For pre/post-restore scripts:* You can specify pre/post-restore scripts only when restoring files to a virtual machine.
- *Only if restoring files to an external distributed SMB share.* Make sure that the folder for the restore is precreated on the share and the shared path leads to this folder.

Recommendation

Only if restoring a large number of files. Instead of restoring individual files, it is highly recommended to restore disks hosting these files by using the Clone vDisks option. For instructions, see “[Cloning virtual disks](#)” on page 195.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click the virtual machine that contains the files that you want to restore to open the Detail view.

 **Note** The Detail view appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Detail view.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Files**.
4. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic:** Ensures the fastest restore to the latest state.
 - **Backup**

- **Copy**
- **Archive**
- **Snapshot**

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

Tip If there are too many files to be displayed on one page, you can move between the pages by clicking **>** and **<**.

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

7. Depending on where you want to restore the selected files (to the same or a different virtual machine, an external SMB or NFS share, or the local machine), select the preferred restore option, click **Next**, and then follow the instructions:

Restore option	Instructions
Restore to virtual machine	<p>a. On the General tab, do the following:</p> <ol style="list-style-type: none"> i. From the Virtual machine drop-down menu, select the virtual machine to which you want to restore the files. You can restore the files to the same or a different virtual machine. ii. Select whether you want to restore the files to the original location or an alternate location. If you select an alternate location, specify the path in the following format: <div style="background-color: #f0f0f0; padding: 5px; text-align: center;"> <code>C:\<Path></code> </div> <ol style="list-style-type: none"> iii. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file). iv. Use the Restore ACL switch if you want to restore the original access control list. <p>Important If the virtual machine is not accessible due to various reasons (for</p>

Restore option	Instructions
Restore to external share	<p>example, credentials are not assigned to it, discovery was not successful, or it is turned off or deleted from the source), you cannot select it for restoring the individual files.</p> <p>b. <i>Optional.</i> Click the Pre/Post Scripts tab, and then do the following:</p> <ol style="list-style-type: none"> In the Run pre-restore script field, enter the path to the script that HYCU will run before the restore is performed. In the Run post-restore script, enter the path to the script that HYCU will run after the restore is performed. <p>c. Click Restore.</p>
Download	<p>a. Select NFS or SMB for the share type, and then specify the path to a shared folder in the following format:</p> <div data-bbox="616 1096 870 1134" style="background-color: #f0f0f0; padding: 2px; border-radius: 5px; display: inline-block;"> <code>\server\<Path></code> </div> <p>b. <i>For SMB:</i> Optionally, provide user credentials to access the SMB share.</p> <p>c. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file).</p> <p>d. Click Restore.</p> <p>Click Download to restore the selected files to the local machine.</p> <p>Important Do not refresh the page or navigate away from the page until the download process job finishes.</p>

Chapter 5

Protecting applications

HYCU enables you to protect your application data with fast and reliable backup and restore operations. After you enable HYCU to access an application running on a virtual machine, complete the required preparatory steps, and back up the application, you can choose to restore either the whole application or only specific application items.

 **Note** The instructions for protecting applications residing on virtual machines apply also to applications residing on servers except where specifically stated otherwise.

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

- “Enabling access to application data” below
- “Planning application protection” on page 212
- “Backing up applications” on page 222
- “Restoring whole applications” on page 223
- “Restoring SQL Server databases” on page 252
- “Restoring Exchange Server databases, mailboxes, and public folders” on page 256
- “Restoring Oracle database instances and tablespaces” on page 260

Enabling access to application data

After you assign credentials to virtual machines as described in “Enabling access to data” on page 146, the process of application discovery starts automatically.

When the application discovery job completes, the discovered applications are listed in the Applications panel. HYCU supports different types of applications on virtual machines and servers. For a list of supported applications, see the *HYCU Compatibility Matrix*.

Depending on the Discovery status of the applications that you want to protect, do one of the following:

✓	<p>HYCU can access the discovered applications that you want to protect with the virtual machine credentials and you can start protecting such applications. For instructions, see “Backing up applications” on page 222.</p> <p>Note Access to Active Directory and SAP HANA is always granted with the virtual machine credentials.</p>
✗	<p>The virtual machine credentials do not have proper permissions and HYCU cannot access applications. To enable HYCU to access applications, do one of the following:</p> <ul style="list-style-type: none"> • If you want to use virtual machine credentials, reassign credentials to virtual machines so that they have proper permissions. For instructions on how to assign credentials to a virtual machine, see “Enabling access to data” on page 146. • If you want to use application-specific credentials, follow the procedure described in this section.

Prerequisites

Windows virtual machines	<ul style="list-style-type: none"> • For Windows 8 and 10: WinRM is enabled and configured by using the <code>winrm quickconfig</code> command. • A Windows user account with WinRM permissions exists. This account should have access to the application and be a member of the virtual machine's local Administrators group. • Access to the virtual machine file system is enabled. For instructions, see “Enabling access to data” on page 146
Linux virtual machines	<ul style="list-style-type: none"> • Access to the virtual machines through SSH is enabled. • Access to the virtual machine file system is enabled. For instructions, see “Enabling access to data” on page 146
Nutanix ESXi clusters	<p>VMware Tools and NGT are installed on the client virtual machine.</p> <p>For detailed information about installing VMware Tools, see VMware documentation. For detailed information about installing NGT, see Nutanix documentation.</p>

Application-specific prerequisites

SQL Server	<ul style="list-style-type: none"> Access should be enabled on all virtual machines where the SQL Server failover cluster and SQL Server Always On Availability Group instance resides. <i>For SQL Server Always On Availability Group:</i> An availability group is created using automatic seeding.
Oracle	<ul style="list-style-type: none"> The OS user must have sudo privileges and the NOPASSWD option set.

Consideration

For an Oracle application: When an operating system is used to authenticate Oracle database users, the Oracle database can be accessed with the OS user credentials, which allows you to skip the procedure of providing access to application data. To enable such authentication mode, contact the Oracle database administrator.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, select the applications that you want to protect.
2. Click  **Configuration**. The Configuration dialog box opens.
3. On the Credentials tab, depending on the credentials that you want to use, do one of the following:
 - If you want to use virtual machine credentials, click **Save**.
 - If you want to use the application-specific credentials, do the following:
 - a. Disable the **Use VM credentials with access to the application** switch.
 - b. Enter credentials for a user account with required permissions and access to the applications. Make sure the following requirements are met:
 - *For applications running on Windows virtual machines:* The specified account must be a member of the virtual machine's local Administrators group.

- *For SQL Server:* The specified account must have the sysadmin role on the SQL Server application instance. The SQL Server account that connects by using SQL Server Authentication is not supported.
- *For Exchange Server:* The specified account must be a member of the View-Only Organization Management role group, and it must have the Databases and Disaster Recovery roles assigned.

c. Click **Save**.

A new process of application discovery is started with the modified credentials for all virtual machines that have these credentials assigned. After this is done, the status of your applications should be  and you can continue with protecting application data as described in “Backing up applications” on [page 222](#).

You can later unassign the credentials from a virtual machine by clicking **Unassign** or delete the virtual machine credentials that you do not need anymore by clicking  **Delete**.

 **Important** You can unassign or delete credentials from a virtual machine only if the discovered applications running on it do not have assigned policies or available restore points. Therefore, before unassigning or deleting credentials, make sure to unassign policies or expire restore point tiers.

Planning application protection

Before performing an application backup, get familiar with prerequisites, limitations, considerations, and recommendations to determine if your environment is ready for application data protection. As part of planning application protection, you can also do the following:

- Set several application-specific options to make sure the actions specified by these options are performed automatically as part of the application backup. For details, see “[Setting application-specific options](#)” on [page 218](#).
- Exclude one or more databases from the SQL Server application backup. For details, see “[Excluding databases from the SQL Server application backup](#)” on [page 221](#).

Prerequisites

- *For vSphere environments:* VMware Tools of the latest version must be installed on virtual machines on which the applications you want to protect are running.
- *For Linux virtual machines with the installed NGT:* The following scripts must be available on the system, must be owned by root, and must have permissions set to 0700: `/usr/local/sbin/pre_freeze` and `/usr/local/sbin/post_thaw`.
- *Only if using the SpinUp functionality.* Make sure all the prerequisites for migrating virtual machines and applications to cloud listed in “[SpinUp specifics](#)” on page 142 are fulfilled.
- *For servers:*
 - *For Windows:*
 - The VSS service must be enabled and running, and the VSS writer status must be stable.
 - WinRM must be enabled and configured by using the `winrm quickconfig` command.
 - *For Linux:* Access to the server through SSH must be enabled.

Application-specific prerequisites

Application type	Prerequisites
SQL Server	<ul style="list-style-type: none"> • Databases must reside on the local disks in the Nutanix environment. • <i>Only if NGT is installed on a virtual machine with an SQL Server application.</i> Taking application-consistent snapshots must be disabled. For details, see Nutanix documentation. • <i>For restoring an SQL Server database to a point in time:</i> The database must be online and must be set to the full or bulk-logged recovery model during the backup. • <i>For restoring databases that are part of an Always On Availability Group:</i> Either all nodes in the Always On Availability Group must be protected by HYCU or only the node with the synchronized databases of the Always On Availability Group (must be online when being protected).

Application type	Prerequisites
	<p>In the latter case, the risk of data loss is increased if the node goes offline or the databases get out of sync.</p> <ul style="list-style-type: none"> • <i>For using a separate disk volume as backup storage for SQL Server temporary files:</i> Make sure that a dedicated disk of a sufficient size is allocated. The volume must be able to store temporary files that are generated between two backups of your SQL Server database. • <i>For SQL Server failover cluster:</i> <ul style="list-style-type: none"> ◦ All virtual machines where an SQL Server failover cluster resides must be discovered by HYCU. ◦ Policies must be assigned to all virtual machines on which the application instance is running.
Active Directory	<ul style="list-style-type: none"> • NGT must be installed and enabled on the client virtual machine. For details on how to do this, see Nutanix documentation. • No volume groups must be attached to the client virtual machine.
Exchange Server	<ul style="list-style-type: none"> • NGT must be installed and enabled on the client virtual machine. For details on how to do this, see Nutanix documentation. • No volume groups must be attached to the client virtual machine. • All databases must be mounted. • The Active Directory application must be protected. <p>Because Exchange Server stores all configuration information in Active Directory, make sure that you also back up your Active Directory application so that you can retrieve the information about the configuration if required. For example, if an entire database is deleted by accident and you want to restore it, you must first restore the Active Directory application, and then you can restore this database by performing the Exchange Server restore. However, if only the contents of the database are deleted, you must restore only the Exchange Server application.</p>

Application type	Prerequisites
Exchange Server	<ul style="list-style-type: none"> For Exchange Server 2019 and 2016 with the November 2023 security update installed: You must run a custom PowerShell session configuration. Run the following PowerShell commands on the virtual machine: <pre data-bbox="536 586 1117 662">New-PSSessionConfigurationFile -Path .\HycuJEA.pssc -RunAsVirtualAccount</pre> <pre data-bbox="536 714 1229 790">Register-PSSessionConfiguration -Path .\\HycuJEA.pssc -Name hycu.powershell -Force</pre>
Oracle	<ul style="list-style-type: none"> The SSH service must be enabled on the Oracle server and must be listening on port 22 for incoming connections. The Oracle database user must have the SYSDBA privilege. The database must be running in ARCHIVELOG mode. Tablespaces must be online. Additional disk space must be provided for temporary files created between two database backups. For optimal restore performance, separate disks must be specified for the temporary and database files.
SAP HANA	<ul style="list-style-type: none"> SAP HANA savepoints are enabled. For multiple volume groups: All data volumes and log volumes must belong to the same volume group. For distributed (multi-host) environments: <ul style="list-style-type: none"> All virtual machines where SAP HANA resides must be discovered by HYCU. Policies must be assigned to all virtual machines on which the application instance is running.

Limitations

- Backing up multiple application types running on a virtual machine is not supported.
- Backing up multiple instances of the same application type running on a virtual machine is supported only for SQL Server and Oracle.

- Backing up applications running on virtual machines in ROBO environments is not supported.
- For Nutanix ESXi clusters:* If you enabled the Backup from replica policy option, backing up virtual machines that have disks on different containers is not supported.
- For vSphere environments:* Protecting applications running on virtual machines that have vSphere Fault Tolerance enabled is not supported.

Application-specific limitations

Application type	Limitations
SQL Server	<ul style="list-style-type: none"> The <code>tempdb</code> SQL Server system database is excluded from all backups. Only a full backup of the <code>master</code>, <code>model</code>, and <code>msdb</code> SQL Server system databases is supported. You can restore an SQL Server system database only as a whole instance. A point-in-time restore of the <code>master</code>, <code>model</code>, <code>msdb</code>, or <code>tempdb</code> SQL Server system database is not possible. Backing up a database that is set to single-user mode is not possible if it is already in use. <i>For Always On Basic Availability Groups:</i> No backups on a secondary replica are possible.
Active Directory	<ul style="list-style-type: none"> Backing up the applications running on the volume groups or on the virtual machines with the attached volume groups is not supported. <i>For Nutanix clusters:</i> Protecting applications that are running on virtual machines with IDE disks is not possible. <i>For AWS GovCloud (US) and Azure Government environments:</i> Protecting Active Directory applications is not supported.
Exchange Server	<ul style="list-style-type: none"> Backing up the applications running on the volume groups or on the virtual machines with the attached volume groups is not supported. <i>For Nutanix clusters:</i> Protecting applications that are running on virtual machines with IDE disks is not

Application type	Limitations
	<p>possible.</p> <ul style="list-style-type: none"> • <i>For AWS GovCloud (US) and Azure Government environments:</i> Protecting Exchange Server applications is not supported.
Oracle	<ul style="list-style-type: none"> • Backing up Oracle Real Application Clusters (RAC) databases is not supported. Consequently, assigning policies to such databases is not possible.

Considerations

- *For Nutanix ESXi clusters:* If a full backup snapshot is missing on a Nutanix cluster (for example, because the HYCU protection domain is deleted from Prism), the next backup will be a full backup.
- *For protection domains configured with NearSync:* Although snapshots in a protection domain are created in a 1-15 minute interval, HYCU uses only the snapshots that are created on an hourly basis for backing up and restoring from snapshots. This applies to the following environments:
 - Nutanix ESXi clusters
 - Nutanix clusters when using the Backup from replica option
- *For SQL Server:*
 - *Only if you have upgraded your SQL Server to a newer version.* HYCU recognizes the upgraded application as a new application and at the same time changes the status of the old one to Protected deleted. Therefore, to ensure data protection for the upgraded application, do the following:
 1. Assign credentials to the upgraded application to enable HYCU to access it. For details, see “[Enabling access to application data](#)” on [page 209](#).
 2. Assign a policy to the upgraded application to protect it. For details, see “[Backing up applications](#)” on [page 222](#).
 - Backing up transaction logs of an SQL Server database with the AUTO CLOSE option set to TRUE may fail if the database has the RECOVERING status.
- *For vSphere environments:* If you use HotAdd, keep in mind the following:
 - When restoring from a snapshot, using HotAdd is supported if all the prerequisites are met. For details about HotAdd prerequisites, see

VMware documentation.

- Using HotAdd is supported only for snapshots that were created automatically by HYCU, not for snapshots that were created by using the Recreate Snapshot option.

Recommendation

For SQL Server and Oracle: It is recommended to use a dedicated disk of a sufficient size for storing temporary files generated during a backup.

Otherwise, this data will be stored on the biggest disk or an operating system disk volume which may affect the restore performance.

Setting application-specific options

HYCU enables you to set several application-specific options before you start backing up your applications. By doing so, you make sure the actions specified by these options are performed automatically as part of the application backup.

Accessing the Options tab

To access the Options tab, follow these steps:

1. In the navigation pane, click  **Applications**.
2. From the list of discovered applications, select the one for which you want to specify the application-specific option, and then click  **Configuration**.
3. Click the **Options** tab.

The following application-specific options are available:

SQL Server	<ul style="list-style-type: none"> • Back up and truncate SQL transaction logs (<i>enabled by default</i>) <p>Use the switch if you want your SQL Server transaction logs to be backed up and truncated in the SQL Server database automatically as part of the HYCU application backup. In this case, you can use HYCU to recover the SQL Server database.</p> <p>If disabled, HYCU does not back up and truncate the SQL Server transaction logs. In this case, to recover the SQL Server database, you should apply the transaction logs manually after restoring data.</p> • Enter path for temporary transaction log backup and
------------	---

	<p>metadata files (<i>optional</i>)</p> <p>If specified, the backup copies of the SQL Server temporary files (transaction logs and metadata files) are stored to this location. Otherwise, these backup copies are stored to the .hycu folder on the root of the disk with the largest amount of free space.</p> <p>Note For better restore performance, it is recommended to use a dedicated disk for storing backup copies of temporary files.</p> <ul style="list-style-type: none"> Optimized SQL Server HADR protection <p><i>This option is available for Windows virtual machines hosting SQL Server databases that are part of an Always On Availability Group, and is not available for such virtual machines residing in an AWS GovCloud (US) or Azure Government environment.</i></p> <p>Enable this option if you want to run backups only on the secondary replica with the highest backup priority. If only the primary replica is available, the backups are run on the primary replica.</p> <p>Important If you plan to enable the Optimized SQL Server HADR protection option, take into account the following:</p> <ul style="list-style-type: none"> The primary replica should not share disks with the secondary replicas or the local databases of the SQL Server instance. The secondary replica with the highest backup priority should not share disks with the secondary replicas with lower backup priority or the local databases of the SQL Server instance. <i>For servers hosting SQL Server databases that are part of an Always On Availability Group:</i> Storage exclusion is performed at the volume level, and not at the disk level.
Exchange Server	<ul style="list-style-type: none"> Priority for Exchange Server restore requests <p>Specifies the priority in which the restore requests for a mailbox restore are processed on the Exchange Server:</p>

	<p>Lowest, Lower, Low, Normal (the default value), High, Higher, Highest, Emergency.</p> <ul style="list-style-type: none"> Optimized Exchange Server DAG protection <p><i>Available for Windows servers running Exchange mailbox servers that are members of a database availability group (DAG). Enable this option if you want to back up only the volumes hosting the passive database copies with the highest activation preference number (including the system volume). Only the DAG members with the assigned policies are taken into account when searching for the database copies with the highest activation preference number. If no passive database copies are available, active database copies will be backed up.</i></p> <p>! Important Optimized Exchange Server DAG protection is effective only if separate databases are stored on separate volumes.</p>
Oracle	<ul style="list-style-type: none"> Back up and truncate Oracle archive logs (<i>enabled by default</i>) <p>Use the switch if you want your Oracle archive logs to be backed up and truncated in the Oracle database automatically as part of the HYCU application backup. In this case, you can use HYCU to recover the Oracle database.</p> <p>If disabled, HYCU does not back up and truncate the Oracle archive logs. In this case, to recover the Oracle database, you should apply the transaction logs manually after restoring data.</p> <ul style="list-style-type: none"> Enter path for temporary Oracle files (<i>optional</i>) <p>If specified, the backup copies the temporary Oracle files will be stored to this location.</p> <p>! Note For better restore performance, it is recommended to use a dedicated disk for storing backup copies of temporary files.</p>

Excluding databases from the SQL Server application backup

By default, all databases belonging to SQL Server applications are backed up during the application backup. However, if you want specific databases to be excluded from the backup, HYCU enables you to select these databases before the application backup is performed. You can exclude databases on standalone SQL Server instances, SQL Server Always On Availability Group instances, and SQL Server failover cluster instances.

Considerations

- *Only if you exclude a database that is part of an Always On Availability Group from the backup.* The database is excluded from all SQL Server instances participating in the group.
- *Only if you exclude a database on a failover cluster from the backup.* The database is excluded from the active node and all passive nodes.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, select the SQL Server application whose databases you want to exclude from the backup, and then click  **Configuration**.
2. Click the **Exclude databases** tab.
3. From the list of all databases, select the ones that you want to exclude from the backup.

 **Important** The master, model, and msdb SQL Server system databases cannot be excluded from the backup and are grayed out.

4. Click **Save**.

You can later make changes to the selection of the excluded databases.

Backing up applications

An application-aware backup allows a consistent backup of discovered applications.

Accessing the Applications panel

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

Limitations

- A tape target cannot be used for storing application data.
- *For vSphere environments:* If you enable the Archiving option and select Snapshot as the backup target type in your policy, such a policy cannot be assigned to the application.

Considerations

- If during virtual machine synchronization, a virtual machine cannot be found in a source environment, the status of this virtual machine and any discovered applications running on it is set to PENDING_REMOVAL. The policy is still assigned to the virtual machine and the applications, but you cannot perform any data protection actions (they are grayed out in HYCU). Depending on whether this virtual machine is found in the source environment during the time interval of two automatic virtual machine synchronization processes, the following happens:
 - *The virtual machine is found in the source environment:* Its status and the status of the applications running on it is changed to Protected.
 - *The virtual machine is not found in the source environment:* If the virtual machine still has at least one valid restore point available, its status and the status of the applications running on it is changed to Protected deleted. This means that the virtual machine that is deleted from the source is still considered protected and is not removed from HYCU.

Procedure

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

 **Tip** To narrow down the list of all displayed applications, you can use the filtering options described in “[Filtering and sorting data](#)” on [page 317](#).

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

 **Note** When you assign the policy to the selected applications, the same policy is also assigned to the virtual machines on which they are running. If these virtual machines already have an assigned policy, the policy assigned to the applications takes precedence over the policy assigned to the virtual machines and is automatically assigned to the virtual machines.

The backup is scheduled according to the values that you defined for your policy. If required, you can also perform a manual backup of any application at any time. For details, see “[Performing a manual backup](#)” on page 333.

Restoring whole applications

With HYCU, you can restore a whole application to its original or a new location by restoring the virtual machine and attached volume groups on which the application is running.

 **Note** *For Active Directory:* HYCU does not perform an authoritative restore.

Prerequisites

- *For Nutanix ESXi clusters and vSphere environments:* You must have the required restore privileges assigned. For details, see “[Assigning privileges to a vSphere user](#)” on page 456.
- *For applications with the Protected deleted status whose backups are stored on the imported targets:* Such applications must be discovered. For details, see “[Enabling access to application data](#)” on page 209.
- *For servers:* At least one Nutanix cluster, vCenter Server, AWS GovCloud (US) region, or Azure Government subscription must be added to HYCU to provide a storage container for storing the restore data. For details on how to add sources to HYCU, see “[Adding sources](#)” on page 53.

Considerations

- A restore is performed from the snapshot only if you are restoring to the same source (the source where the original virtual machine was running). If you are restoring to a different source, depending on the tier that you select for the restore, the following will happen:
 - If you select Snapshot, the restore will fail.
 - If you select Automatic, the restore will be performed from the target if there is an available target. Otherwise, it will fail.
- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- You cannot perform a restore of an application whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up` configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see [“Customizing HYCU configuration settings” on page 523](#).
- *For SQL Server:* If you disabled the Back up and truncate SQL transaction logs option, you should apply transaction logs manually after restoring data to recover the SQL Server database.
- *For Oracle:* If you disabled the Back up and truncate Oracle archive logs option, you should apply archive logs manually after restoring data to recover the Oracle database.
- *For SQL Server failover cluster and SAP HANA multi-host environments:* Make sure to select the virtual machine with the latest backup of attached volume groups. To identify the appropriate virtual machine, you can use the Jobs panel. For details, see [“Managing HYCU jobs” on page 300](#).
- *For AWS GovCloud (US) virtual machines with encrypted volumes:* Depending on whether you are restoring such a virtual machine to the same or a different region, the following applies:
 - To the same region: The restored volumes will be encrypted with the same KMS key as the original ones.
 - To a different region: The restored volumes will be encrypted with the default KMS managed key for EBS encryption.

- Only if you are cloning a server running an application for which the Optimized Exchange Server DAG protection option was enabled during the backup. After cloning the server, the volumes that were excluded during the backup will show up as RAW partitions in the Disk Management console. To recover data from the RAW partitions, do the following:
 - Reformat the RAW partitions to NTFS or REFS.
 - Reseed the database copies.

Restore options

You can select between the following restore options:

I want to...	Restore option	Description
Restore a virtual machine to the original location.	Restore VM	<p>Enables you to restore an application by restoring the virtual machine on which the application is running to the same source. Select this option if you want to replace the original virtual machine with the restored one. For instructions, see “Restoring a virtual machine” below.</p> <p>Important You cannot restore applications running on servers by using this option.</p>
Restore a virtual machine by cloning it to a new location.	Clone VM	<p>Enables you to restore an application by creating a clone of the virtual machine on which the application is running to the same or a different source. Select this option if you want to keep the original virtual machine. For instructions, see “Cloning a virtual machine” on page 238.</p>

Restoring a virtual machine

You can restore an application by restoring the virtual machine on which the application is running to its original or a new location on the same source. In this case, the original virtual machine will be overwritten.

Caution When you are restoring the application to the original location, the restored data overrides the data in the original location. To avoid data

loss, make sure that you back up the potentially unprotected data—the data that appeared between the last successful backup and the restore. To start a manual backup, see “[Performing a manual backup](#)” on page 333.

For details on how to restore a virtual machine, depending on your data protection environment, see one of the following sections:

- “[Restoring a virtual machine to a Nutanix cluster or a vSphere environment](#)” below
- “[Restoring a virtual machine to an AWS GovCloud \(US\) environment](#)” on page 230
- “[Restoring a virtual machine to an Azure Government environment](#)” on page 234

Restoring a virtual machine to a Nutanix cluster or a vSphere environment

Considerations

- *Only if volume groups are attached to the virtual machine that you are restoring.* You can choose to restore the volume groups together with the virtual machine if they were attached to it at backup time. In this case, the original volume groups are deleted and the restored ones are automatically attached to the restored virtual machine as well as all other virtual machines to which they were attached at backup time.
- The restored virtual machine retains the original MAC address.
- *Only if you plan to restore a vSphere virtual machine.* Depending on how you plan to restore data, consider the following:
 - *From a target:* The original virtual machine and all its snapshots will be deleted as part of the restore process.
 - *From a snapshot:* The entire virtual machine will be reverted to the selected snapshot and any excluded or included disk configuration will be ignored.
- *Only if you plan to restore vSphere virtual machine data to the original storage container.* If the storage container is mounted to several hosts and the original host is powered off or in maintenance mode at restore time, data will be restored to the same storage container on a different host.
- *Only if you plan to restore a vSphere virtual machine to a datacenter that was not added to HYCU.* The protection status of such a virtual machine will be Protected deleted after the restore.

- *Only if you plan to restore a vSphere virtual machine from a datacenter that was removed from HYCU.* After you remove the datacenter, the protection status of the virtual machine changes from Protected to Protected deleted. When restoring such a virtual machine, consider the following:
 - If restoring to a datacenter that is added to HYCU, the protection status of the virtual machine changes back to Protected.
 - If restoring to any datacenter that is not added to HYCU, the protection status of the virtual machine stays Protected deleted.
- *Only if you plan to restore a virtual machine running on a Nutanix ESXi cluster.* If Snapshot is selected as the backup target type in your policy, the NVRAM file will not be restored.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

 **Note** The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open the Detail view.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

 **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.
3. Click  **Restore**.
4. Select **Restore whole server**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. In the General section, do the following:
 - a. From the Storage container drop-down menu, select where you want to restore the virtual machine. By default, the original storage container is selected.

 **Note** If you decide to restore the virtual machine to another storage container, keep in mind the following:

- Restore from the Snapshot tier cannot be performed to another storage container.
- If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.

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

- **Automatic:** Ensures the fastest restore to the latest state.
- **Backup**
- **Copy**
- **Archive**
- **Snapshot**

c. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
- In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

 **Note** The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.

d. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore. The original virtual machine will be deleted automatically.

 **Important** Only if you are restoring a vSphere virtual machine to a vSphere environment and you have disabled the Power virtual machine on switch. When you try to power on the virtual machine and you are prompted to answer whether the virtual machine has been moved or

copied, make sure to answer **I Moved It**.

- e. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
- f. *For volume groups attached to the virtual machine:* Use the **Restore volume groups** switch if you want to restore also the volume groups that are attached to the virtual machine.

7. In the Network section, review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

Depending on whether the original networks are available, proceed as follows:

- If the original networks are available, you can leave the default values and restore the virtual machine with the original network settings, or you can modify the network settings.
- If the original networks are not available, you must modify the network settings.

Modifying network settings

Original networks are...	Instructions
Available	<p>You can do the following:</p> <ul style="list-style-type: none"> • Add a new network adapter by clicking  New, and then selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting it, and then clicking  Edit and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking  Delete.
Unavailable	You can do the following:

Original networks are...	Instructions
	<ul style="list-style-type: none"> • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking  Edit and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking  Delete. • Add a new network adapter by clicking  New, and the selecting the preferred network.

 **Note** You can restore the virtual machine without a network adapter. Make sure to configure the network settings on the virtual machine afterward.

8. Click **Restore**.

 **Note** *For Nutanix ESXi clusters:* Because the minimum RAM required for restoring a virtual machine is 256 MiB, any virtual machine with less RAM is automatically set to 256 MiB during the restore.

During the restore, the original application instance is offline and not accessible.

Restoring a virtual machine to an AWS GovCloud (US) environment

Considerations

- Make sure that the virtual machine you are restoring is not deleted from AWS GovCloud (US). If you delete a virtual machine from AWS GovCloud (US), you cannot restore it even if it still has a valid restore point available in HYCU (that is, even if its status is Protected deleted).
- When restoring a virtual machine, the original virtual machine disks are deleted and replaced with the restored ones.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

 **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click  **Restore**.
4. Select **Restore whole server**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. The following information is displayed and preselected:
 - The AWS GovCloud (US) account to which the virtual machine will be restored.
 - The account ID of the AWS GovCloud (US) account to which the virtual machine will be restored.
 - The region to which the virtual machine will be restored.
 - *Only if found by HYCU*. The key pair name for connection to the restored virtual machine.

 **Important** If HYCU does not find the name of your key pair and you want to use it, you can select it from the Key pair name drop-down menu.

 - The availability zone to which the virtual machine will be restored.
7. Click **Next**.
8. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**

- **Archive**
- **Snapshot**

9. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU threads field, enter the number of CPUs for the restored virtual machine multiplied by the number of cores per CPU and the number of threads per core.
- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- From the Virtual machine type drop-down menu, select the virtual machine type.

 **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

10. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.

11. Under Network interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the Virtual Private Cloud (VPC) to which the original virtual machine belongs. If required, you can also modify network settings.

Modifying network settings

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

 **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same Virtual Private Cloud (VPC). The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
 - The Virtual Private Cloud (VPC) to which the network interface will be added is displayed and preselected.
 - From the Subnets drop-down menu, select the subnet to which the network interface should be assigned.
 - From the Security groups drop-down menu, select one or more security groups that will be associated with the network interface. If you want to select all the available security groups, select **Select all**.
 - In the Public address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Auto-assign	An automatically allocated public IP address will be assigned to the network interface on the restored virtual machine.
Elastic IP (Reserved)	An elastic public IP address that you reserved in AWS GovCloud (US) will be assigned to the network interface on the restored virtual machine.
Elastic IP (New)	An elastic public IP address will be assigned to the network interface on the restored virtual machine.

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

Option	Description
Auto-assign	An automatically allocated private IP address will be assigned to the network interface on the restored virtual machine.
Custom	A private IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

12. Only if the virtual machine operating system has not been discovered yet. Select the virtual machine operating system:

- Linux**
- Windows**

13. Under Operating system license, select one of the following options:

OS license option	Select this option if you want to...
Keep existing license	Keep the existing OS license on the restored virtual machine. Important Make sure that the existing license is applicable also in AWS.
<i>Available only for the Windows Server OS. Replace existing license with AWS license</i>	Replace the existing OS license with an AWS license on the restored virtual machine.

14. Click **Restore**.

Restoring a virtual machine to an Azure Government environment

Consideration

If you want the restored virtual machine to have the same static IP address as the original virtual machine, do one of the following:

- Before the restore, in Azure Government, disassociate the IP address from the original virtual machine, and then select this IP address for the network interface during the restore in HYCU.
- During the restore, select a different IP address for the network interface. After the restore, in Azure Government, assign the preferred IP address to the restored virtual machine.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

 **Note** The Detail view appears only if you click an application.

Selecting the check box before the name of the application will not open the Detail view.

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

 **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click  **Restore**.
4. Select **Restore whole server**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. From the Location drop-down menu, select the geographic region for the restored virtual machine.
7. From the Availability zone drop-down menu, select the zone for the restored virtual machine.

 **Note** The selected geographic region and the size of the virtual machine determine to which zones you can restore data. If you do not want to restore data to any zone, select **None**.

8. Click **Next**.
9. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**

10. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU cores field, enter the number of virtual CPUs for the restored virtual machine.
- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- From the Virtual machine type drop-down menu, select the virtual machine type.

 **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

11. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
12. Under Network interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the subscription to which the original virtual machine belongs. If required, you can also modify network settings.

Modifying network settings

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

 **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:

- a. *Only if you are adding a network interface.* From the Network drop-down menu, select the network for the network interface.

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

- b. Select the subnet to which the network interface should be assigned.
- c. In the Public IP address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	A static IP address will be assigned to the network interface on the restored virtual machine.
Existing	A preferred public IP address resource that you have created in Azure Government will be assigned to the network interface on the restored virtual machine.

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

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the restored virtual machine.

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

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

13. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
 - **Linux**
 - **Windows**
14. Click **Restore**.

Cloning a virtual machine

You can restore an application by creating a clone of the virtual machine on which the application is running to its original or a new location on the same or a different source. In this case, the original virtual machine will not be overwritten.

For details on how to clone a virtual machine, depending on your data protection environment, see one of the following sections:

- “[Cloning a virtual machine to a Nutanix cluster or a vSphere environment](#)” below
- “[Cloning a virtual machine to an AWS GovCloud \(US\) environment](#)” on [page 244](#)
- “[Cloning a virtual machine to an Azure Government environment](#)” on [page 248](#)

① Important After you clone the virtual machine, make sure that everything works as expected by going through the considerations and the recommendations listed in “[After cloning a virtual machine](#)” on [page 185](#).

Cloning a virtual machine to a Nutanix cluster or a vSphere environment

Prerequisites

- *For virtual machines that you plan to clone to a new location:* A Nutanix cluster or a vCenter Server for a vSphere environment to which you plan to clone the virtual machine must be added to HYCU. For details on how to do this, see “[Adding a Nutanix cluster](#)” on [page 54](#) or “[Adding a vCenter Server](#)” on [page 57](#).
- *For Linux servers:* In the `/etc/fstab` system configuration file of the server, UUIDs (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`) must be used instead of device names for file system device identification.

Limitation

For vSphere environments: Attaching the ISO image to the restored virtual machine is not supported.

Considerations

- *Only if volume groups are attached to the virtual machine that you are cloning.* You can choose to restore the volume groups together with the virtual machine if they were attached to it at backup time. In this case, the original volume groups are kept alongside of the restored ones. If the volume groups are also attached to other virtual machines, the following applies (depending on how they are attached to the virtual machines):
 - Directly: Volume groups are automatically attached only to the cloned virtual machine.
 - By using iSCSI: Volume groups are automatically attached to all virtual machines to which they were attached at backup time.
- *For restoring a virtual machine running on a Nutanix AHV cluster to a Nutanix ESXi cluster:* If virtual machine disks are attached to the PCI bus, the bus type will be automatically changed to SCSI after the restore. Because of this configuration change, the restore finishes with a warning.
- *For Linux virtual machines running on a Nutanix ESXi cluster:* If after restoring a virtual machine that was created through the vSphere (Web) Client, the virtual machine does not boot, follow the steps described in [“After restoring a virtual machine to a Nutanix ESXi cluster” on page 532](#).
- After you restore a virtual machine, it might happen that the order of virtual disks differs from the one on the original virtual machine if you performed the restore:
 - From a Nutanix AHV cluster to a Nutanix ESXi cluster or a vSphere environment
 - From a Nutanix ESXi to another Nutanix ESXi cluster
 - From a vSphere environment to a Nutanix ESXi cluster
 In this case, make the necessary adjustments, including the selection of the correct boot disk.
- *Only if you plan to restore vSphere virtual machine data to the original storage container.* If the storage container is mounted to several hosts and the original host is powered off or in maintenance mode at restore time, data will be restored to the same storage container on a different host.

- *Only if ownership is set for the virtual machine.* The same owner is automatically assigned to the restored virtual machine.
- *Only if you plan to restore a virtual machine running on a Nutanix ESXi cluster.* If Snapshot is selected as the backup target type in your policy, the NVRAM file will not be restored.
- *Only if the original virtual machine resides on a source other than a vSphere environment.* Make sure to modify the virtual machine configuration by specifying the appropriate guest operating system.

Recommendation

For Linux virtual machines: It is recommended that the use of persistent network device names based on MAC addresses is disabled. Otherwise, you will have to configure the network manually. For details on how to disable the use of persistent network device names, see your Linux distribution documentation.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

 **Note** The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open the Detail view.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

 **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.
3. Click  **Restore**.
4. Select **Restore whole server**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.

7. In the General section, do the following:

- a. From the Storage container drop-down menu, select the storage container where you want to restore the virtual machine.

Note By default, the original storage container is selected. If you decide to restore the virtual machine to another storage container, keep in mind the following:

- Restore from the Snapshot tier cannot be performed to another storage container.
- If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.
- If the selected storage container is on a different source, additional prerequisites apply. For details, see [“After restoring a virtual machine to a different source” on page 530](#).

- b. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic:** Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**

- c. In the New VM name field, specify a new name for the virtual machine.
- d. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
- In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

Note The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.

e. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore.

➊ **Important** Make sure to consider the following:

- This option is disabled for virtual machines that have volume groups attached by using iSCSI. For details on what needs to be done before turning on the restored virtual machine, see “[After cloning a virtual machine](#)” on page 185.
- *Only if you are cloning a virtual machine from a Nutanix cluster or a vSphere environment.* If you turn the restored virtual machine on, the original virtual machine will be turned off automatically.
- *Only if you are cloning a vSphere virtual machine to a vSphere environment and you have disabled the Power virtual machine on switch.* When you try to power on the virtual machine and you are prompted to answer whether the virtual machine has been moved or copied, make sure to answer **I Copied It**.

f. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.

g. *For volume groups attached to the virtual machine:* Use the **Clone volume groups** switch if you want to restore also the volume groups that are attached to the virtual machine.

8. In the Network section, do the following:

- a. Review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

Depending on whether the original networks are available, proceed as follows:

- If the original networks are available, you can leave the default values and clone the virtual machine with the original network settings, or you can modify the network settings.
- If the original networks are not available, you must modify the network settings.

Modifying network settings

Original networks are...	Instructions
Available	<p>You can do the following:</p> <ul style="list-style-type: none"> • Add a new network adapter by clicking  New and selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting the virtual adapter, and then clicking  Edit and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking  Delete.
Unavailable	<p>You can do the following:</p> <ul style="list-style-type: none"> • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking  Edit and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking  Delete. • Add a new network adapter by clicking  New, and then selecting the preferred network.

 **Note** You can clone the virtual machine without a network adapter. Make sure to configure the network settings on the virtual machine afterward.

b. *Only if you are restoring the virtual machine to a different Nutanix cluster or vSphere environment. Use the **Keep original MAC address** switch if you want the restored virtual machine to keep the original MAC address. Keep in mind that this is applicable only if at least one network adapter has a MAC address assigned.*

9. Click **Restore**.

During the restore, the original application instance is offline and not accessible.

There are some considerations that you should be aware of after cloning a virtual machine. For details, see [“After cloning a virtual machine” on page 185](#).

Cloning a virtual machine to an AWS GovCloud (US) environment

Prerequisite

For virtual machines that you plan to restore to a new location: The AWS GovCloud (US) region to which you plan to restore the virtual machine must be added to HYCU. For instructions, see [“Adding an AWS GovCloud \(US\) region” on page 59](#).

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for restoring data to a new location.
- *For virtual machines that have BitLocker volumes encrypted with TPM-based keys:* Restoring such volumes is not supported.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

 **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click  **Restore**.
4. Select **Restore whole server**, and then click **Next**.

5. Select **Clone VM**, and then click **Next**.
6. From the Destination source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. From the AWS GovCloud (US) account, select the account to which the virtual machine will be restored.

The following information is displayed and preselected:

- The account ID of the AWS GovCloud (US) account to which the virtual machine will be restored.
- The region to which the virtual machine will be restored.

8. *Optional.* From the Key pair name drop-down menu, select the key pair name that you want to use for connection to the restored virtual machine.

! Important *For Windows virtual machines:* The key pair name that you select can be used only if the EC2Config or EC2Launch service was configured on the original virtual machine or if you configure it later on the restored virtual machine.

9. From the Availability zone drop-down menu, select the availability zone to which the virtual machine will be restored.
10. Click **Next**.
11. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic:** Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
12. In the New VM name field, specify a name for the restored virtual machine.
13. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.
If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:
 - In the vCPU threads field, enter the number of CPUs for the restored virtual machine multiplied by the number of cores per CPU and the number of threads per core.

- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- From the Virtual machine type drop-down menu, select the virtual machine type.

 **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

14. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
15. Under Network interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the Virtual Private Cloud (VPC) to which the original virtual machine belongs. If required, you can also modify network settings.

Modifying network settings

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

 **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same Virtual Private Cloud (VPC). The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
 - a. The Virtual Private Cloud (VPC) to which the network interface will be added is displayed and preselected.
 - b. From the Subnets drop-down menu, select the subnet to which the network interface should be assigned.
 - c. From the Security groups drop-down menu, select one or more security groups that will be associated with the network interface. If you want to select all the available security groups, select **Select all**.

d. In the Public address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Auto-assign	An automatically allocated public IP address will be assigned to the network interface on the restored virtual machine.
Elastic IP (Reserved)	An elastic public IP address that you reserved in AWS GovCloud (US) will be assigned to the network interface on the restored virtual machine.
Elastic IP (New)	An elastic public IP address will be assigned to the network interface on the restored virtual machine.

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

Option	Description
Auto-assign	An automatically allocated private IP address will be assigned to the network interface on the restored virtual machine.
Custom	A private IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

16. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:

- **Linux**
- **Windows**

17. Under Operating system license, select one of the following options:

OS license option	Select this option if you want to...
Keep existing license	Keep the existing OS license on the restored virtual machine. ① Important Make sure that the existing license is applicable also in AWS.
<i>Available only for the Windows Server OS. Replace existing license with AWS license</i>	Replace the existing OS license with an AWS license on the restored virtual machine.

18. Click **Restore**.

During the restore, the original application instance is offline and not accessible.

Cloning a virtual machine to an Azure Government environment

Prerequisites

- *For virtual machines that you plan to restore to a new location:* The Azure Government subscription to which you plan to restore the virtual machine must be added to HYCU. For details on how to do this, see “[Adding an Azure Government subscription](#)” on page 61.
- *For virtual machines that have Azure Disk Encryption enabled:* The key vault must be available on the location to which you are restoring the virtual machine.

Limitation

If a restore point contains only a Snapshot tier, you cannot use it for restoring data to a new location.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

Note The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open the Detail view.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

Important If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.
3. Click  **Restore**.
4. Select **Restore whole server**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. From the Service principal drop-down menu, select the service principal that has access to the required resources (the source from which and to which you are restoring the virtual machine).
8. From the Subscription drop-down menu, select the subscription for the restored virtual machine.
9. From the Resource group drop-down menu, select the resource group for the restored virtual machine.
10. From the Location drop-down menu, select the geographic region for the restored virtual machine.
11. From the Availability zone drop-down menu, select the zone for the restored virtual machine.

Note The selected geographic region and the size of the virtual machine determine to which zones you can restore data. If you do not want to restore data to any zone, select **None**.
12. Click **Next**.
13. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:

- **Automatic:** Ensures the fastest restore to the latest state.
- **Backup**
- **Copy**
- **Archive**
- **Snapshot**

14. In the New VM name field, specify a name for the restored virtual machine.
15. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- From the Virtual machine type drop-down menu, select the virtual machine type.

Note The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

16. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
17. Under Network interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the subscription to which the original virtual machine belongs. If required, you can also modify network settings.

Modifying network settings

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

Note When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on

the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:

- Only if you are adding a network interface.* From the Network drop-down menu, select the network for the network interface.

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

- Select the subnet to which the network interface should be assigned.
- In the Public IP address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	A static IP address will be assigned to the network interface on the restored virtual machine.
Existing	A preferred public IP address resource that you have created in Azure Government will be assigned to the network interface on the restored virtual machine.

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

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	The static IP address that you specify will be

Option	Description
	assigned to the network interface on the restored virtual machine.

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

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

- Only if the virtual machine operating system has not been discovered yet. Select the virtual machine operating system:
 - Linux**
 - Windows**
- Click **Restore**.

During the restore, the original application instance is offline and not accessible.

Restoring SQL Server databases

With HYCU, you can restore SQL Server databases to the original or a different SQL Server instance.

Prerequisites

- For point-in-time restore: The database recovery model must be set to full or bulk-logged.
- For restoring the whole SQL Server failover cluster instance: The SQL Server service must be stopped by using the Failover Cluster Manager. For details on how to do this, see SQL Server documentation.
- For improved restore performance, the startup type of the Microsoft iSCSI Initiator Service may not be set to **Disabled**.
- Only if you are restoring data that is stored in the archive storage tier on an Azure target. You must recreate a snapshot and use this snapshot for restoring data, or manually rehydrate data. For instructions on how to recreate a snapshot, see “[Recreating snapshots](#)” on page 340. For instructions on how to manually rehydrate data, see Azure documentation.

- *For AWS GovCloud (US) and Azure Government environments:* The virtual machine that hosts the SQL Server instance to which you plan to restore the SQL Server databases must be in the same virtual network as the HYCU backup controller.

Limitations

- Restoring SQL Server databases to another SQL Server application instance is supported only if you are restoring to the same or later version of the application.
- Databases that are part of an Always On Availability Group can be restored only to a primary node (from a secondary or primary node). However, keep in mind that in the case of an Always On Basic Availability Group, the databases can be restored only from a primary node.
- *Only if you plan to use the Archive restore point.* Performing the point-in-time restore is not supported.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- If you are restoring the databases to a different SQL Server instance, they will be renamed and copied to the default SQL Server location of the selected target.
- If a virtual machine is deleted from the source, but it still has at least one valid restore point available, it is considered protected. In this case, the status of the virtual machine or any discovered applications running on it is Protected deleted. When restoring application items of such an application, keep in mind that you cannot restore them to the original application instance.
- You cannot perform a restore of an application whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up` configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see [“Customizing HYCU configuration settings” on page 523](#).

- If any databases were excluded from the backup, you cannot select them for the restore.
- *For SQL Server failover clusters:*
 - The restore needs to be redirected to the active SQL Server failover cluster instance.
 - The Overwrite existing databases option can be enabled for a redirected restore only if the database location also exists on the target virtual machine.
- *For restoring an SQL Server database to a different SQL Server instance:* The Overwrite existing databases option should be enabled only when restoring to an SQL Server instance which is on a different server and has identical database paths.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application whose databases you want to restore to open the Detail view. The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open the Detail view.

 **Note** With the SQL Server Always On Availability Group, you can expand the application item to view the discovered Availability Groups.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

 **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the databases.
3. Click  **Restore**. The Restore MS SQL Server dialog box opens.

 **Note** If the Back up and truncate SQL transaction logs option was disabled during the backup, you are prompted that database recovery must be performed after the restore.
4. Select **Restore databases**, and then click **Next**.
5. From the Target instance drop-down menu, select where you want to restore the databases.

6. *For SQL Server Always On Availability Group:* From the Destination availability group drop-down menu, select one of the available Availability Groups to restore the databases to this group or leave the field empty to restore the databases to the SQL Server.
7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic:** Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
8. *Only if you want to restore the whole instance.* Select the **Whole instance** check box.
9. *Only if you want to restore databases.* Select the **Database name** check box to restore all databases or select the databases that you want to restore.
10. *Optional.* Specify a point in time to which you want to restore data. The databases will be restored to the state they were in at the specified time.

Note To perform a point-in-time restore, select a backup that was performed before the specified point in time so that the database instance can be brought to the appropriate state by applying the transaction log files from the next backup.
11. Click **Next**.
12. Use the **Leave databases in restoring state** switch if you want to leave the databases in the restoring state. By doing so, you can apply transaction logs to the databases after the restore and perform a manual point-in-time restore.
13. Use the **Overwrite existing databases** switch if you want to overwrite existing databases when performing a restore. In this case, the backups will be restored to their original location and all data will be overwritten. Keep in mind that if you are restoring the databases to another SQL Server instance, all the databases that have the same names (and not necessarily the contents) will be overwritten.

Otherwise, to restore data to a different location on the same or another SQL Server instance, specify a database prefix that will be given to the databases, a new database file location, and a new database log location.

ⓘ Important If you are restoring the whole instance, you can only overwrite existing databases. In this case, the Overwrite existing databases option is enabled by default and you cannot disable it.

14. Click **Restore**.
15. *Only if the Back up and truncate SQL transaction logs option was disabled during the backup* Recover the SQL Server databases by applying the transaction logs manually.
16. *Only if using SQL Server 2014 Always On Availability Groups*. Join the restored databases to an Always On Availability Group by using SQL Server Management Studio. For details on how to do this, see Microsoft documentation.

ⓘ Note After you join the restored databases to the Always On Availability Group, it is recommended to perform a new backup of your Always On Availability Group.

17. *Only if restoring the whole SQL Server failover cluster instance*. Start the SQL Server service and all other related services by using the Failover Cluster Manager. For details on how to do this, see SQL Server documentation.

Restoring Exchange Server databases, mailboxes, and public folders

With HYCU, you can restore Exchange Server databases, mailboxes, and public folders. When restoring Exchange Server databases, you can choose between restoring to the original mailbox server and, if the mailbox server is a member of a Database Availability Group (DAG), to another mailbox server inside the DAG. When restoring mailboxes and public folders, the recovery database can be restored to the original mailbox server or any other mailbox server that is part of your Exchange Server organization. From there, the actual restore is performed to any mailbox or public folder within the organization.

Prerequisites

- *For restoring mailboxes:*
 - The mailbox to which you are restoring data must exist on the server and be initialized.
 - *Only if the original mailbox to which you plan to restore data was deleted from the server.* You must create a new mailbox with the same or a different name, and make sure it is initialized (to do so, sign in to it with your Exchange client).
- *For restoring public folders:* The public folder must exist in the public folder mailbox. If it does not exist, recreate it manually with the same name it had at backup time.
- For improved restore performance, the startup type of the Microsoft iSCSI Initiator Service must not be set to Disabled.
- *Only if you are restoring data that is stored in the archive storage tier on an Azure target.* You must recreate a snapshot and use this snapshot for restoring data, or manually rehydrate data. For instructions on how to recreate a snapshot, see “[Recreating snapshots](#)” on page 340. For instructions on how to manually rehydrate data, see Azure documentation.

Limitations

- Restoring data to the `hycu` subfolder (the Restore to subfolder option) is currently not supported for public folders.
- *For Exchange Server 2019 or 2016 that has the November 2023 security update installed and was backed up with HYCU version 4.9.0-5310 or earlier:* Restoring the mailboxes and/or public folders from the corrupted backups is not possible. As a result, the Restore mailboxes and/or public folders option in the Restore MS Exchange Server dialog box is disabled.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- You cannot perform a restore of an application whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up`

configuration setting in the HYCU config.properties file to true. For details on how to customize the HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on page 523.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application whose application items you want to restore to open the Detail view.

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

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

 **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring application items.

3. Click  **Restore**. The Restore Exchange dialog box opens.

4. Select which application items you want to restore:

- **Restore databases**

- a. From the Recovery database server drop-down menu, select the server for restoring the data. When specifying a recovery database server, keep in mind that you can select it only if your mailbox server is a member of a DAG and you want to restore data to another mailbox server inside the DAG. Otherwise, you can restore only to the original mailbox server.

 **Important** For restoring a mailbox server that is a member of a DAG: Make sure to select the recovery database server on which the databases are currently active.

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

- **Automatic:** Ensures the fastest restore to the latest state.

- **Backup**

- **Copy**

- **Archive**

- **Snapshot**

- Select the **Database name** check box to restore all databases or select the databases that you want to restore.

- Use the **Enable restore to recovery database** switch if you want to enable restoring data to a recovery database. If enabled, provide a recovery database path. The default one is C:\ProgramData\Hycu.

- **Restore mailboxes and/or public folders**

- From the Recovery database server drop-down menu, select the mailbox server for restoring the data. You can select among the mailbox servers that are part of your Exchange Server organization.

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

- **Automatic:** Ensures the fastest restore to the latest state.

- **Backup**

- **Copy**

- **Archive**

- **Snapshot**

- From the list of mailboxes and/or public folders that are available for a restore, select the ones that you want to restore.

Tip If there are too many mailboxes and/or public folders to be displayed on one page, you can move between the pages by clicking **>** and **<**. You can also use **▼** to set the number of mailboxes and/or public folders to be displayed per page.

You can search for a mailbox and/or public folder by entering its name and then pressing **Enter** in the Search field.

- Enable the **Use non-default domain controller** switch if you want to use a domain controller other than the default one, and then, in the Domain controller field, enter the FQDN or IP address of the preferred domain controller.

- Click **Next**.

- Select where you want to restore data:

- **Original mailbox**
- **Alternate mailbox**, and then enter an alternate mailbox name.

g. Select the mode for restoring data:

- **Restore in place**
Enables you to restore data to the original location.
- **Restore to subfolder** (*not supported for public folders*)
Enables you to restore data to the hycu subfolder that is created automatically.

h. *For restoring data to the original location:* Use the **Conflict resolution** switch if you want to resolve any potential data conflict by keeping the most recent version of the items in conflict. Otherwise, HYCU will overwrite the existing items with the ones from the backup.

i. Enter a temporary recovery database path. The default one is C:\ProgramData\Hycu.

5. Click **Restore**.

Restoring Oracle database instances and tablespaces

With HYCU, you can restore the whole Oracle database instance or the selected tablespaces to the original location.

Prerequisites

- On the original virtual machine, references in the /etc/fstab system configuration file entries must use universally unique identifiers (for example, UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5) rather than device names (for example, /dev/sda1) unless they refer to logical volumes (for example, /dev/mapper/ol-root).
- The `bashrc` and `.bash_profile` scripts may not write to standard output (STDOUT) or standard error (STDERR) for the user whose credentials are used for application discovery.
- *Only if you are restoring data that is stored in the archive storage tier on an Azure target.* You must recreate a snapshot and use this snapshot for restoring data, or manually rehydrate data. For instructions on how to

recreate a snapshot, see “[Recreating snapshots](#)” on page 340. For instructions on how to manually rehydrate data, see Azure documentation.

Limitations

- Tablespaces can be restored only from the latest restore point in the backup chain and cannot be restored to a point in time.
- Only if you plan to use the Archive restore point.* Performing the point-in-time restore is not supported.

Considerations

- When performing a database instance or tablespace restore, you can perform a complete or point-in-time restore:

- Complete restore

HYCU performs a complete restore of the whole database instance or tablespaces from the latest backup in the backup chain.

When performing the complete restore, the control file and archive log files are not restored, and only the existing archive log files are applied. If the control file or the existing archive log files are lost, a complete restore is not possible and a point-in-time restore must be performed.

- Point-in-time restore

To perform a point-in-time restore, you must select a backup that was performed before the specified point in time so that the database instance can be brought to the point in time by applying the archive log files from the next backup.

When performing the point-in-time restore, the control file, database files, and required archive log files are restored.

Important After a successful point-in-time restore, the archive log files are reset. Therefore, it is highly recommended to perform a backup immediately after performing the point-in-time restore because the database will not be protected in terms of a complete restore until a new backup is performed.

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- You cannot perform a restore of an application whose retention period specified in the policy has been exceeded (such restore points are grayed

out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up` configuration setting in the HYCU `config.properties` file to true. For details on how to customize the HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on page 523.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application whose database you want to restore to open the Detail view.

 **Note** The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open the Detail view.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

 **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the database instance.
3. Click  **Restore**. The Restore Oracle Database dialog box opens.

 **Note** If the Back up and truncate Oracle archive logs option was disabled during the backup, you are prompted that database recovery must be performed after the restore.
4. Select **Restore databases**, and then click **Next**.
5. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
6. *Only if you want to restore the whole database instance.* Select the **Whole instance** check box.

7. *Only if you want to restore tablespaces.* Select the **Tablespace name** check box to restore all tablespaces or select the tablespaces that you want to restore.
8. *Only if restoring the whole database instance.* Optionally, specify a point in time to which you want to restore data. The database instance will be restored to the state it was in at the specified time.
9. Click **Restore**.
10. *Only if the Back up and truncate Oracle archive logs option was disabled during the backup.* Recover the Oracle databases by applying the archive logs manually.

Chapter 6

Protecting file shares

HYCU enables you to protect your file share data with fast and reliable backup and restore operations. After you back up a file share, you can choose to restore either the whole file share or individual files.

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

- “Configuring file share backup options” below
- “Backing up file shares” on page 266
- “Restoring file share data” on page 268

Configuring file share backup options

Before you start protecting file shares, you can adjust file share protection to the needs of your data protection environment by configuring backup options.

Backup option	Description
Exclude files or folders from backup	You can specify any file share folder or file to be excluded from the file share backup. Enter the full path (from the root of the file share) to the files or the folders that you want to exclude.
Incremental forever backup	If you want all your file share backups after the initial full backup to be only incremental backups, you can configure HYCU to track and store only the changes since the last backup (as incremental backups), allowing you to operate with a single backup chain.

Consideration

Only if you enable the Incremental forever backup option. To maintain a single backup chain, HYCU analyzes the whole backup chain after every backup (that is, it searches for any restore points with the Backup and Copy tiers whose retention period expired) and then performs the automatic merge of the relevant restore points, keeping only a limited number of restore points. This means that data from several incremental backups that are marked for expiration is combined and merged into the next restore point in the backup chain. However, you can at any time merge restore points also manually by expiring the restore point tiers. For instructions, see “[Managing data retention](#)” on page 396.

Accessing the Shares panel

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

Procedure

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

I want to...	Instructions
Exclude one or more file share folders or files from the backup.	<p>In the Exclude files or folders from backup field, enter the full path (from the root of the file share) to the file he files or the folders that you want to exclude from the backup (for example, <code>/backup</code>), and then click  Add. Repeat this step to add additional file share folders or files.</p> <p> Note The paths to all the file share folders or files that you excluded from the backup are added to the Exclude folder paths list. If you want to remove any of them from the exclude list, click .</p>
Enable the incremental forever backup.	Enable the Incremental forever backup switch.

4. Click **Save**.

Backing up file shares

A file-share backup allows a rapid backup of file shares by using parallel backup streams.

Prerequisite

For Dell PowerScale OneFS SMB file shares: The backup operator must have full permissions on all of the file shares that you plan to protect.

Limitations

- The iSCSI, Nutanix, and tape targets cannot be used for storing file share data.
- Backing up from a replica is not supported for Nutanix Files. Therefore, if a policy that you plan to assign to file shares has the Backup from replica option enabled, this option will be ignored.
- Backing up file shares to cloud targets is supported if the file system item names contain only characters in the Unicode Basic Multilingual Plane (BMP).
- If Snapshot is defined as the backup target type in your policy, such a policy cannot be assigned to a file share.
- *For NFS file shares:* Backing up files whose file names contain non-UTF-8 multilingual characters (for example, those created by Windows clients) is not supported. Therefore, such files will be skipped during the backup.
- If you use Smart disaster recovery (DR) for Nutanix Files protection, HYCU enables you to protect replicated file share data. After you add a recovery file server as a source to HYCU, you can back up the corresponding file shares by assigning policies to them, and later also restore them. Keep in mind that you cannot restore data to replicated file shares. For details on how to configure Smart DR, see Nutanix documentation.
- Backing up files on the file shares whose file names are longer than 255 bytes is supported only for Nutanix Files and Dell PowerScale OneFS.
- *For S3 file shares:* HYCU does not support backing up files whose names contain the leading slash character.

Considerations

- You can change the number of incremental file share backups after which a full reindex is performed by customizing the `afs.reindex.interval.count`

configuration setting, which allows you to speed up the process of searching for the relevant files when you are restoring them. For details on how to do this, see [“Customizing HYCU configuration settings” on page 523](#).

- By default, if up to 10,000 and 1% file backups fail during the backup of a file share, the backup status of the file share is marked as Completed with errors (and not as Failed). You can customize these values by editing the `afs.partial.success.threshold.count` and `afs.partial.success.max.fail.fraction` configuration settings. For details on how to do this, see [“Customizing HYCU configuration settings” on page 523](#).
- When backing up a file share, HYCU also backs up any nested shares that are inside the selected file share. Keep in mind that backing up nested file shares individually is not supported.
- *Only if you enabled the Incremental forever backup option for a file share.* When you assign a policy to the file share, the new backup chain settings defined in the policy will be ignored and full backups will not be performed.
- *For Nutanix Files:* Backing up connected file shares is supported. Keep in mind that connected file shares must be backed up individually because the backup of a parent file share does not include the contents of child file shares.
- Backing up file shares with tiered files is supported. However, consider the following:
 - The backup operators or the HYCU instance IP addresses must not be set up as zero users or clients because this could cause backup data corruption on tiered files.
 - Additional fees may apply for backup and restore operations due to data egress.
- *For Dell PowerScale OneFS:* When the incremental backup is run, the status of the Create snapshot job shows the values obtained from the Dell PowerScale OneFS file server.

Recommendations

- Using an NFS target for storing file share data requires you to enable public access to the target. For security purposes, it is recommended that you avoid such a configuration.
- *For S3 and generic file shares:* Avoid making any changes to your file shares during the backup process.

Accessing the Shares panel

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

Procedure

1. In the Shares panel, select the file shares that you want to back up.

 **Tip** You can update the list of file shares by clicking  **Refresh**. To narrow down the list of displayed file shares, you can use the filtering options described in “Filtering and sorting data” on page 317.

2. Click  **Set Policy**.
3. From the list of available policies, select the preferred policy.
4. Click **Assign** to assign the policy to the selected file shares.

After you assign the policy, the backup is scheduled according to the values that you defined for your policy. If required, you can also perform a manual backup at any time. For details, see “Performing a manual backup” on page 333.

 **Tip** If you have more than one HYCU instance in your data protection environment, you can see which HYCU instance performed a backup by clicking the preferred backup job in the Jobs panel and checking the HYCU instance IP address in the Detail view.

Restoring file share data

You can restore a whole file share or individual files to the original or a different file server, to an external SMB or NFS file share, or to a local machine.

The file share data can be restored from a target or a snapshot. Restoring data from the snapshot is possible only if the `afs.restore.snapshot.enabled` configuration setting is set to `true` (the default value is `false`). In this case, the restore is always performed from the snapshot if the snapshot is available. Otherwise, the restore is performed from the target. For details on how to customize HYCU configuration settings, see “Customizing HYCU configuration settings” on page 523.

 **Important** Restoring the file share data from a snapshot is not supported for S3 and generic file shares.

Prerequisite

For restoring data to a different file server share: The file server with the file share to which you want to restore data must be added to HYCU. For details on how to do this, see “[Adding a file server](#)” on page 62.

Limitations

- The restore of alternate data streams (ADS) is supported only if you are restoring data from one file server SMB share to another file server SMB share.
- *Only if restoring Nutanix Files shares that contain alternate data streams (ADS) in top-level directories to distributed file shares.* Restoring ADS to top-level directories of distributed file shares is not supported. ADS can be restored to subdirectories of distributed file shares or to standard file shares.
- Symbolic links are restored only when restoring data from one NFS file share to another NFS file share, or from an SMB file share to an NFS file share.
- *Only if restoring files to an external file share.* Restoring files or folders with newlines in their names is supported only for an NFS share set up on Unix.
- *Only if restoring files to a local machine:*
 - The files can be restored only if the size of the uncompressed files is less than or equal to 2 GiB.
 - Restoring the original access control list for the files is not supported.
- *Only if restoring S3 file shares to an alternate location.* The file destination path must not contain the leading slash character.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- *Only if restoring a large number of files from the file share backup.* The HYCU instance may require more RAM than is available by default. In this case, increase the default value by using the `afs.instance.memory.mb` configuration setting. For details on how to customize the HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on page 523.

- *Only if restoring files to a local machine.* The restored files are downloaded as a .zip file. To avoid any potential issues with unzipping the restored files and to make sure that the files or folders with newlines in their names are properly restored, always use 7-Zip when extracting the files.
- If the number of files that could not be restored during the file share restore is less than or equal to 100 (the default value), the status of the file share restore is Warning. You can edit this default value by customizing the `afs.restore.partial.success.threshold.count` configuration setting. For details on how to do this, see “[Customizing HYCU configuration settings](#)” on page 523.

Recommendation

For optimal restore performance, it is recommended that you restore data to a file server share instead of an external file share whenever possible.

Accessing the Shares panel

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

Procedure

1. In the Shares panel, click the file share that contains the files that you want to restore to open the Detail view.

 **Note** The Detail view appears only if you click a file share. Selecting the check box before the name of the file share will not open the Detail view.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Files**. The Restore Files dialog box opens.
4. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
5. Click **Next**.

6. In the Folder section, select the uppermost check box (the one in front of the  icon) if you want to restore the whole file share. Otherwise, from the list of available folders and files, select the ones that you want to restore. Click **Next**.

 **Tip** If there are too many files to be displayed on one page, you can move between the pages by clicking  and . You can also use  to set the number of files to be displayed per page.

7. Depending on where you want to restore the selected files (to the original or a different file server share, an external SMB or NFS file share, or the local machine), select the preferred restore option, click **Next**, and then follow the instructions:

Restore option	Instructions
Restore to file server share	<p>a. From the Share drop-down menu, select the file server share to which you want to restore the files.</p> <p>b. Select whether you want to restore the files to the original location or an alternate location on the same file server share. If you select an alternate location, in the Path field, specify the part of the path that is appended to the location of the selected share.</p> <p> Note Under the Path field, the full alternate location path is displayed and updated as you type.</p> <p>c. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file).</p> <p> Important If you plan to rename the original files, you must be a file server admin. For all other operations, you can be either a file server or a backup admin.</p> <p>d. <i>Only if restoring files from one SMB file share to another SMB file share or from one S3 file share to another S3 file share.</i> Enable the Restore ACLs switch</p>

Restore option	Instructions
	<p>if you want to restore the original access control lists for the files.</p> <p>e. <i>Only if restoring files from one S3 file share to another S3 file share.</i> Enable the Restore S3 object tags switch if you want to restore the original object tags for the files.</p> <p>f. Click Restore.</p>
Restore to external share	<p>From the Share type drop-down menu, select where you want to restore the files, and then provide the required information:</p> <ul style="list-style-type: none"> • NFS <ol style="list-style-type: none"> a. Enter the path to the NFS shared folder in the following format: <div style="background-color: #f0f0f0; padding: 5px; text-align: center;"><code>\server\<Path></code></div> <ol style="list-style-type: none"> b. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file). c. Click Restore. • SMB <ol style="list-style-type: none"> a. Enter the path to the SMB shared folder in the following format: <div style="background-color: #f0f0f0; padding: 5px; text-align: center;"><code>\server\<Path></code></div> <ol style="list-style-type: none"> b. <i>Optional.</i> Provide user credentials to access the SMB file share. c. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file).

Restore option	Instructions
	<p>d. <i>Only if restoring files from one SMB file share to another SMB file share.</i> Enable the Restore ACLs switch if you want to restore the original access control list for the files.</p> <p>Important If you enable the Restore ACLs switch, keep in mind that the restored files might not be accessible due to the ACLs not being recognized on the destination SMB file share.</p> <p>e. Click Restore.</p>
Download	<p>Click Download to restore the selected files to the local machine.</p> <p>Important Do not refresh the page or navigate away from the page until the download process job finishes.</p>

Chapter 7

Protecting volume groups

HYCU enables you to protect Nutanix volume groups with fast and reliable backup and restore operations. After you back up a volume group, you can choose to restore either the whole volume group or only individual virtual disks by exporting them to an NFS or SMB share.

① Important If the volume groups are attached to one or more virtual machines at backup time, they are backed up automatically during the virtual machine backup. For details, see “[Protecting virtual machines](#)” on [page 128](#).

For details on how to protect volume groups efficiently, see the following sections:

- “[Backing up volume groups](#)” below
- “[Restoring volume groups](#)” on the next page

Backing up volume groups

With HYCU, you can back up Nutanix volume groups in a fast and efficient way.

Prerequisite

A Nutanix cluster on which the volume group that you want to protect resides must be added to HYCU. For instructions, see “[Adding a Nutanix cluster](#)” on [page 54](#).

Consideration

The volume groups that HYCU creates automatically and uses for data protection purposes are not shown in the Volume Groups panel. The names of these volume groups start with the `NTNX-`, `hycu-vg-`, and `HYCU-` prefixes, therefore make sure not to create your own volume groups with the same prefixes.

Accessing the Volume Groups panel

To access the Volume Groups panel, in the navigation pane, click  **Volume Groups**.

Procedure

1. In the Volume Groups panel, select the volume groups that you want to back up.
- Tip** You can update the list of volume groups by clicking  **Refresh**. To narrow down the list of displayed volume groups, you can use the filtering options described in “[Filtering and sorting data](#)” on page 317.
2. Click  **Set Policy**.
3. From the list of available policies, select the preferred policy.
4. Click **Assign** to assign the policy to the selected volume groups.

The backup is scheduled according to the values that you defined for your policy. If required, you can also perform a manual backup of any volume group at any time. For details, see “[Performing a manual backup](#)” on page 333.

Restoring volume groups

HYCU enables you to restore either a whole volume group or only individual virtual disks that became corrupted.

Consideration

If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.

Restore options

You can select among the following restore options:

Restore option	Description
Restore Volume Group	Enables you to restore a volume group. Select this option if you want to replace the original volume group with the restored one. For instructions, see

Restore option	Description
	“Restoring a volume group” below.
Clone Volume Group	Enables you to restore a volume group by creating its clone. Select this option if you want to keep the original volume group. For instructions, see “ Cloning a volume group ” on the next page.
Export vDisks	Enables you to restore virtual disks to an NFS or SMB share. Select this option if you want to make the virtual disks available to users with specific access permissions, or if you want to use the virtual disks later to restore data to an environment with a source not supported by HYCU or not added to HYCU. For instructions, see “ Exporting virtual disks ” on page 278.

Accessing the Volume Groups panel

To access the Volume Groups panel, in the navigation pane, click  **Volume Groups**.

Restoring a volume group

You can restore a volume group to its original or a new location. In this case, the original volume group will be overwritten.

Consideration

Only if the volume group is attached to one or more virtual machines. The virtual machines to which the volume group is attached must be turned off.

Procedure

1. In the Volume Groups panel, click the volume group that you want to restore. The Detail view appears at the bottom of the screen.

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

2. In the Detail view, select the preferred restore point.
3. Click  **Restore**.
4. Select **Restore Volume Group**, and then click **Next**.

5. From the Storage container drop-down menu, select where you want to restore the volume group. By default, the original storage container is selected.
6. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
7. *Only if the volume group you are restoring is attached to one or more virtual machines.* Enable the **Attach volume group** switch if you want the volume group to be attached to the virtual machines after the restore.
8. Click **Restore**.

Cloning a volume group

You can create a clone of the original volume group by restoring the volume group to its original or a new location. In this case, the original volume group will not be overwritten.

Procedure

1. In the Volume Groups panel, click the volume group that you want to restore. The Detail view appears at the bottom of the screen.

Note The Detail view appears only if you click a volume group. Selecting the check box before the name of the volume group will not open the Detail view.
2. In the Detail view, select the preferred restore point.
3. Click  **Restore**.
4. Select **Clone Volume Group**, and then click **Next**.
5. From the Storage container drop-down menu, select where you want to restore the volume group. By default, the original storage container is selected.
6. In the New volume group name field, specify a new name for the volume group.

7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
8. *Only if the volume group you are restoring is attached to one or more virtual machines.* Enable the **Attach volume group** switch if you want the volume group to be attached to the virtual machines after the restore.
9. Click **Restore**.

Exporting virtual disks

You can restore virtual disks to an NFS or SMB share.

Procedure

1. In the Volume Groups panel, click the volume group whose virtual disks you want to restore. The Detail view appears at the bottom of the screen.

Note The Detail view appears only if you click a volume group. Selecting the check box before the name of the volume group will not open the Detail view.
2. In the Detail view, select the preferred restore point.
3. Click  **Restore**.
4. Select **Export vDisks**, and then click **Next**.
5. From the list of virtual disks that are available for the restore, select the ones that you want to restore, and then click **Next**.
6. From the Type drop-down menu, select where you want to restore the virtual disks, and then provide the required information:

Type	Instructions
SMB	<ol style="list-style-type: none"> a. <i>Optional.</i> Enter the domain and user credentials. b. Enter the SMB server name or IP address. c. Enter the path to the SMB shared folder from the root of the server (for example, <code>/backups/HYCU</code>).

Type	Instructions
NFS	<ol style="list-style-type: none">a. Enter the NFS server name or IP address.b. Enter the path to the NFS shared folder from the root of the server (for example, /backups/HYCU).

7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic:** Ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
8. Click **Restore**.

Chapter 8

Recovering your data protection environment

If a disaster occurs in your data protection environment and your data is corrupted or unavailable, HYCU provides an effective approach to recovering data by importing targets on which backup data is stored. You can decide to recover the following:

- HYCU backup controller and use it later to restore data
- Virtual machines, applications, file shares, and volume groups

The procedures described in this section are applicable if your backup data is stored on targets. If you selected Snapshot as the backup target type in your policy and no data archives exist, you can use the snapshots created by HYCU to perform disaster recovery through the Nutanix Prism web console, the vSphere (Web) Client, the AWS GovCloud (US) console, or the Azure Government portal. For details, see Nutanix, VMware, AWS, or Azure documentation.

Considerations

- After performing the disaster recovery, not all your restore points might be used for restoring data because of the expected discrepancy between data in the database and data on targets.
- *For DR-ready virtual machines and applications:* You can recover your HYCU backup controller and protected data to cloud by using the SpinUp functionality. For more information, see [“Protecting data across on-premises and cloud environments” on page 485](#).

Procedures

1. Prepare for disaster recovery. For instructions, see [“Preparing for disaster recovery” on the next page](#).
2. Perform disaster recovery. For instructions, see [“Performing disaster recovery” on page 286](#).

3. *Only if HYCU is used for file share protection:* Create HYCU instances. For instructions, see “[HYCU instances](#)” on page 63.

Preparing for disaster recovery

Prerequisites

- You must know configuration parameters of the targets that store the backup of your original HYCU backup controller or backups of other entities you want to recover. For details, see “[Preparing for disaster recovery](#)” on page 132.
- The targets that store backup data of the entities you want to recover must be accessible to the source where you plan to deploy a recovery HYCU backup controller.
- *Only if the backup of the original HYCU backup controller is stored on an iSCSI target.* The iSCSI storage device must be dedicated to a single HYCU backup controller and no other appliances than HYCU.
- *Only if the backup of the original HYCU backup controller or virtual machines, applications, file shares, and volume groups you want to recover is stored on a Google Cloud target.* A Google Cloud service account must be created and added to HYCU. For instructions on how to add a cloud account to HYCU, see “[Adding a Google Cloud service account](#)” on page 367.
- *Only if the backup of the original HYCU backup controller or other entities you want to recover is stored on a target with enabled target encryption.* The encryption target key from the original HYCU backup controller must be exported and the file containing the encryption key must be available.

When preparing for disaster recovery, you must perform the following tasks:

Task	Instructions
1. Deploy a recovery HYCU backup controller.	“ Deploying a recovery HYCU backup controller ” on the next page
2. Import the targets that store the backup of the original HYCU backup controller. The imported targets may also contain backups of virtual machines, applications, file	“ Importing targets ” on page 284

Task	Instructions
shares, and volume groups.	
3. Add a source to which you plan to restore your HYCU backup controller. If you plan to restore also virtual machines, applications, file shares, and volume groups, add the sources to which you plan to restore them.	“Adding sources” on page 53

Deploying a recovery HYCU backup controller

Procedure

1. Sign in to the Nutanix Prism web console (for Nutanix AHV clusters), the vSphere (Web) Client (for Nutanix ESXi clusters and vSphere environments), the AWS GovCloud (US) console (for AWS GovCloud (US) environments), or the Azure Government portal (for Azure Government environments).
2. Deploy a recovery HYCU backup controller that you will use for restoring the original HYCU backup controller or other entities:

I want to deploy the HYCU backup controller to...	Instructions
Nutanix AHV cluster	“Deploying HYCU to a Nutanix AHV cluster” on page 35
Nutanix ESXi cluster or vSphere environment	“Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment” on page 40
AWS GovCloud (US) environment	“Deploying HYCU to an AWS GovCloud (US) environment” on page 42
Azure Government environment	“Deploying HYCU to an Azure Government environment” on page 44

3. *Only if you plan to restore the HYCU backup controller to a different source.*
Enable the creation of a clone of the HYCU backup controller. To do so, in

the HYCU config.properties file, set the `clone.enabled.for.hycu.dr` configuration setting to `true`. For instructions on how to customize HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on [page 523](#).

⚠ Caution Make sure that a clone of the HYCU backup controller is not activated while the original HYCU backup controller is still active. Otherwise, data loss may occur.

4. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
5. *Only if the backup of the original HYCU backup controller or backups of the entities you want to restore is stored on a target with enabled target encryption.* Import the encryption key that you have exported from the original HYCU backup controller. For instructions, see “[Configuring target encryption](#)” on [page 372](#).

After deploying a recovery HYCU backup controller

Depending on your data protection needs, after you perform disaster recovery, you can decide to keep or delete the recovery HYCU backup controller. If you delete the recovery HYCU backup controller, you will have to deploy a new one the next time you perform disaster recovery.

Limitation

For Nutanix and iSCSI targets: Keeping the recovery HYCU backup controller is not supported. If you want to use such targets for disaster recovery, you must deploy a new recovery HYCU backup controller every time.

Considerations

If you decide to keep the recovery HYCU backup controller, consider the following:

- After a successful import of targets, the recovery HYCU backup controller is automatically put in recovery mode and the following applies:
 - HYCU automatically refreshes the imported targets every 60 minutes to get the information about the latest restore points (the backups stored on the targets), as well as the information about the targets that are available for importing or that have been deleted.

 **Note** You can at any time refresh the imported targets also manually. To do this, in the Targets panel, click  **Refresh**.

- Backup operations are disabled. This means that you cannot assign policies, perform manual backups, or expire backups manually.
- Setting power options is disabled.
- Only limited target options can be edited.
- Adding targets is disabled.
- For successful target synchronization, the recovery HYCU backup controller must be deployed with HYCU version 4.5.0.
- Deactivated targets are excluded from target synchronization.
- The default automatic target synchronization value can be adjusted to your data protection needs. For details on how to customize HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on page 523.

Importing targets

Prerequisites

- The activities on the original HYCU backup controller (if it still exists) must be suspended and no jobs may be running. For instructions, see “[Setting power options](#)” on page 403.
- No targets or only imported targets can exist on the recovery HYCU backup controller. Otherwise, importing targets is disabled.
- *For importing iSCSI or Nutanix targets:* The targets must be unmounted on any other powered on HYCU backup controller.

Limitation

Backing up data to imported targets is not supported.

Considerations

- The targets you import should contain the complete backup chains of the entities you want to recover.
- Make sure not to make any changes to HYCU until the import job is finished.

Accessing the Targets panel

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

Procedure

1. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
2. In the Targets panel, click  **Import**. The Import Target dialog box opens.
3. From the Type drop-down menu, select the type of target.
4. Specify the values so that they match the original target configuration, and then click **Next**.
5. In the Import Target Catalog dialog box, select the name of the HYCU backup controller whose backup you want to import, and then click **Next**.
6. In the Multiple Targets dialog box, one or more targets that store backup data of the selected HYCU backup controller and other entities is displayed. If any additional targets are found, select them one by one and specify the values so that they match the original target configuration. For each target, click **Validate** to check the configuration.
7. After you validated all the targets required for your restore, click **Import**.

 **Note** It is recommended to import all the targets from the list to ensure that complete backup chains are available for the restore. If you do not import some targets and backup chains are not complete, you can import missing targets later by repeating the import procedure.

After a successful import of targets

- The imported targets are listed in the Targets panel and their mode is set to Read-Only, which prevents you from storing backup data to these targets.
- The HYCU backup controller is listed in the Virtual Machines panel, and its status is Protected deleted.
- For recovering virtual machines, applications, file shares, and volume groups, consider the following:
 - The self-service groups existing in the original data protection environment are recreated on the recovery HYCU backup controller. The recreated self-service groups do not contain any users. To restore virtual machines, applications, file shares, and volume groups, you need to create users and add them to the recreated user groups that have ownership over the virtual machines, file shares, and volume groups you want to restore. For instructions, see “[Setting up a user environment](#)” on

page 347.

- The virtual machines whose backups are stored on the imported targets are listed in the Virtual Machines panel, and their status is Protected deleted. To restore virtual machines other than the HYCU backup controller, see “Restoring virtual machines” on page 155.
- Applications whose backups are stored on the imported targets are listed in the Applications panel, and their status is Protected deleted. To restore applications, see “Restoring whole applications” on page 223.
- File shares whose backups are stored on the imported targets are listed in the Shares panel, and their status is Protected deleted. To restore file shares, see “Restoring file share data” on page 268.
- Volume groups whose backups are stored on the imported targets are listed in the Volume Groups panel, and their status is Protected deleted. To restore volume groups, see “Restoring volume groups” on page 275.

Performing disaster recovery

Perform disaster recovery by using one of the following approaches:

I want to recover...	Instructions
The HYCU backup controller to the original source by using a restore point created with HYCU version 4.0.0 or later.	“Restoring the HYCU backup controller to the original source” on the next page
The HYCU backup controller to a different source by using a restore point created with HYCU version 4.0.0 or later.	“Restoring the HYCU backup controller to a different source” on page 290
The HYCU backup controller to the original or a different source by using a restore point created with a HYCU version earlier than 4.0.0.	“Exporting virtual disks” on page 197
Virtual machines	“Restoring virtual machines” on page 155
Applications	“Restoring whole applications” on

I want to recover...	Instructions
	page 223
File shares	“Restoring file share data” on page 268
Volume groups	“Restoring volume groups” on page 275

Restoring the HYCU backup controller to the original source

Depending on the source on which the original HYCU backup controller was running, see one of the following sections:

- “Restoring the HYCU backup controller to a Nutanix cluster or a vSphere environment” below
- “Restoring the HYCU backup controller to an AWS GovCloud (US) environment” on the next page
- “Restoring the HYCU backup controller to an Azure Government environment” on page 289

Restoring the HYCU backup controller to a Nutanix cluster or a vSphere environment

Use this procedure when the original cluster of the HYCU backup controller is not damaged.

Prerequisites

- The recovery HYCU backup controller must have network access to the cluster of the original HYCU backup controller.
- *Only if the backup of the original HYCU backup controller is stored on an iSCSI or a Nutanix target:* The target must be deactivated and detached from the recovery HYCU backup controller before you power on the restored HYCU backup controller.

Procedure

1. Sign in to the HYCU web user interface of the recovery HYCU backup controller.

2. In the Virtual Machines panel, select the HYCU backup controller.
3. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

4. Click  **Restore**.
5. Select **Restore VM**, and then restore the HYCU backup controller by following the instructions described in “[Restoring a virtual machine to a Nutanix cluster or a vSphere environment](#)” on page 160. The activities of the restored HYCU backup controller are suspended automatically.
6. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
7. Sign in to the HYCU web user interface of the restored HYCU backup controller.
8. Resume the activities of the restored HYCU backup controller. For instructions, see “[Setting power options](#)” on page 403.
9. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see the relevant documentation.
10. *For Nutanix ESXi clusters:* If the original HYCU backup controller does not exist, configure settings for the new network adapter that was assigned to the HYCU backup controller. For instructions, see “[Configuring your network](#)” on page 392.

 **Important** Make sure to enter the original IP address of the HYCU backup controller. After editing the connection, delete the old network adapter.

Restoring the HYCU backup controller to an AWS GovCloud (US) environment

Use this procedure when you want to restore the HYCU backup controller to the original AWS GovCloud (US) region.

Procedure

1. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
2. In the Virtual Machines panel, select the HYCU backup controller.
3. In the Detail view that appears at the bottom of the screen, select the latest restore point.

Note The Detail view appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Detail view.
4. Click  **Restore**,
5. Select **Restore VM**, and then restore the HYCU backup controller by following the instructions described in [“Restoring a virtual machine to an AWS GovCloud \(US\) environment” on page 164](#).

The activities of the restored HYCU backup controller are suspended automatically.
6. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
7. Sign in to the HYCU web user interface of the restored HYCU backup controller.
8. Resume the activities of the restored HYCU backup controller. For instructions, see [“Setting power options” on page 403](#).
9. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see the relevant documentation.

Restoring the HYCU backup controller to an Azure Government environment

Use this procedure when you want to restore the HYCU backup controller to the original Azure Government subscription.

Procedure

1. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
2. In the Virtual Machines panel, select the HYCU backup controller.

3. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

4. Click  **Restore**.
5. Select **Restore VM**, and then restore the HYCU backup controller by following the instructions described in “Restoring a virtual machine to an Azure Government environment” on page 168.
The activities of the restored HYCU backup controller are suspended automatically.
6. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
7. Sign in to the HYCU web user interface of the restored HYCU backup controller.
8. Resume the activities of the restored HYCU backup controller. For instructions, see “Setting power options” on page 403.
9. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see the relevant documentation.

Restoring the HYCU backup controller to a different source

Depending on the source to which you want to restore the HYCU backup controller, see one the following sections:

- “Restoring the HYCU backup controller to a Nutanix cluster or a vSphere environment” on the next page
- “Restoring the HYCU backup controller to an AWS GovCloud (US) environment” on page 292
- “Restoring the HYCU backup controller to an Azure Government environment” on page 294

Restoring the HYCU backup controller to a Nutanix cluster or a vSphere environment

Use this procedure when the cluster of the original HYCU backup controller is damaged or inoperable, or if you want to relocate the HYCU backup controller.

Prerequisites

- The recovery HYCU backup controller must have network access to the cluster you plan to restore the original HYCU backup controller to.
- *Only if the backup of the original HYCU backup controller is stored on an iSCSI or a Nutanix target:* The target must be deactivated and detached from the recovery HYCU backup controller before you power on the restored HYCU backup controller.

Procedure

1. *Only if the original HYCU backup controller still exists.* Suspend the activities of the original HYCU backup controller.

⚠ Caution Make sure that a clone of the HYCU backup controller is not activated while the original HYCU backup controller is still active. Skipping this step may result in data loss.

To suspend the activities of the original HYCU backup controller, follow these steps:

- a. *Only if the HYCU backup controller is turned off.* Turn the HYCU backup controller (virtual machine) on.
- b. Sign in to the HYCU web user interface.
- c. Suspend the activities of the HYCU backup controller. For instructions, see [“Setting power options” on page 403](#).
- d. Wait for the running jobs to complete. You can check this by filtering the Jobs list by the Executing job status. For instructions, see [“Filtering and sorting data” on page 317](#).

2. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
3. In the Virtual Machines panel, select the original HYCU backup controller.
4. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

5. Click  **Restore**.
6. Select **Clone VM**, and then restore the HYCU backup controller by following the instructions described in “[Cloning a virtual machine to a Nutanix cluster or a vSphere environment](#)” on page 172.
The activities of the restored HYCU backup controller are suspended automatically.
7. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
8. Sign in to the HYCU web user interface of the restored HYCU backup controller.
9. Resume the activities of the restored HYCU backup controller. For instructions, see “[Setting power options](#)” on page 403.
10. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see the relevant documentation.
11. *Only if you want to use network settings of the original HYCU backup controller.* Configure settings for the network adapter of the HYCU backup controller. For instructions, see “[Configuring your network](#)” on page 392.

 **Note** Make sure to enter the original IP address of the HYCU backup controller.

Restoring the HYCU backup controller to an AWS GovCloud (US) environment

Use this procedure if you want to relocate the HYCU backup controller.

Procedure

1. *Only if the original HYCU backup controller still exists.* Suspend the activities of the original HYCU backup controller.

 **Caution** Make sure that a clone of the HYCU backup controller is not activated while the original HYCU backup controller is still active. Skipping this step may result in data loss.

To suspend the activities of the original HYCU backup controller, follow these steps:

- a. *Only if the HYCU backup controller is turned off.* Turn the HYCU backup controller virtual machine on.
- b. Sign in to the HYCU web user interface.
- c. Suspend the activities of the HYCU backup controller. For instructions, see “[Setting power options](#)” on page 403.
- d. Wait for the running jobs to complete. You can check this by filtering the Jobs list by the Executing job status. For instructions, see “[Filtering and sorting data](#)” on page 317.

2. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
3. In the Virtual Machines panel, select the original HYCU backup controller.
4. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

5. Click  **Restore**.
6. Select **Clone VM**, and then restore the HYCU backup controller by following the instructions described in “[Cloning a virtual machine to an AWS GovCloud \(US\) environment](#)” on page 177.

The activities of the restored HYCU backup controller are suspended automatically.

7. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
8. Sign in to the HYCU web user interface of the restored HYCU backup controller.
9. Resume the activities of the restored HYCU backup controller. For instructions, see “[Setting power options](#)” on page 403.
10. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see the relevant documentation.

Restoring the HYCU backup controller to an Azure Government environment

Use this procedure if you want to relocate the HYCU backup controller.

Procedure

1. *Only if the original HYCU backup controller still exists.* Suspend the activities of the original HYCU backup controller.

⚠ Caution Make sure that a clone of the HYCU backup controller is not activated while the original HYCU backup controller is still active. Skipping this step may result in data loss.

To suspend the activities of the original HYCU backup controller, follow these steps:

- a. *Only if the HYCU backup controller is turned off.* Turn the HYCU backup controller (virtual machine) on.
- b. Sign in to the HYCU web user interface.
- c. Suspend the activities of the HYCU backup controller. For instructions, see [“Setting power options” on page 403](#).
- d. Wait for the running jobs to complete. You can check this by filtering the Jobs list by the Executing job status. For instructions, see [“Filtering and sorting data” on page 317](#).

2. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
3. In the Virtual Machines panel, select the original HYCU backup controller.
4. In the Detail view that appears at the bottom of the screen, select the latest restore point.

💡 Note The Detail view appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Detail view.

5. Select **Clone VM**, and then restore the HYCU backup controller by following the instructions described in [“Cloning a virtual machine to an Azure Government environment” on page 181](#).

The activities of the restored HYCU backup controller are suspended automatically.

6. Sign out of the HYCU web user interface of the recovery HYCU backup controller..

7. Sign in to the HYCU web user interface of the restored HYCU backup controller.
8. Resume the activities of the restored HYCU backup controller. For instructions, see [“Setting power options” on page 403](#).
9. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see the relevant documentation.

Chapter 9

Performing daily tasks

To ensure the secure and reliable performance of the data protection environment, HYCU provides various mechanisms to support your daily activities.

I want to...	Procedure
Get an at-a-glance overview of the data protection environment state, identify eventual bottlenecks, and inspect different areas of the data protection environment.	“Using the HYCU dashboard” on the next page
Track jobs that are running in my environment, get an insight into a specific job status, generate a job report, and cancel a currently running job.	“Managing HYCU jobs” on page 300
View all events that occurred in my environment.	“Managing HYCU events” on page 301
Configure HYCU to send notifications when events occur.	“Configuring event notifications” on page 302
Enable purging of events and jobs.	“Enabling the purge of events and jobs” on page 305
Obtain reports on different aspects of the data protection environment.	“Using HYCU reports” on page 306
View entity details.	“Viewing entity details” on page 313
Narrow down the list of displayed items by applying filters.	“Filtering and sorting data” on page 317
Export data that you can view in a table in any of the panels to a JSON or CSV file.	“Exporting the contents of the panel” on page 324

I want to...	Procedure
View target information, activate or deactivate a target, increase the size of an iSCSI target, or edit or delete a target.	“ Managing targets ” on page 325
View policy information, or edit or delete a policy.	“ Managing policies ” on page 330
Back up data manually.	“ Performing a manual backup ” on page 333
Set up a validation policy and schedule the backup validation.	“ Setting up a validation policy ” on page 334
Archive data manually.	“ Archiving data manually ” on page 339
Recreate a snapshot.	“ Recreating snapshots ” on page 340

In case of recognized problems in your data protection environment that can degrade the efficiency and reliability of data protection (for example, when storage, vCPU, or memory utilization is exceeded), you can make adjustments to better meet your data protection goals. For details, see “[Adjusting the HYCU backup controller resources](#)” on page 341.

You can customize the HYCU web user interface to your needs. For details, see “[Customizing your HYCU web user interface](#)” on page 342.

Using the HYCU dashboard

The HYCU dashboard provides you with an at-a-glance overview of the data protection status in your environment. This intuitive dashboard enables you to monitor all data protection activity and to quickly identify areas that need your attention. You can use this dashboard as a starting point for your everyday tasks because it enables you to easily access the area of interest by simply clicking the corresponding widget.

Accessing the Dashboard panel

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

① Important Your user role defines which widgets you are allowed to see and access.

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

Dashboard widget	Description
Virtual Machines	<p>Shows the percentage of protected and unprotected virtual machines in the data protection environment, the percentage of compliant and non-compliant virtual machines, and the percentage of protected virtual machines that have the DR-ready status. A virtual machine is considered:</p> <ul style="list-style-type: none"> Protected: If it has at least one valid backup available and does not have the Exclude policy assigned. Compliant: If the time since the last successful backup is lower than its RPO and the estimated time to recover is lower than its RTO. DR-ready: If all backups in the current backup chain are stored on one of the cloud targets and a successful platform readiness check was performed during its latest backup. <p>For detailed information on protecting virtual machines, see “Backing up virtual machines” on page 153.</p>
Applications	<p>Shows the percentage of protected and unprotected applications in the data protection environment, and the percentage of compliant and non-compliant applications. An application is considered:</p> <ul style="list-style-type: none"> Protected: If it has at least one valid backup available and does not have the Exclude policy assigned. Compliant: If the time since the last successful backup is lower than its RPO and the estimated time to recover is lower than its RTO. <p>For detailed information about protecting applications, see “Backing up applications” on page 222.</p>
HYCU Controller*	Shows the details about your HYCU backup controller, the

Dashboard widget	Description
	average vCPU and memory utilization in the last 48 hours, and the percentage of used storage. For details about what to do if any of these values reaches a critical value (that is, if any of the values becomes red), see “ Adjusting the HYCU backup controller resources ” on page 341.
Backups	Shows the number of backups in the last 24 hours and the backup success rate for the last 7 days.
Targets*	Shows list of all targets in the data protection environment, and the information on how much space is used and available for storing data on each target and on all targets in the data protection environment combined. For detailed information about setting up targets, see “ Setting up targets ” on page 78.
Policies	Shows the number of entities in the data protection environment, the number of entities that have no policy assigned, and the number of entities that are compliant and non-compliant with the RPO and RTO set in their assigned policy. The number of compliant and non-compliant entities for specific policies is also shown. For detailed information about policies, see “ Defining your backup strategy ” on page 107.
Jobs	Shows the number of jobs in the data protection environment in the last 56 hours according to their status (Success, Warning, Failed, In progress, and Queued). For details on jobs, see “ Managing HYCU jobs ” on the next page.
Events	Shows the number of events in the data protection environment in the last 56 hours according to their status (Success, Warning, and Failed). For details on events, see “ Managing HYCU events ” on page 301.

* An infrastructure group administrator only.

Managing HYCU jobs

In the Jobs panel, you can do the following:

- Check the processes that are currently running.
- Check the completed and stopped processes.
- Check more details about a specific job in the Detail view that appears at the bottom of the screen after you select the job.

 **Tip** By pausing on the progress bar of a particular task (for example, Backup data), additional information about the task is available, such as how much data has already been backed up and when the progress time has been last updated.

- *For virtual machines with attached volume groups:* Check the backup and restore process statuses of the volume groups attached to the virtual machines. To do so, click the arrow next to the backup or restore job of a virtual machine with attached volume groups, and a list of attached volume group processes and their statuses will be expanded. Keep in mind that volume group processes will not appear all at once, but one after another, as the job progresses.
- Generate a report about a specific job by selecting it, and then clicking  **View Report**. To copy the report to the clipboard, in the Job Report dialog box that opens, click **Copy to Clipboard**.
- Cancel a currently running or queued job by selecting it, and then clicking  **Abort Job**.
- Enable purging of jobs. For details, see “[Enabling the purge of events and jobs](#)” on page 305.

Consideration

If a backup, backup copy, or archive job fails, HYCU automatically schedules job retries. Consider the following:

- If the backup job fails, the time interval between two successive retries is doubled with each retry until the RPO value is reached (for example, by default, the first retry occurs after 15 minutes, the second one after 30 minutes, the third one after 1 hour, and so on). When the RPO value is reached, the time interval for retrying the backup job becomes the same as the one specified for the RPO.

- If the backup copy job fails, HYCU retries the failed job two times with the time interval of 15 minutes (by default). If these retries fail, the retry job is suspended for 24 hours.
- If the archive job fails, HYCU retries the failed job once after 15 minutes (by default). If this retry fails, the retry job is suspended for 12 hours.

Accessing the Jobs panel

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

 **Tip** You can update the list of jobs by clicking  **Refresh**.

The following information is available for each job:

Job information	Description
Name	Name of a job that was performed (for example, adding a source, adding a target, running a backup, and so on).
Status	Current status of a job (for example, Queued, a progress bar indicating the Executing status, OK, or Error).
Created	When a job was created.
Finished	When a job finished.

Managing HYCU events

In the Events panel, you can do the following:

- View all events that occurred in your environment.
- Check details about the selected event.
- List events that match the specified filter.
- Configure HYCU to send notifications when the events occur. For details, see “[Configuring event notifications](#)” on the next page.
- Enable purging of events. For details, see “[Enabling the purge of events and jobs](#)” on page 305.

Accessing the Events panel

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

 **Tip** You can update the list of events by clicking  **Refresh**.

The following information is available for each event:

Event information	Description
Status	Status of the event (Success, Warning, Failed)
Message	Description of the event
Category	Category to which the event belongs (for example, Policies, Backup, Credentials, System for an internal event, and so on)
Timestamp	Event creation time

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

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

Configuring event notifications

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

Accessing the Notifications dialog box

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

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

- “Setting up email notifications” below
- “Setting up webhook notifications” on the next page

Setting up email notifications

Prerequisite

Because HYCU uses SMTP to send email notifications, an SMTP server must be configured. For details, see “Configuring an SMTP server” on page 407.

Procedure

1. In the Notifications dialog box, click the **Email** tab, and then click  **New**.
2. In the Subject field, enter a subject for the email notification.
3. From the Category drop-down menu, select one or more categories to which the events belong (for example, Policies, Backup, Credentials, System, and so on). To include all categories, click **Select all**.
4. From the Status drop-down menu, select the status of the events (Success, Warning, Failed). To include all statuses, click **Select all**.
5. From the Language drop-down menu, select the preferred language for email notifications.
6. In the Email address field, enter one or more email addresses to which you want the notifications to be sent. If you are entering more than one email address, make sure to press the spacebar after entering each one.
7. Click **Save**.

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

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

Setting up webhook notifications

Procedure

1. In the Notifications dialog box, click the **Webhooks** tab, and then click  **New**.
2. Enter a name for the webhook notification and, optionally, its description.
3. From the Category drop-down menu, select one or more categories to which the events belong (for example, Policies, Backup, Credentials, System, and so on). To include all categories, click **Select all**.
4. From the Status drop-down menu, select the status of the events (Success, Warning, Failed). To include all statuses, click **Select all**.
5. From the Language drop-down menu, select the preferred language for webhook notifications.
6. In the Post URL field, enter the URL of the endpoint the webhook notifications should be sent to in one of the following formats:

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

For details on the format of the data that HYCU sends to the specified URL, see “[Webhook data format](#)” below.

7. *Only if the receiving endpoint requires sender's identification.* From the Authentication Type drop-down menu, select one of the following authentication types:
 - **Basic authentication**, and then enter the user name and password associated with your webhook endpoint.
 - **Authentication by secret**, and then enter the secret to connect to your webhook endpoint.
8. Click **Next**.
9. *Optional.* Customize the request body that is sent by HYCU. You can click the appropriate fields in the HYCU fields list to easily insert event variables into the body.

! Important Make sure the format you define in the body is supported by the platform to which webhook notifications will be sent.

For details on the format of the webhook request body, see “[Webhook data format](#)” below.
10. Click **Save**.

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

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

Webhook data format

The webhook data format is defined by:

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

HTTP request headers

The request headers are sent in the following format:

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

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

HTTP request body

The request body is sent in the following format:

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

 **Note** Null values are ignored.

HTTP response code

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

Enabling the purge of events and jobs

You can configure HYCU to periodically delete events and/or jobs (as well as all associated job reports) that are no longer needed for daily business operations by enabling the purge of data from the HYCU database.

Prerequisite

You must be an infrastructure group administrator.

Consideration

Jobs related to backups, copies of backups, and archives will be deleted only if the corresponding restore points no longer exist or are expired.

Depending on whether you want to purge events or jobs, access one of the following panels:

- Accessing the Events panel

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

- Accessing the Jobs panel

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

Procedure

1. In the Events or Jobs panel, click  **Purge Configuration**.
2. Depending on your context, use the **Enable purging of events** or **Enable purging of jobs** switch.
3. Specify the number of years, months, weeks, or days to retain the data. Events or jobs that are older than the specified value will be purged. The maximum value is 99 years.
4. Click **Save** to start purging the HYCU database based on the specified value.

 **Important** This action cannot be undone. When your event or job data is deleted, you cannot retrieve it.

After you enable purging of events and/or jobs, you can at any later time edit the purge configuration or disable purging.

Using HYCU reports

HYCU reports provide you with a visual presentation of data protection environment resources and jobs. This comprehensive and precise presentation allows you to have an optimum view for analyzing data and therefore making the best 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, a heatmap, a line chart, an area chart, or a scatter chart.

Consideration

Keep in mind that your user group and user role determine what kind of report data you can view and what report actions you can perform.

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

- View reports. For details, see “[Viewing reports](#)” on page 309.
- Generate reports. For details, see “[Generating reports](#)” on page 310.
- Schedule reports. For details, see “[Scheduling reports](#)” on page 311.

 **Note** When scheduling the reports, you can also choose to send them by email.

- Export and import reports. For details, see “[Exporting and importing reports](#)” on page 312.

Accessing the Reports panel

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

Getting started with reporting

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

For a list of predefined reports, see “[Predefined reports](#)” below. For instructions on how to create reports, see “[Creating reports](#)” on the next page.

Predefined reports

The predefined reports represented by the  icon enable you to obtain reports on the key aspects of your data protection environment such as data transfer, job status, the number of backups, and the amount of protected data. These reports cannot be edited or deleted.

Predefined report	Description
Entity compliance status	List of virtual machines, applications, shares, and servers that are compliant and non-compliant with backup requirements.
Hourly activities per policy	List of assigned policies with the corresponding number of jobs that were running during each of the last 24 hours.
Hourly activities per target ^a	List of targets with the corresponding number of jobs that were running during each of the last 24 hours.

Predefined report	Description
Protected data	Total amount of protected data calculated on a daily basis.
Protected data per policy	Amount of data protected in the last 24 hours per policy.
Protected data per owner ^a	Total amount of protected data per owner.
Protected data per target ^a	Amount of the data protected in the last 24 hours per target.
Protected data timeline per target ^a	Daily amount of protected data per target.
Protected VM size per target ^a	List of protected virtual machine and servers, and distribution of the corresponding protected data between targets.
VM backup status	List of backups that occurred in the last 24 hours including information such as status and duration of backups, backup size, and so on.
VM backup status per target ^a	List of targets and related backups that occurred in the last 24 hours including information such as status and duration of backups, backup size, and so on.

^a Available only to an infrastructure group administrator.

Creating reports

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

Prerequisite

You have the Administrator user role assigned.

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

I want to...	Procedure
Create a new report	1. In the Reports panel, click  New . The New

I want to...	Procedure
from scratch.	<p>Report dialog box opens.</p> <ol style="list-style-type: none"> 2. Enter a report name and, optionally, its description. 3. Select the type of report. 4. Select the aggregation calculation to be used for your report. 5. Select the report tag for which the aggregation calculation should be performed. 6. Specify the time range for the report. You can select one of the predefined time ranges, or select Custom, and then use the calendar to select a start date and an end date of the time range. 7. Distribute the report tags for the collected data that you want to include in your report between x-axis and y-axis to determine how the collected data will be presented in the report. 8. Click Save.
Edit an existing report and save it as a new report.	<ol style="list-style-type: none"> 1. In the Reports panel, from the list of reports, select the one that you want to edit and save as a new report, and then click  Edit. The Report Configuration dialog box opens. 2. Enter a new name for the report, and then make the required modifications. 3. Click Save As.

You can later edit any of the created reports (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**). You cannot edit or delete the predefined reports represented by the  icon.

Viewing reports

You can view the reports on the current state of your data protection environment or the saved reports that were generated either manually or automatically.

I want to...	Procedure
View a report on the current state of my data protection environment.	In the Reports panel, from the list of reports, select the preferred report, and then click  Preview .
View a saved report.	<ol style="list-style-type: none"> 1. In the Reports panel, from the list of reports, select the preferred report. 2. In the Detail view that appears at the bottom of the screen, select the preferred report version, and then click  View. <p>For details on how to generate reports 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 do the following:

- Switch between the reports.
- Download the report in the PDF, PNG, or CSV format. To do so, click  **Download**, and then select one of the available formats.
- *For users with the Administrator user role assigned:* If you view a report on the current state of the data protection environment, you can save this version of the report by clicking **Generate**. The saved report is added to the list of report versions.

Generating reports

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

Prerequisite

You have the Administrator user role assigned.

Procedure

1. In the Reports panel, from the list of reports, select the one that you want to generate.

 **Note** If none of the available reports meets your reporting requirements, you can create a new report. For details, see “[Creating reports](#)” on page 308.

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

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

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

You can later do the following:

- View the saved reports. For details, see “[Viewing reports](#)” on page 309.
- Delete the saved reports that you do not need anymore. To do so, select the preferred report version, and then click  **Delete**.

Scheduling reports

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

Prerequisites

- You have the Administrator user role assigned.
- *For sending reports by email:* An SMTP server is configured. For details, see “[Configuring an SMTP server](#)” on page 407.

Procedure

1. In the Reports panel, from the list of reports, select the one that you want to be generated on a regular basis, and then click  **Scheduler**. The Report Scheduler dialog box opens.

 **Note** If none of the available reports meets your reporting requirements, you can create a new report. For details, see “[Creating reports](#)” on page 308.

2. In the Schedule date field, specify the date and the time of day when you want the report generation to begin.
3. From the Interval drop-down menu, select how often you want the reports to be generated (daily, weekly, or monthly).
4. Use the **Send report to email** switch if you want to schedule the automatic delivery of the reports to email recipients, and then do the following:
 - a. From the Report format drop-down menu, select a file format for your report (PDF, PNG, or CSV).
 - b. In the Email address field, enter one or more email recipients that should receive the reports. If you are entering more than one email address, make sure to press the spacebar after entering each one.
 - c. In the Email text field, enter the custom information to be included in the report email. For example, you can specify the purpose of the report, list further actions to be done by the recipient, or add contact information for additional report details.
5. Click **Save**.

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

You can later do the following:

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

Exporting and importing reports

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

 **Important** Your permissions determine what kind of reports you can view and edit, and therefore also define a different level of access to the reports, which you should consider before copying reports from one HYCU deployment to another.

Exporting reports

Procedure

In the Reports panel, from the list of all reports, select the one that you want to export, and then click  **Export**. The selected report will be exported to a JSON file and saved to the download location on your system.

Importing reports

Procedure

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

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

Viewing entity details

You can view the details about each virtual machine, discovered application, file share, server, and volume group in the Detail view of the Virtual Machines, Applications, Shares, or Volume Groups panel. The following details are available:

Summary	Shows detailed information about the selected entity.
Restore point	<p>You can view the following information about each restore point:</p> <ul style="list-style-type: none"> • Date and time when the restore point was created. • Tiers: <ul style="list-style-type: none"> ◦ BCKP Backup: Available by default unless a backup is expired. ▪ FULL Full: Visible if a full backup was performed. ▪ INCR Incremental: Visible if an incremental backup was performed. ◦ Archive: <ul style="list-style-type: none"> ▪ D ARCH Daily archive: Available if a daily data

	<p>archive was created.</p> <ul style="list-style-type: none"> ■ W ARCH Weekly archive: Available if a weekly data archive was created. ■ M ARCH Monthly archive: Available if a monthly data archive was created. ■ Y ARCH Yearly archive: Available if a yearly data archive was created. <p>By pausing on the archive icon, you can see the total number of data archives and the archive expiration time. If any of the archive jobs failed, the number of failed archive jobs is shown.</p> <ul style="list-style-type: none"> ○ COPY Copy: Available if a copy of backup data was created. By pausing on the icon, you can see the total number of backup copies and the backup copy expiration time. If any of the backup copy jobs failed, the number of failed backup copy jobs is shown. ○ SNAP Snapshot: Available if the source contains a local snapshot that enables you to perform a fast restore. By pausing on the icon, you can see whether the snapshot was recreated and its expiration time. <p>If any virtual disks were excluded from a backup, the corresponding tier label is marked with a line in the upper left corner. For example, FULL.</p> <p>Important If any of the tiers is colored red, it cannot be used for a restore.</p>
Compliance	<p>Shows the compliance status of an entity:</p> <ul style="list-style-type: none"> • ✓ Success • ✗ Failure • ? Undefined <p>An entity is considered to be compliant with backup requirements if the time since the last successful backup is lower than the RPO set in the HYCU policy and the estimated time to recover is lower than the RTO set in the HYCU policy.</p> <p>By pausing on the compliance status indicated by a respective</p>

	icon, additional information about the backup is available. You can see backup frequency, the elapsed time since the last successful backup, the time limit you set for the restore, and the estimated time required for the restore. In addition, if the compliance status of your entity is Failure, this list will also include a reason why it is not compliant.
Backup status	For details, see “ Viewing the backup status of entities ” below.
Restore status	Shows a progress bar indicating the progress of the entity restore. Tip If you double-click a progress bar, you are directed to the Jobs panel where you can check details about the related job.

Tip If there are too many items to be displayed on one page, you can move between the pages by clicking **>** and **<**. You can also use **▼** to set the number of items to be displayed per page.

Viewing the backup status of entities

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

Limitation

For virtual machines with attached volume groups: The Completed with errors backup status is available only for virtual machines that have volume groups attached directly.

Backup status of the entity	Restore a VM, a VG, or vDisks?	Restore VM files?	Restore an application?	Restore a file share?
✓ Completed successfully	✓	✓	✓	✓
! Completed with warnings	✓	✓	✓ ^a	✓
! Completed with errors	✓ ^b	✓ ^c	✓ ^d	✓ ^e

Backup status of the entity	Restore a VM, a VG, or vDisks?	Restore VM files?	Restore an application?	Restore a file share?
✖ Failed	✗	✗	✗	✗
○ Expired	✗	✗	✗	✗
?] Skipped ^f	✓	✓	✗	N/A

^a You cannot specify a point in time to which you want to restore data. This backup status may occur because disk mapping failed or a virtual machine does not have an NIC, or, in case of applications, at least one database log backup failed (whereas all other databases are in a consistent state).

^b Because not all virtual machine disk files were backed up successfully, the virtual machine can be partially restored. It may not be possible to turn it on if one of the system disks was not backed up.

^c Because not all virtual machine disk files were backed up successfully, the individual files can be partially restored (only the files that are displayed in the Restore Files dialog box).

^d An application can be partially restored (only the databases that are displayed in the respective restore dialog boxes).

^e Because not all files were backed up successfully, the file share can be partially restored. The files whose backup was unsuccessful are listed in the Job Report in their corresponding subtasks.

^f Applicable only for backups of passive nodes of failover clusters with shared storage.

Note By pausing on the backup status indicated by an icon, additional information about the backup is available. You can see the backup type, backup consistency, the duration and size of the backup, which target was used, and the backup UUID. For volume groups, you can also see if the volume group has been backed up both as part of the virtual machine backup and by assigning a policy directly to it.

If you double-click a backup status icon, you are directed to the Jobs panel where you can check details about the related jobs.

Tier statuses

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

Tier status	Restore an entity?
 or  (Done)	✓
 or  (Done with warnings)	✓ For details on what data can be restored if one of these backup statuses is shown, see “Viewing the backup status of entities” on page 315 .
 or  (Done with errors)	✗
 or  (Inaccessible on source)	✗
 or  (Deleted from the source)	✗
 or  (Failed)	✗
 or  (Expired)	✗

Filtering and sorting data

HYCU enables you to filter data in the panels so you can easily find what you need. You can apply two types of filters—the main view filter (to focus on certain aspects of your data protection environment) or the detail view filter (to focus on the information about the backup and restore data of the selected entity). After you apply any of the filters, only data that matches the filter criteria is displayed and you can easily find what you need.

In addition, to make it easier to work with the tables in the panels that have a large number of columns, you can also sort the data in ascending or descending order.

Depending on whether you want to filter or sort your data, see one of the following sections:

- “Filtering data in panels” below
- “Sorting data in panels” on page 324

Filtering data in panels

I want to apply...	Available in panels	Instructions
Main view filter	Applications, Virtual Machines, Volume Groups, Shares, Policies, Targets, Jobs, Events, and Self-Service	“Applying the main view filter” below
Detail view filter	Applications, Virtual Machines, Volume Groups, and Shares	“Applying the detail view filter” on page 323

Applying the main view filter

Apply the main view filter when you want to focus on certain aspects of your data protection environment (for example, filtering data in the Virtual Machines panel helps you to focus only on the virtual machines that you are interested in or responsible for).

 **Note** You can filter the items also by using the Search field on the left side of the panel. Typing text in this field automatically filters and displays only the matching items.

Procedure

1. In the selected panel, click  **Filters**.
2. In the side panel that opens, select your filter criteria.
3. Click **Apply Filters**.

For the details about the available main view filtering options, see one of the following sections:

- “Main view filtering options in the Applications panel” on the next page
- “Main view filtering options in the Virtual Machines panel” on the next page
- “Main view filtering options in the Volume Groups panel” on page 321
- “Main view filtering options in the Shares panel” on page 321
- “Main view filtering option in the Policies panel” on page 322
- “Main view filtering options in the Targets panel” on page 322

- “Main view filtering options in the Jobs panel” on page 322
- “Main view filtering options in the Events panel” on page 323
- “Main view filtering option in the Self-Service panel” on page 323

Main view filtering options in the Applications panel

Filtering option	Filter applications by...
Sources	Sources that host the virtual machines on which the applications are running.
Policy assignment	Policies assigned to the applications (Unassigned, Assigned, and/or Specific policies). Note If you filter applications by the Assigned option, the ones to which the Exclude policy is assigned will not be listed.
Owners	Owners that are assigned to the virtual machines on which the applications are running.
Application types	Application types.
Compliance	Compliance statuses of the applications: <ul style="list-style-type: none"> Success Failure Undefined: The Exclude policy is assigned to the applications or the applications do not have a policy assigned.
Protection	Protection statuses of the applications (Protected, Unprotected, and/or Protected deleted).
Discovery	Discovery statuses of the applications: <ul style="list-style-type: none"> Success Failure Warning: Application discovery failed because the virtual machine is offline or is not reachable.

Main view filtering options in the Virtual Machines panel

Filtering option	Filter virtual machines by...
Sources	Sources that host the virtual machines.

Filtering option	Filter virtual machines by...
Credential groups	Credential groups assigned to the virtual machines.
Policy assignment	Policies assigned to the virtual machines (Unassigned, Assigned, and/or Specific policies). <p>Note If you filter virtual machines by the Assigned option, the ones to which the Exclude policy is assigned will not be listed.</p>
Validation policy assignment	Validation policies assigned to the virtual machines (Unassigned, Assigned, and/or Specific policies).
Owners	Owners assigned to the virtual machines.
Compliance	Compliance statuses of the virtual machines: <ul style="list-style-type: none"> Success Failure Undefined: The Exclude policy is assigned to the virtual machines or the virtual machines do not have a policy assigned.
Discovery	Discovery statuses of the applications running on the virtual machines: <ul style="list-style-type: none"> Success Failure Warning: Application discovery failed because the virtual is offline or is not reachable. Undefined: The information about the application discovery status is not available.
Protection	Protection statuses of the virtual machines (Protected, Unprotected, and/or Protected deleted).
Validation status	Backup validation statuses of the virtual machines.
Disaster recovery readiness	DR readiness statuses of the virtual machines.

Main view filtering options in the Volume Groups panel

Filtering option	Filter volume groups by...
Sources	Sources that host the volume groups.
Policy assignment	Policies assigned to the volume groups (Unassigned, Assigned, and/or Specific policies). Note If you filter volume groups by the Assigned option, the ones to which the Exclude policy is assigned will not be listed.
Owners	Owners assigned to the volume groups.
Compliance	Compliance statuses of the volume groups: <ul style="list-style-type: none">• Success• Failure• Undefined: The Exclude policy is assigned to the volume groups or the volume groups do not have a policy assigned.
Protection	Protection statuses of the volume groups (Protected, Unprotected, and/or Protected deleted).

Main view filtering options in the Shares panel

Filtering option	Filter file shares by...
Sources	Sources that host the file shares.
Protocol	Protocols of the file shares (SMB, NFS, or S3).
Policy assignment	Policies assigned to the file shares (Unassigned, Assigned, and/or Specific policies). Note If you filter file shares by the Assigned option, the ones to which the Exclude policy is assigned will not be listed.
Owners	Owners assigned to the file shares.
Compliance	Compliance statuses of the file shares: <ul style="list-style-type: none">• Success• Failure• Undefined: The Exclude policy is assigned to the file

Filtering option	Filter file shares by...
	shares or the file shares do not have a policy assigned.
Protection	Protection statuses of the file shares (Protected, Unprotected, and/or Protected deleted).
Incremental forever backup	Enabled or disabled Incremental forever backup option.

Main view filtering option in the Policies panel

Filtering option	Filter policies by...
Compliance	<p>Compliance statuses of the policies:</p> <ul style="list-style-type: none"> Success: All entities to which the policies are assigned are compliant with the policy settings. Failure: Not all entities to which the policies are assigned are compliant with the policy settings. Undefined: The Exclude policy is assigned to the entities or the entities do not have a policy assigned.

Main view filtering options in the Targets panel

Filtering option	Filter targets by...
Target type	Target types.
Health	Health statuses of the targets (OK, Warning, Error, and/or Undefined).

Main view filtering options in the Jobs panel

Filtering option	Filter jobs by...
Status	Job statuses (OK, Warning, Queued, Executing, Aborted, and/or Failed).
Time range	Time ranges: You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, and/or Last week), or use the calendar to select a start date and hour and an end date and hour of the time range for the jobs to be displayed.

Main view filtering options in the Events panel

Filtering option	Filter events by...
Category	Event categories.
Username	User names to include only the events started by the selected users and/or user groups.
Status	Event statuses (Success, Warning, and/or Failed).
Time range	Time ranges: You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, and/or Last week), or use the calendar to select a start date and hour and an end date and hour of the time range for the events to be displayed.

Main view filtering option in the Self-Service panel

Filtering option	Filter users and/or user groups by...
Status	User or user group statuses—which users or user groups are allowed to sign in to HYCU and which are not (Active or Inactive).

Applying the detail view filter

Apply the detail view filter when you want to focus on the information about the backup and restore data of the selected entity.

Procedure

1. From the list of all entities in the selected panel, select the entity that you want to filter by backup and restore data.
2. In the Detail view that appears at the bottom of the screen, click  **Filters - Detail View**.
3. In the side panel that opens, select your filter criteria.
4. Click **Apply Filters**.

 **Tip** If there are too many filtered items to be displayed on one page, you can move between the pages by clicking  and . You can also use  to set the number of filtered items to be displayed per page.

The following detail view filtering options are available:

Filtering option	Filter entities by...
Tiers	Restore point tiers of the entities (Archive daily, Archive weekly, Archive monthly, Archive yearly, Backup, Copy, and/or Snapshot).
Backup type	Entity backup types (Incremental or Full).
Restore point date	Time when the entity restore points were created. You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, and/or Last week), or use the calendar to select a start date and hour and an end date and hour of the time range within which the restore points were created.
Backup status	Entity backup statuses (Success, Failure, Expired, and/or Warning).
Compliance	Compliance statuses of the entities: <ul style="list-style-type: none"> • Success • Failure • Undefined: The Exclude policy is assigned to the entities or the entities do not have a policy assigned.

Sorting data in panels

Procedure

1. In the selected panel, click the column header of the column that you want to sort. The column is sorted in ascending order, which is indicated by the  icon.
2. Click the column header again to sort the data in descending order, which is indicated by the  icon.

Exporting the contents of the panel

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

Consideration

If you want to export only specific data, click  **Filters**, select your filter criteria based on what kind of data you want to export to a file, and then click **Apply Filters**.

Procedure

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

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

Managing targets

If you have the proper permissions, you can view target information, edit target properties, activate or deactivate a target, or delete a target if you do not want to use it for storing protected data anymore.

Accessing the Targets panel

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

Viewing target information

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

Target information	Description
Name	Name of the target.

Target information	Description
Type	<p>Type of target (NFS, SMB, Nutanix, Nutanix Objects, iSCSI, Amazon S3 / S3 Compatible, AZURE, Google Cloud, QStar NFS, QStar SMB, or Data Domain).</p> <p>Note A tape target is represented by the  icon, and a target that has WORM enabled is represented by the  icon.</p>
Health	<p>Health status of the target:</p> <ul style="list-style-type: none"> Gray: Shows the initial target status before a health test. It also indicates an inactive target. Green: The target is in a healthy state with target utilization of less than the configured value (by default, 90%). Yellow: Target utilization is over the configured value (by default, 90%). Red: Target utilization is over the configured value (by default, 95%). It also indicates a target error state after a test task (for example, an I/O error occurred, the target is not accessible, the permission is denied, and so on).
	<p>HYCU calculates if there is enough space on the target for storing backup data based on the following:</p> <ul style="list-style-type: none"> <i>If no previous backup is stored on the target:</i> The total provisioned space of all disks included in the virtual machine backup, regardless of whether the backup is full or incremental. <i>If a previous backup is stored on the target:</i> The size of the last incremental backup for incremental backups, or the size of the last full backup for full backups or incremental backups if no previous incremental backup exists.
Size	<p>Estimation of the amount of storage space that should be reserved for the backup files (in MiB, GiB, or TiB).</p>
Utilization	<p>Percentage of the specified target size that is already used for storing protected data.</p>

Target information	Description
Mode	<p>Mode of the target:</p> <ul style="list-style-type: none"> • Read/Write: You can use this target for backing up and restoring data. • Read Only: You can use this target only for restoring data. <p>ⓘ Important The Read-Only mode is automatically set on an imported target to prevent you from performing backups. Make sure not to change the mode of the imported targets.</p>
Status	<p>Status of the target:</p> <ul style="list-style-type: none"> • Active: You can use this target for backing up and restoring data. • Inactive: You cannot use this target for backing up and restoring data. This status indicates that the target is deactivated due to maintenance tasks (for example, adding new disks). <p>For details on how to change the status of the target, see “Activating or deactivating a target” on page 329.</p>

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

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

Editing a target

⚠ Caution Making any changes to the target location may result in data loss. Therefore, before specifying a new target location, make sure you have already moved the existing backup data to this new location on the same or a different server.

Prerequisites

Only if you plan to increase the size of an iSCSI target.

- The size of the target must be increased on the iSCSI server.
- No backup or restore job may be in progress on the selected target.
- No other maintenance task may be already running on the selected target (such as editing the target and updating the iSCSI Initiator secret or resetting mutual CHAP authentication sessions for the targets with CHAP authentication enabled).
- No other size increase of the selected target may be started.

Considerations

- If you change the target settings in the policy assigned to the HYCU backup controller, make sure to update the note of the target's configuration.
- *For QStar tape targets:* If the status of the Integral Volume set is offline, the corresponding tape target is automatically deactivated in HYCU. When the Integral Volume set is remounted in QStar, make sure to activate the target. For details on how to do this, see “Activating or deactivating a target” on the [next page](#).

Procedure

1. In the Targets panel, select the target that you want to edit, and then click  **Edit**. The Edit Target dialog box opens.
2. Edit the selected target as required. For detailed information about target properties, see “[Setting up targets](#)” on page 78.

ⓘ Important If you want to make specific changes to the NFS, SMB, Nutanix, iSCSI, or tape target, make sure you first detach the storage. For a list of possible changes and instructions, see “[Detaching storage and changing target data](#)” on the [next page](#).
3. *Only if you are editing an iSCSI target and you want to increase its size.* Select the **Extend target** check box.
4. Click **Save**.

After you edit the target, you will receive a message indicating whether the target was edited successfully. If you increased the size of the iSCSI target, you will also receive a message indicating whether the increase of the target size completed successfully.

Detaching storage and changing target data

If you want to change data for the NFS, SMB, Nutanix, iSCSI, or tape target, make sure that the storage is first detached from the HYCU backup controller to be able to perform the required changes:

Target type	Possible changes
NFS	Server name, IP address, or path to the shared folder
SMB	Server name, IP address, or path to the shared folder
Nutanix	URL
iSCSI	Portal IP address
Tape (QStar NFS and QStar SMB)	Web service endpoint

Procedure

1. Deactivate the target and detach the storage from the HYCU backup controller as described in “[Activating or deactivating a target](#)” below.
2. Make the required changes first on the server where the target is located, and then also in the HYCU web user interface as described in “[Editing a target](#)” on page 327.
3. Activate the target as described in “[Activating or deactivating a target](#)” below.

Activating or deactivating a target

Procedure

1. In the Targets panel, select the target that you want to activate or deactivate.
2. Change the status of the selected target by clicking **Activate** or **Deactivate**.
3. If you are deactivating the target to change the data related to the NFS, SMB, Nutanix, iSCSI, or tape target, enable the **Detach storage** switch. For details on detaching storage from the HYCU backup controller, see “[Detaching storage and changing target data](#)” above.
4. *For target deactivation:* Click **Deactivate** to confirm that you want to deactivate the selected target.

If you deactivate a target, this target will not be used for backup and restore operations anymore.

Deleting a target

You can delete a target if it does not contain protected data. After deleting a target, no backup or restore actions including this target are possible anymore.

Procedure

1. In the Targets panel, select the target that you want to delete, and then click  **Delete**.

 **Note** If the target that you want to delete is used for archiving, make sure that no data archive with the specified archive target is used by any policy.
2. Click **Delete** to confirm that you want to delete the selected target.

Managing policies

If you have the proper permissions, you can view policy information, edit policy properties, or delete a policy if you do not want to use it for protecting data anymore.

Consideration

You cannot delete the Exclude policy.

Accessing the Policies panel

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

Viewing policy information

You can view information about each policy in the list of policies in the Policies panel. This allows you to have an overview of the general status of the policies.

Consideration

The values for the backup RPO, RTO, and retention period that are defined in a policy are rounded to days, weeks, months, or years for display, but are stored

and used internally as defined. For example, 30 days are rounded to one month in the HYCU web user interface.

The following information is available for each policy:

Policy information	Description
Name	Name of the policy.
Compliance	<p>Compliance status of the policy:</p> <ul style="list-style-type: none"> ✓ Success ✗ Failure ⌚ Undefined <p>A policy is considered compliant if all entities to which this policy is assigned are compliant with the policy settings. For detailed information about the compliance status of entities, see “Viewing entity details” on page 313.</p>
VM Count	Total number of virtual machines and servers that have the particular policy assigned to them.
App Count	Total number of applications that have the particular policy assigned to them.
Description	Description of the policy (how often backup and restore jobs are performed).

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

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

Editing a policy

Limitations

- If editing a policy that is assigned to the HYCU backup controller, you cannot select the Backup from replica policy option because HYCU does not support backing up the HYCU backup controller from a replica in the remote office/branch office (ROBO) environment.

- *For vSphere environments:* When editing a policy that is assigned to a virtual machine or an application, the following limitations apply:
 - You cannot enable the Backup from replica option.
 - You can enable the Fast restore option or select Snapshot as the backup target type only if the virtual machine is residing on a vVols or vSAN datastore.

Consideration

If you edit a policy in such a way that you enable the Copy option, the next backup of the virtual machines and volume groups to which this policy is assigned will be a full backup.

Procedure

1. In the Policies panel, select the policy that you want to edit, and then click  **Edit**. The Edit Policy dialog box opens.
2. Edit the selected policy as required. For detailed information about policy properties, see [“Creating a policy” on page 110](#).
3. Click **Save**.

Deleting a policy

Considerations

- A policy that is assigned to one or more entities for which backups are scheduled cannot be deleted. If you want to delete such a policy, you must first abort the scheduled backups. For details on how to abort queued jobs, see [“Managing HYCU jobs” on page 300](#).
- If you delete a policy that is assigned to one or more entities, keep in mind that no further backups will be performed for these entities.

Procedure

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

Performing a manual backup

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

Prerequisite

Only if backing up a volume group manually. Make sure a policy is assigned directly to the volume group. If the policy is assigned only to the virtual machine to which the volume group is attached, performing a manual backup for the selected volume group is not possible.

Considerations

- You can prevent your manual backups from interfering with the scheduled backups determined by the RPO specified in the policy. To do so, set the `exclude.manually.run.backups.regarding.rpo` configuration setting to `true`. This is especially important if you define backup windows because performing a manual backup can prevent the backup scheduled in the backup window from starting, which can result in data not being protected until the next backup window or the next manual backup. For details on how to customize HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on page 523.
- *Only if you enabled the Incremental forever backup option for a file share.* Make sure not to enable the **Force full backup** switch for such a file share. Otherwise, a full backup will be performed instead of an incremental one.

Procedure

1. In the Applications, Virtual Machines, Volume Groups, or Shares panel, select which entities you want to back up.
2. Click  **Backup** to perform the backup of the selected entities.
3. *Only if you selected Target as the backup target type in your policy.* Enable the **Force full backup** switch if you want to perform a full backup. Otherwise, HYCU will perform a full or incremental backup based on the settings defined in your policy.
4. Click **Yes** to confirm that you want to start the manual backup.

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

Setting up a validation policy

As an alternative to manually performing the backup validation for a virtual machine and verifying that the virtual machine has no corrupted backups, you can set up a validation policy and schedule the backup validation according to the values that you define in your validation policy. For details on how to validate the virtual machine backup by creating a virtual machine clone, see “[Validating the virtual machine backup](#)” on page 188.

① Important HYCU automatically creates a clone of the virtual machine while performing the backup validation.

Prerequisites

- If you plan to select a vSphere storage container for the virtual machine copy, the latest version of VMware Tools must be installed on the virtual machine.
- *Only if you plan to specify the Advanced validation type.*
 - Credentials must be assigned to the virtual machine. For prerequisites, limitations, considerations, and instructions, see “[Enabling access to application data](#)” on page 209.
 - A network card must be added to the virtual machine.

Limitation

Performing the backup validation is not supported for the following:

- The HYCU backup controller
- AWS GovCloud (US) environments
- Azure Government environments

Considerations

- Network conflicts may occur during the backup validation if the virtual machine is configured with a static IP address, resulting in unreliable backup validation data.
- *Only if you plan to specify the Advanced validation type when performing the backup validation for a Windows virtual machine.* Checking for disk errors may fail in some cases, which does not mean that your virtual machine is corrupted. However, it is highly recommended that you check the status of such a virtual machine manually.

- After the backup validation is performed, consider the following:
 - You can view the backup validation status of a virtual machine in the Validation column in the Virtual Machines panel (represented by an icon). By pausing on the icon, you can also see which validation policy is assigned to the virtual machine.
 - The Exclude policy is automatically assigned to the cloned virtual machine.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, select one or more virtual machines for which you want the backup validation to be performed.

 **Tip** You can update the list of virtual machines by clicking  **Refresh**. To narrow down the list of displayed virtual machines, you can use the filtering options described in “[Filtering and sorting data](#)” on [page 317](#).

2. Click  **Validation**. The Validation Policies dialog box opens.
3. Click  **New**.
4. Enter a name for your validation policy and, optionally, its description.
5. From the Storage container drop-down menu, select where you want to clone the virtual machine for which you are performing the backup validation.
6. From the Restore from drop-down menu, select which tier you want to use for the backup validation. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**
 - **Backup**
 - **Copy**
 - **Archive**

 **Note** If you select Automatic, the tier for the backup validation is by default selected in the following priority order: Backup > Copy > Archive. This means that HYCU will always use the first available tier in the specified order for the backup validation. However, you can at any time

change this default behavior by customizing the `backup.validation.restore.source.priority.order` configuration setting in the HYCU `config.properties` file and adjusting the tier order to your data protection needs. For details on how to customize HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on [page 523](#).

7. From the Keep VM after validation drop-down menu, depending on whether you want to keep the virtual machine after the backup validation is performed, select one of the following options:

Option	Description
Always	The virtual machine will be kept after the backup validation is performed.
On validation error	The virtual machine will be kept after the backup validation is performed only if a validation error occurs during the validation.
Never	The virtual machine will be automatically deleted after the backup validation is performed.

8. From the Validation type drop-down menu, select one of the following types:

Validation type	Description
Basic	During the backup validation, the following tasks will be performed: <ul style="list-style-type: none"> • The virtual machine will be cloned and turned on. • The guest OS will be shut down.
Advanced	During the backup validation, the following tasks will be performed: <ul style="list-style-type: none"> • The virtual machine will be cloned and turned on. • Any applications running on the virtual machine will be discovered. • Virtual disks will be validated, which includes checking the virtual machine file system and existing disks on the virtual machine. For Windows virtual machines,

Validation type	Description
	<p>checking for disk errors is also performed.</p> <ul style="list-style-type: none"> • The custom scripts will be run, if specified. • The guest OS will be shut down.

9. *Only if you selected the Advanced validation type.* Do the following:
 - a. Enable the **Run custom script** switch if you want the custom script to be run on the virtual machine as part of the backup validation process, and then make sure that the proper path to the script is specified.

Note The script returns an exit code of 0 for success and any other value for failure.
 - b. From the Network drop-down menu, select the network for the virtual machine.
10. Click **Next**.
11. Depending on whether you want backup validation for the virtual machine to be performed on a daily, weekly, monthly, and/or yearly basis, add any of the preferred backup validation options to the list of the enabled options by clicking it:
 - **Daily**
 - **Weekly**
 - **Monthly**
 - **Yearly**
12. In the Start at fields, specify the hour and the minute when the backup validation job should start.
13. From the Time zone drop-down menu, select the appropriate time zone for the backup validation job.

Note All backup validation jobs are by default started based on the HYCU backup controller time zone.
14. Depending on the selected backup validation options, specify at what intervals you want backup validation to be performed:

Backup validation option	Instructions
Daily	<ol style="list-style-type: none"> a. In the Recur every field, specify whether you want

Backup validation option	Instructions
	<p>backup validation to be performed every day or every few days.</p> <p>b. Use the Apply only on weekdays switch if you want backup validation to be performed only on weekdays.</p>
Weekly	<p>a. In the Recur every field, specify whether you want backup validation to be performed every week or every few weeks.</p> <p>b. Select one or more days of the week on which you want backup validation to be performed.</p>
Monthly	<p>a. In the Recur every field, specify whether you want backup validation to be performed every month or every few months.</p> <p>b. Select whether you want backup validation to be performed on the same day of the month (for example, on the fifth day of the month), or on a specific day of the month (for example, on the second Friday of the month).</p>
Yearly	<p>a. In the Recur every field, specify whether you want backup validation to be performed every year or every few years.</p> <p>b. Select whether you want backup validation to be performed on the same day of the preferred month (for example, on the fifth day of January), or on a specific day of the preferred month (for example, on the second Friday of April).</p>

15. Click **Save**.

16. Click **Assign**.

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

Archiving data manually

HYCU archives your data automatically after you enable the Archiving policy option. However, you can archive data manually at any time (for example, if you want to archive data for a specific restore point or if an archiving job fails).

Prerequisites

- You must have the Administrator, Backup and Restore Operator, or Backup Operator user role assigned.
- The Archiving option must be specified in the assigned policy and a data archive must be created.

Considerations

- Retention time for archives is calculated from the date and time when the restore point for the entity whose data you are archiving was created.
- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for archiving data manually.
- *Only if you selected Snapshot as the backup target type in your policy.* The configuration settings that HYCU uses for archiving are the ones that the virtual machine has at the time when archiving starts.

Depending on the type of data that you want to archive, access one of the following panels:

- Accessing the Applications panel
To access the Applications panel, in the navigation pane, click  **Applications**.
- Accessing the Virtual Machines panel
To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.
- Accessing the Shares panel
To access the Shares panel, in the navigation pane, click  **Shares**.
- Accessing the Volume Groups panel

To access the Volume Groups panel, in the navigation pane, click  **Volume Groups**.

Procedure

1. In the Applications, Virtual Machines, Shares, or Volume Groups panel, click the entity whose data you want to archive.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

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

3. Click  **Run Archiving**. The Run Archiving dialog box opens.
4. Select the preferred archiving option.
5. Click **Run**.

Recreating snapshots

Recreating snapshots is required in the following scenarios:

- If you plan to restore files from a snapshot (and not directly from a target) and no snapshot is available for the selected virtual machine restore point.
- If you plan to restore applications, export virtual disks, or restore files that are stored in the archive storage tier on an Azure target.

Limitation

Recreating snapshots is not supported for AWS GovCloud (US) and Azure Government environments.

Consideration

If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for recreating snapshots.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, select the virtual machine whose snapshot you want to recreate.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

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

3. Click  **Recreate Snapshot**. The Recreate Snapshot dialog box opens.
4. From the Storage container drop-down menu, select where you want to recreate the snapshot.
5. From the Restore from drop-down menu, select which tier you want to use for recreating the snapshot. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest snapshot creation.
 - **Backup**
 - **Copy**
 - **Archive**
6. Click **Recreate**.

Adjusting the HYCU backup controller resources

When storage, vCPU, or memory utilization is exceeded (that is, when the utilization of any of these resources is greater than 90 percent), their values that are indicated by circles become red in the HYCU Controller widget in the Dashboard panel. You can adjust these resources to better meet your data protection goals.

Procedure

1. Depending on the environment in which the HYCU backup controller resides, sign in to the Nutanix Prism web console, the vSphere Web Client or the vSphere Client, the AWS GovCloud (US) console, the Google Cloud console, or the Azure Government portal.
2. Shut down the HYCU backup controller.

3. Modify the storage, vCPU, or memory configuration as required.
4. Turn on the HYCU backup controller.

For instructions on how to perform these steps, see Nutanix, VMware, AWS, Google Cloud, or Azure documentation.

Customizing your HYCU web user interface

The HYCU web user interface is designed to be customized to match your needs. When customizing your HYCU web user interface, you can do the following:

- Adjust the table density to determine how close or far apart the rows in the tables should be.
- Choose to show or hide the separator line between the rows in the tables.
- Switch your HYCU web user interface to light or dark mode. HYCU by default uses the color mode exposed by your browser.

Accessing the Customization dialog box

To access the Customization dialog box, click  at the upper right of the screen, and then select **Customization**.

Procedure

1. In the Customization dialog box, do the following:
 - Under Table density, select **Default density** or **High density** depending on how close or far apart you want the rows in the tables to be.
 - Enable the **Row dividers** switch if you want to show the separator line between the rows in the tables.
 - Do one of the following:
 - *If you want to use the dark mode:* Enable the **Dark mode** switch.
 - *If you want to use the light mode:* Disable the **Dark mode** switch.
2. Click **Close**.

The changes take place immediately without the need to sign out and sign in again to the HYCU web user interface. The preferred customization is remembered for the next time you sign in to HYCU.

Chapter 10

Managing users

The HYCU user management system provides security mechanisms to help prevent unauthorized users from accessing protected data. Only users that are given specific rights have access to the data protection environment. These users can be authenticated either by HYCU or any of the supported identity providers. For details on identity providers, see “[Integrating HYCU with identity providers](#)” on page 373.

Each user that signs in to HYCU must belong to one of the HYCU groups—an infrastructure group or a self-service group—and have a user role assigned.

For details on HYCU groups and user roles, see “[HYCU groups](#)” below and “[User roles](#)” on page 345.

 **Note** User management concepts and procedures apply to both virtual machines and servers.

HYCU groups

For a consolidated user management experience, HYCU provides two types of groups to which users can belong.

Group	Description
Infrastructure group	<p>Created by default during the deployment of the HYCU virtual appliance and already includes a built-in user with the Administrator user role assigned (represented by )—cannot be edited, deactivated, and deleted.</p> <p>Users can be added to this group by an infrastructure group administrator (an infrastructure group user with the Administrator user role assigned).</p>
Self-service	Must be created by an infrastructure group administrator

Group	Description
group	<p>and represents a customer or a department responsible for a specific set of entities in the data protection environment.</p> <p>Users can be added to this group by an infrastructure group administrator.</p> <p>Important If a specific self-service group is deleted, all data that is backed up by this group is deleted from the database.</p>

You can manage users only if you have an Administrator role assigned. However, keep in mind that the scope of user management actions that you can perform differs depending on whether you belong to the infrastructure or self-service group. As an infrastructure group administrator, you can manage users and groups throughout the whole data protection environment, whereas as a self-service group administrator, you can manage only the group you belong to. The following diagram shows which user-related actions you can perform:

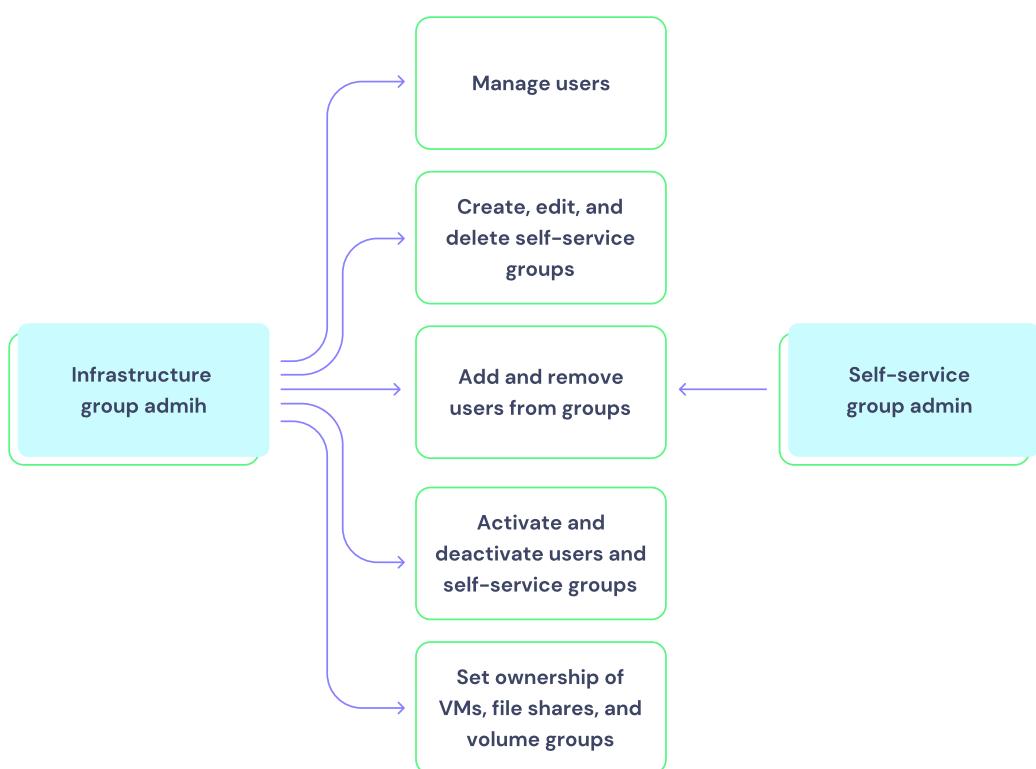


Figure 10-1: User management actions performed by the infrastructure and self-service group administrators

Depending on the HYCU group to which you belong and the assigned user role, you can perform only specific actions in the data protection environment. For details on user roles, see “[User roles](#)” below.

User roles

Each user in a group has an assigned role that determines the scope of actions the user can perform in the data protection environment. This means that access to data and information within the data protection environment is limited based on the role that the user has assigned. If a user is a member of multiple groups, this user can have different roles assigned in different groups, depending on the business needs, and can switch between these groups while being signed in to HYCU.

Depending on the group to which a user belongs, the user can perform the following actions:

Role	Infrastructure group	Self-service group
Administrator	<ul style="list-style-type: none"> Perform all actions in the data protection environment. 	<ul style="list-style-type: none"> Assign policies. Back up and restore virtual machines, applications, file shares, and volume groups. Manage data retention. Perform virtual machine backup validation by using the Validate VM backup option. Assign and unassign validation policies. Add and remove users from groups. Perform all report management actions. Add, edit, and remove cloud accounts.

Role	Infrastructure group	Self-service group
Viewer	<ul style="list-style-type: none"> View information about applications, virtual machines, file shares, volume groups, policies, targets, jobs, events, users, generated report versions, and settings available through the  Administration menu in the data protection environment. 	<ul style="list-style-type: none"> View information about applications, virtual machines, file shares, volume groups, policies, jobs, events, and generated report versions in the data protection environment.
Backup Operator	<ul style="list-style-type: none"> View the same information as Viewer. Define a backup strategy. Back up virtual machines, file shares, and volume groups that are not owned by any self-service group, and back up applications. 	<ul style="list-style-type: none"> View the same information as Viewer. Assign policies. Back up virtual machines, applications, file shares, and volume groups.
Restore Operator	<ul style="list-style-type: none"> View the same information as Viewer. Restore virtual machines, file shares, and volume groups that are not owned by any self-service group, and restore applications. Perform virtual machine backup validation by using the Validate VM backup option. Assign and unassign validation policies. 	<ul style="list-style-type: none"> View the same information as Viewer. Restore virtual machines, applications, file shares, and volume groups. Perform virtual machine backup validation by using the Validate VM backup option. Assign and unassign validation policies.
Backup and Restore Operator	<ul style="list-style-type: none"> View the same information as Viewer. Define a backup strategy. 	<ul style="list-style-type: none"> View the same information as Viewer. Assign policies.

Role	Infrastructure group	Self-service group
	<ul style="list-style-type: none"> • Back up and restore virtual machines, file shares, and volume groups that are not owned by any self-service group, and back up and restore applications. • Perform virtual machine backup validation by using the Validate VM backup option. • Assign and unassign validation policies. 	<ul style="list-style-type: none"> • Back up and restore virtual machines, applications, file shares, and volume groups. • Perform virtual machine backup validation by using the Validate VM backup option. • Assign and unassign validation policies.

Setting up a user environment

Before users can start using HYCU for data protection, you must give them rights to access data within the data protection environment. By creating a user and adding the user to a group, you allow the user to access only the defined data protection environment and to perform a set of actions specified by the assigned role:

Task	Performed by...	Instructions
1. Create a new user.	An infrastructure group administrator	“Creating a user” on the next page
2. Add a user to a user group.	An infrastructure or a self-service group administrator	“Adding a user to a group” on page 352

While setting up a user environment, you can tailor it to the user's needs by performing one or more of the following tasks:

Task	Performed by...	Instructions
Create a new self-service group.	An infrastructure group administrator	“Creating a self-service group” on page 353

Task	Performed by...	Instructions
Set ownership of virtual machines, file shares, and volume groups.	An infrastructure group administrator	“Setting ownership” on page 354
Enable or disable specific groups or users from signing in to HYCU.	An infrastructure group administrator	“Activating or deactivating users or self-service groups” on page 358

Accessing the Self-Service panel

To access the Self-Service panel, in the navigation pane, click  **Self-Service**.

Creating a user

Prerequisites

- *For using two-factor authentication:* An appropriate authenticator must be set up. Depending on the authentication method:
 - A time-based one-time password (OTP) authentication application, such as Google Authenticator on your mobile phone.
 - A FIDO-compatible authenticator, such as a hardware key, fingerprint reader, or similar.
- *For integrating HYCU with identity providers:* In an identity provider environment, HYCU must be assigned as an application to users for whom you want to enable signing in to HYCU by using the identity provider. For detailed instructions on how to integrate HYCU with identity providers, see [“Integrating HYCU with identity providers” on page 373](#).

Limitations

- You cannot add the Active Directory primary group (usually the Domain Users group) as an AD group.
- If certificate authentication is enabled, setting up two-factor authentication for AD users is not supported.

Consideration

The members of an identity provider group (Active Directory or OpenID Connect (OIDC)) are listed as individual users. For AD users, this allows you to

enable two-factor authentication and set the preferred language for each user.

Procedure

1. In the Self-Service panel, click  **Manage Users**, and then click  **New**.
2. Depending on what kind of user you are adding, enter one of the following:
 - *For a HYCU user, an AD user, an OIDC user, or an OIDC group:* User name
 - *For an AD group:* Common name
3. From the Authentication Type drop-down menu, select one of the following authentication types, and then follow the instructions:

Authentication type	Instructions
HYCU	<ol style="list-style-type: none"> From the Language drop-down menu, select the preferred language for the user. In the Name field, enter a display name for the user. <i>Optional.</i> In the Email field, enter the email address of the user. In the Password field, enter the user password. <p> Note The minimum password length is six characters.</p>
OIDC User	<ol style="list-style-type: none"> From the Language drop-down menu, select the preferred language for the user. From the Identity Provider drop-down menu, select the identity provider. In the Identity Provider User ID field, enter the

Authentication type	Instructions
	<p>ID of the identity provider user.</p> <p>Note Depending on your identity provider, the user ID corresponds to the following:</p> <ul style="list-style-type: none"> <i>Active Directory Federation Services</i>: Object GUID <i>Google</i>: User email address <i>Keycloak</i>: User ID <i>Microsoft</i>: Object ID <i>Okta</i>: Part of the URL when you navigate to the user's profile <p>For details, see the respective identity provider documentation.</p>
OIDC Group	<ol style="list-style-type: none"> From the Language drop-down menu, select the preferred language for the group. From the Identity Provider drop-down menu, select the identity provider. In the Identity Provider Group ID field, enter the ID of the identity provider group. <p>Note Depending on your identity provider, the group ID corresponds to the following:</p> <ul style="list-style-type: none"> <i>Active Directory Federation Services</i>: Object GUID <i>Keycloak</i>: Group ID <i>Microsoft</i>: Group Object IDs <i>Okta</i>: Group name <p>For details, see the respective identity provider documentation.</p>
AD User	<ol style="list-style-type: none"> From the Language drop-down menu, select the preferred language for the user. From the Identity Provider drop-down menu, select the Active Directory the AD user belongs

Authentication type	Instructions
	to.
AD Group	<p>a. From the Language drop-down menu, select the preferred language for the user.</p> <p>b. From the Identity Provider drop-down menu, select the Active Directory the AD group belongs to.</p>

4. *Only if you are adding a HYCU user, an AD user, or an AD group.* Use the **Two-factor authentication** switch if you want to enable two-factor authentication for the user, and then select one of the following two-factor authentication methods:

- **Time-based one-time password**

This option enables the use of a time-based one-time password (OTP) generated by an OTP application. The user needs to set up an OTP during the first sign-in after two-factor authentication is enabled.

- **FIDO**

This option enables the use of an authenticator complying with FIDO protocols (FIDO authenticator). The user needs to register a FIDO authenticator. For details, see “[Managing FIDO authenticators](#)” on [page 450](#).

5. *Only if you enabled two-factor authentication.* To prevent the user from disabling two-factor authentication, make sure the **User cannot disable two-factor authentication** check box is selected. If you clear the check box, the user can disable two-factor authentication. An infrastructure group administrator can disable two-factor authentication even if this option is enabled.

 **Note** If a user disables two-factor authentication, the administrator is notified with a security warning.

6. Click **Next**.

7. *Only if you want to add the user to a user group.* Do the following:

a. From the User Group drop-down menu, select the user group to which you want to add the user.

 **Important** You can add the user to the user group also later by following the procedure described in “[Adding a user to a group](#)” on [page 450](#).

the next page. In any case, keep in mind that the user must be added to at least one user group before they can sign in to HYCU.

- b. From the User Role drop-down menu, select the role that you want to assign to the user (**Administrator**, **Backup and Restore Operator**, **Restore Operator**, **Backup Operator**, or **Viewer**).
8. Click **Save**.

The user is added to the list of all users.

You can later do the following:

- Edit any of the existing HYCU or identity provider users by clicking  **Edit** and making the required modifications. Keep in mind that the built-in user, AD users, and AD groups cannot be edited.
- Enable or disable specific users from signing in to HYCU. For details, see “Activating or deactivating a user” on page 358.
- Delete any of the existing users by clicking  **Delete**. Keep in mind that the built-in user cannot be deleted.

Adding a user to a group

Prerequisite

Only if you want to add a user to a self-service group. The self-service group must be created. For instructions, see “Creating a self-service group” on the next page.

Considerations

- You can add a user to multiple groups in which the user can have different user roles assigned. For details on user roles, see “User roles” on page 345.
- If an AD user has multiple user roles assigned based on membership in several AD groups, the user acquires the role with the highest privilege level. User roles are prioritized in the following order: Administrator > Backup and Restore Operator > Restore Operator > Backup Operator > Viewer. However, keep in mind that a role assigned to an AD user independently of an AD group always takes precedence over a role within an AD group.

Procedure

1. In the Self-Service panel, in the Detail view, select the group to which you want to add a user.
2. Click  **Add to Group**.

 **Note** You can add the user to the infrastructure group that is created by default or a self-service group that you must create yourself.

3. In the Username field, enter a user name.
-  **Important** *For AD user and AD group:* Enter a user name in one of the following formats: *user@domain* or *domain\name*.
4. From the User Role drop-down menu, select a role that you want to assign to the user (**Administrator**, **Backup and Restore Operator**, **Restore Operator**, **Backup Operator**, or **Viewer**).
5. Click **Add User**.

Depending on the needs of a specific data protection environment, you can at any time remove a user from a group by selecting the user that you want to remove and clicking  **Remove from Group**.

Creating a self-service group

Prerequisite

Only if you plan to use the Advanced option that allows you to automatically assign virtual machines to the self-service group that you are creating. Tags must be assigned to virtual machines on their source. For details, see “[Setting up automatic virtual machine assignment by using tags](#)” on page 355.

Procedure

1. In the Self-Service panel, click  **New**.
2. Enter a self-service group name and, optionally, its description.
3. *Only if you want virtual machines to be automatically assigned to the self-service group you are creating.* Click **Advanced**, enter a key and a value, and then click **Add**. If required, repeat this step for all the keys and the values that you want to add.
For details, see “[Setting up automatic virtual machine assignment by using tags](#)” on page 355.
4. Click **Save**.

You can later do the following:

- Add users to groups. For details, see “[Adding a user to a group](#)” on page 352.
- Edit any of the existing self-service groups by clicking  **Edit** and making the required modifications.
- Allow users belonging to a specific self-service group to see only policies whose names start with their group name followed by an underscore (for example, HYCUGroup_Policy1) and the Exclude policy (alongside of other policies already assigned to the virtual machines, file shares, and volume groups whose owners they are). To do so, in the HYCU config.properties file, set the policies.group.specific.synchronized configuration setting to true. Keep in mind that such policies can be edited or deleted only if they are not assigned to any entity. For details on how to customize the HYCU configuration settings, see “[Customizing HYCU configuration settings](#)” on page 523.
- Enable or disable specific self-service groups from signing in to HYCU. For details, see “[Activating or deactivating a self-service group](#)” on page 359.
- Delete any of the existing self-service groups by clicking  **Delete**.

Setting ownership

By setting ownership of virtual machines, file shares, and volume groups, you enable specific groups to protect only the assigned virtual machines, file shares, and volume groups. Depending on the entity to which you want to assign an owner, see one of the following sections:

- “[Setting ownership of virtual machines](#)” below
- “[Setting ownership of file shares](#)” on page 357
- “[Setting ownership of volume groups](#)” on page 357

Setting ownership of virtual machines

When setting ownership of virtual machines, you can use one of the following approaches:

Approach to setting ownership	Instructions
Assign tags to virtual machines on their source, and then specify the corresponding keys and values in HYCU.	“ Setting up automatic virtual machine assignment by using tags ” on the next page

Approach to setting ownership	Instructions
Select virtual machines in HYCU, and then assign these virtual machines to a specific self-service group.	“ Setting up virtual machine assignment by using the Assign option ” on the next page

Considerations

- Assigning virtual machines to self-service groups by using tags takes precedence over assigning policies by using the Assign option. This means that the tag added to the virtual machine defines to which self-service group the virtual machine will be assigned (even if this virtual machine is already assigned to another self-service group by using the Assign option).
- A virtual machine is not unassigned from its self-service group if a tag is removed from the virtual machine.

Setting up automatic virtual machine assignment by using tags

By setting up automatic virtual machine assignment, you ensure that virtual machines are automatically assigned to self-service groups.

After you assign tags to virtual machines and specify the matching keys and values, and the comparison of these values shows that the specified values match, the corresponding virtual machines are automatically assigned to the groups during the next virtual machine synchronization.

HYCU performs the automatic synchronization of virtual machines every five minutes. However, you can at any time update the list of virtual machines also manually by clicking  **Refresh** in the Virtual Machines panel.

 **Note** HYCU uses the term tags to refer also to categories and custom attributes that are assigned to virtual machines on Nutanix clusters and in vSphere environments.

Procedure

1. Sign in to the Nutanix Prism web console, the vSphere (Web) Client, the AWS GovCloud (US) console, or the Azure Government portal.
2. Assign tags to virtual machines for which you want to set up automatic assignment. For instructions, see the relevant documentation.
3. Sign in to the HYCU web user interface.
4. Specify the matching keys and values for the specific self-service group as described in “[Creating a self-service group](#)” on page 353.

ⓘ Important Depending on your data protection environment, the key and the value that you should enter represent the following:

- *For Nutanix AHV clusters:* The name and the value of the category.
- *For Nutanix ESXi clusters or vSphere environments:* The tag name and the category of the tag, or the attribute and the value of the custom attribute.
- *For AWS GovCloud (US) or Azure Government environments:* The name and the value of the tag.

Setting up virtual machine assignment by using the Assign option

Consideration

When changing ownership of virtual machines, you can choose whether you want data protected by a specific owner to be kept or deleted. If you choose to keep data protected by the specific owner, such virtual machines will be kept in HYCU with the Protected deleted status. Restoring these virtual machines by using the Restore VM option is possible only if they are deleted from the source before the restore is performed.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, select the virtual machines to which you want to assign an owner, and then click  **Owner**.
2. From the list of groups, select which group you want to assign as the owner of the selected virtual machines, and then click **Assign**.

ⓘ Important If a virtual machine or an application has backup or restore jobs in progress, or a scheduled backup task in the queue, you cannot assign a new group to the relevant virtual machine.

Depending on the needs of a specific data protection environment, you can at any time remove the owner from the virtual machines by selecting the virtual machines from which you want to remove the owner, and then clicking  **Owner** followed by **Unassign**.

Setting ownership of file shares

Consideration

When changing ownership of file shares, you can choose whether you want data protected by specific owners to be kept or deleted. If you choose to keep data protected by the specific owner, such file shares will be kept in HYCU with the Protected deleted status.

Accessing the Shares panel

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

Procedure

1. In the Shares panel, select file shares to which you want to assign an owner, and then click  **Owner**.
2. From the list of groups, select which group you want to assign as an owner of the selected file shares, and then click **Assign**.

 **Important** If any backup or restore job for a file share is already in progress, or a scheduled backup task is in the queue, you cannot assign a new group to this file share.

Depending on the needs of a specific data protection environment, you can at any time remove an owner from the file shares by selecting the file shares from which you want to remove the owner, and then clicking  **Owner** followed by **Unassign**.

Setting ownership of volume groups

Consideration

When changing ownership of volume groups, you can choose whether you want data protected by specific owners to be kept or deleted. If you choose to keep data protected by the specific owner, such volume groups will be kept in HYCU with the Protected deleted status.

Accessing the Volume Groups panel

To access the Volume Groups panel, in the navigation pane, click  **Volume Groups**.

Procedure

1. In the Volume Groups panel, select volume groups to which you want to assign an owner, and then click  **Owner**.
2. From the list of groups, select which group you want to assign as an owner of the selected volume groups, and then click **Assign**.

! Important If any backup or restore job for a volume group is already in progress, or a scheduled backup task is in the queue, you cannot assign a new group to this volume group.

Depending on the needs of a specific data protection environment, you can at any time remove an owner from the volume groups by selecting the volume groups from which you want to remove the owner, and then clicking  **Owner** followed by **Unassign**.

Activating or deactivating users or self-service groups

Depending on the nature of your business, you can at any time enable or disable specific users or self-service groups from signing in to HYCU by activating or deactivating them. By activating or deactivating a self-service group, you enable or disable all users belonging to the specific self-service group from signing in to HYCU as members of that group.

Activating or deactivating a user

Procedure

1. In the Self-Service panel, click  **Manage Users**.
2. From the list of all users, select the one whose status you want to change.
3. Depending on the status of the user, do one of the following:
 - If the status of the selected user is Inactive and you want to activate it, click  **Activate**.
 - If the status of the selected user is Active and you want to deactivate it, click  **Deactivate**.

Activating or deactivating a self-service group

Procedure

1. In the Self-Service panel, from the list of self-service groups, select the one whose status you want to change.
2. Depending on the status of the self-service group, do one of the following:
 - If the status of the selected self-service group is Inactive and you want to activate it, click **Activate**.
 - If the status of the selected self-service group is Active and you want to deactivate it, click **Deactivate**.

Note If a user is a member of several self-service groups and at least one of these groups has the Active status, the user is automatically switched to it. If there is more than one group with the Active status to which the user belongs, the user is automatically switched to the one that was created first.

Switching to another group

As a user you can belong to one or more groups and sign in to HYCU with all the permissions associated with the group to which you belong. If you are a member of more than one group, you can at any time switch to another group (provided that its status is Active) while being signed in to HYCU. This means that you can select any of the groups to which you belong and use it for a session.

Procedure

1. Click the group under which you are currently signed in to HYCU at the upper right of the screen.

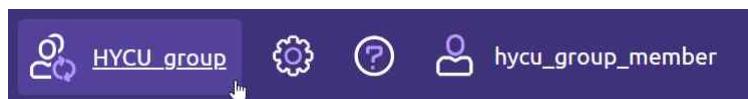


Figure 10-2: Example of a self-service group, HYCU_group, under which a user, hycu_group_member, is signed in to HYCU

2. From the list of all groups to which you belong, select the one to which you want to switch.

The group under which you are currently signed in to HYCU has next to it.

3. Click **Switch**.

You are automatically switched to the group you selected.

Updating your user profile

You can edit your name, email address, preferred language, and authentication settings by using the Update Profile option.

Consideration

As a user with the Administrator role assigned, you can edit other users' information through the Self-Service panel. For details, see “[Creating a user](#)” on [page 348](#).

Accessing the Update Profile dialog box

To access your Update Profile dialog box, click  at the upper right of the screen and then select **Update Profile**.

Procedure

1. In the Name field, specify a new name.
2. In the Email field, enter the email address that you want to be associated with your user profile.
3. From the Language drop-down menu, select the preferred language.
4. *Optional.* Enable two-factor authentication by using the **Two-Factor Authentication** switch. Select the two-factor authentication method:
 - **Time-based one-time password**
This option enables the use of a time-based one-time password (OTP) generated by an OTP application.
 - **FIDO**
This option enables the use of an authenticator complying with FIDO protocols (FIDO authenticator).
5. Click **Save**.
6. *Only if you enabled two-factor authentication.* Perform the initial two-factor authentication setup:
 - *For a time-based one-time password:* The Configure Two-Factor Authentication dialog box opens. Do the following:
 - a. Scan the QR code with an appropriate OTP application (for example Google Authenticator on a mobile phone) or alternatively enter the

OTP backup code in the application manually.

- b. In the Authentication code field, enter the generated six-digit code, and then click **Confirm** to finish the setup process.

 **Note** If you do not set up a one-time password, the Configure Two-Factor Authentication dialog box opens during your next sign-in.

- *For FIDO:* The FIDO Authenticators dialog box opens. Do the following:
 - a. Follow the wizard to register the authenticator (for example, a security key or Windows Hello with a fingerprint reader). The process depends on the type of authenticator you select and the operating system version.
 - b. Enter a name for the authenticator, and then click **Register**.

 **Note** If you do not complete the registration of at least one authenticator, you are prompted to register one at the first sign-in after two-factor authentication is enabled.

Later you can also add additional authenticators or revoke existing ones. For more details, see “[Managing FIDO authenticators](#)” on [page 450](#).

Chapter 11

Administering

After you deploy HYCU, you can perform various administration tasks through the ☰ **Administration** menu to customize HYCU for your data protection environment.

I want to...	Procedure
Add cloud accounts to HYCU.	“Adding a cloud account” on the next page
Configure encryption for targets.	“Configuring target encryption” on page 372
Integrate HYCU with identity providers.	“Integrating HYCU with identity providers” on page 373
Manage HYCU instances.	“Managing HYCU instances” on page 381
Set the iSCSI Initiator secret.	“Setting the iSCSI Initiator secret” on page 384
Obtain a permanent HYCU license.	“Licensing” on page 385
Configure log file settings to troubleshoot problems if HYCU does not perform as expected.	“Setting up logging” on page 389
Change network settings or enable network bandwidth throttling.	“Configuring your network” on page 392
Set data retention for restore point tiers.	“Managing data retention” on page 396
Set power options.	“Setting power options” on page 403
Securely store, access, and manage my credentials (secrets) by employing	“Managing secrets” on page 404

I want to...	Procedure
the Conjur secrets management solution.	
Configure an SMTP server.	“Configuring an SMTP server” on page 407
Upgrade HYCU to a new available version.	“Upgrading HYCU” on page 408
Apply a HYCU update.	“Applying HYCU updates” on page 427
Configure the SSL certificate.	“Configuring SSL certificates” on page 432
Share telemetry diagnostic data with HYCU.	“Sharing telemetry data with HYCU” on page 438

If for whatever reason you decide that you no longer want to use HYCU for protecting your data, you can easily remove it from your system. For details, see “[Removing HYCU](#)” on page 439.

Adding a cloud account

You must add one or more cloud accounts to HYCU before you can migrate virtual machines across different infrastructures, perform disaster recovery of data to cloud, monitor cloud data protection environments, deploy HYCU to cloud, add an AWS GovCloud (US) region to HYCU, add an Azure Government subscription to HYCU, or store data to a Google Cloud target.

The type of cloud account that you must add to HYCU depends on which data protection tasks you want to perform and which HYCU backup and recovery service you use to protect your data. For details on the different types of cloud accounts, see the following sections:

- “[Adding a cloud provider account](#)” below
- “[Adding a HYCU account](#)” on page 371

Adding a cloud provider account

You must add one or more cloud provider accounts to HYCU before performing any of the following data protection tasks:

- Adding an AWS GovCloud (US) region to HYCU.
- Adding an Azure Government subscription to HYCU.
- Storing data to a Google Cloud target.
- Migrating data protected with HYCU from the on-premises environment to cloud.
- Migrating data protected with HYCU R-Cloud, HYCU for AWS, HYCU for Google Cloud, or HYCU for Azure from cloud to the on-premises environment.
- Performing disaster recovery of data to cloud in the event of a disaster.
- Monitoring the HYCU for Google Cloud and HYCU for Azure data protection environments in HYCU Manager.
- Deploying HYCU to Google Cloud.

Consideration

Migrating virtual machines across different infrastructures, performing disaster recovery of data to cloud, monitoring cloud data protection environments, and deploying HYCU to cloud are supported only if you own a HYCU R-Cloud license.

Depending on which data protection tasks you want to perform, add one or more cloud provider accounts to HYCU:

I want to...	Cloud provider account	Instructions
<ul style="list-style-type: none"> • Migrate data protected with HYCU to AWS. • Migrate data protected with HYCU R-Cloud or HYCU for AWS to the on-premises environment. • Perform disaster recovery of data to AWS. 	AWS user account	<p>“Adding an AWS user account” on the next page</p>
<ul style="list-style-type: none"> • Add an AWS GovCloud (US) region to HYCU to be able to protect AWS GovCloud (US) virtual machines and applications. 	AWS GovCloud (US) account	<p>“Adding an AWS GovCloud (US) account” on page 366</p>
<ul style="list-style-type: none"> • Store data to a Google Cloud target. 	Google Cloud service account	<p>“Adding a Google Cloud service</p>

I want to...	Cloud provider account	Instructions
<ul style="list-style-type: none"> • Migrate data protected with HYCU to Google Cloud. • Migrate data protected with HYCU R-Cloud or HYCU for Google Cloud to the on-premises environment. • Perform disaster recovery of data to Google Cloud. • Monitor my HYCU for Google Cloud data protection environment in HYCU Manager. • Deploy HYCU to Google Cloud. 		<p>“Adding a Google Cloud account” on page 367</p>
<ul style="list-style-type: none"> • Migrate data protected with HYCU to Azure. • Migrate data protected with HYCU for Azure to the on-premises environment. • Perform disaster recovery of data to Azure. • Monitor my HYCU for Azure data protection environment in HYCU Manager. 	Azure service principal	<p>“Adding an Azure service principal” on page 369</p>
<ul style="list-style-type: none"> • Add an Azure Government subscription to HYCU. • Migrate data protected with HYCU to Azure Government. • Perform disaster recovery of data to Azure Government. 	Azure Government service principal	<p>“Adding an Azure Government service principal” on page 371</p>

Adding an AWS user account

Prerequisite

A user account must be created in AWS and it must have permissions to perform the following actions in the S3 service: `ListBucket`, `CreateBucket`, `DeleteBucket`, `GetObject`, `PutObject`, `DeleteObject`, and

`PutBucketPublicAccessBlock`. In addition, you must set the `Resources` value to **All resources** for these actions. For more information about S3 permissions, see AWS documentation.

Accessing the Cloud Accounts dialog box

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

Procedure

1. In the Cloud Accounts dialog box, click the **Cloud Provider Accounts** tab, and then click  **New**. The Select Cloud dialog box opens.
2. Select **Add AWS user account**, and then click **Next**. The AWS Authentication dialog box opens.
3. In the Name field, enter a name for your AWS user account.
4. In the Access key ID field, enter the access key ID of your AWS user account.
5. In the Secret access key field, enter the secret access key of your AWS user account.

 **Note** The access key ID and the secret access key are used to authenticate AWS API service calls.

6. Click **Save**.

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

Adding an AWS GovCloud (US) account

Prerequisite

Your AWS GovCloud (US) account must have permissions to perform the following actions in the S3 service: `ListBucket`, `CreateBucket`, `DeleteBucket`, `GetObject`, `PutObject`, `DeleteObject`, and `PutBucketPublicAccessBlock`. In addition, you must set the `Resources` value to **All resources** for these actions.

For more information about S3 permissions, see AWS documentation.

Accessing the Cloud Accounts dialog box

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

Procedure

1. In the Cloud Accounts dialog box, click the **Cloud Provider Accounts** tab, and then click  **New**.
2. Select **Add AWS GovCloud (US) account**, and then click **Next**.
3. In the Name field, enter a name for your AWS GovCloud (US) account.
4. In the Access key ID field, enter the access key ID of your AWS GovCloud (US) account.
5. In the Secret access key field, enter the secret access key of your AWS GovCloud (US) account.

 **Note** The access key ID and the secret access key are used to authenticate AWS GovCloud (US) API service calls.

6. Click **Save**.

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

Adding a Google Cloud service account

The type of Google Cloud service account that you add to HYCU depends on what data protection tasks you want to perform.

 **Important** You must always add a dedicated service account that you have created yourself to HYCU.

I want to...	Service account to add
Add a Google Cloud project as a source to HYCU.	An account that is granted the Compute Admin role or a custom role with the permissions described in “ Adding a Google Cloud project ” on page 60.
Deploy a HYCU instance to Google Cloud.	An account with permissions to access NetApp Cloud Volumes Service for Google Cloud. For details, see “ Enabling HYCU to access NetApp Cloud Volumes Service for Google Cloud ” on page 75.
Protect file shares created by NetApp Cloud Volumes Service for Google Cloud.	An account that has access to the buckets where you want to store your backup data.
Store data to a Google Cloud target.	An account that has access to the buckets where you want to store your backup data.

I want to...	Service account to add
Migrate data protected with HYCU R-Cloud or HYCU for Google Cloud from Google Cloud to the on-premises environment.	An account that is added to HYCU R-Cloud or HYCU for Google Cloud, and has the Storage Admin role assigned on the projects containing the instances.
Migrate data protected with HYCU from the on-premises environment to Google Cloud.	An account that is added to HYCU R-Cloud or HYCU for Google Cloud, and has the Compute Admin, Storage Admin, Service Account User, and Service Account Token Creator roles assigned on the projects where you want to migrate your virtual machines.
Perform disaster recovery of data to Google Cloud in the event of a disaster.	An account that is added to HYCU R-Cloud or HYCU for Google Cloud, and has the Compute Admin, Storage Admin, Service Account User, and Service Account Token Creator roles assigned on the projects where you want to perform disaster recovery.
Monitor my HYCU for Google Cloud data protection environment in HYCU Manager.	An account with permissions to access the protection sets that you want to monitor in HYCU Manager.

Prerequisites

- The service account must be configured in Google Cloud.
- The following APIs must be enabled on the Google Cloud project on which the service account was created:
 - Cloud Resource Manager API
 - Compute Engine API
 - Cloud Storage API
 - Identity and Access Management API

For instructions on how to enable the listed APIs, see Google Cloud documentation.

- You must have access to a valid JSON file that stores the service account information, including its private key.

Accessing the Cloud Accounts dialog box

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

Procedure

1. In the Cloud Accounts dialog box, click the **Cloud Provider Accounts** tab, and then click  **New**. The Select Cloud dialog box opens.
2. Select **Add Google Cloud service account**, and then click **Next**. The Google Cloud Authentication dialog box opens.
3. Browse for the JSON file with the service account information. In the Service account authentication field, the file name is displayed.

 **Note** Only if you are signed in to HYCU as a self-service group administrator. If you use Conjur for managing your HYCU secrets, you can enable the **Retrieve values from secrets manager** switch if you want to provide the secret instead of browsing for the file. For details on managing secrets, see “[Managing secrets](#)” on page 404.

4. In the Name field, enter a name for your service account.
5. Click **Upload or Save**.

You can later edit any of the existing cloud accounts (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**). Keep in mind that you cannot delete a cloud account in the following cases:

- A Google Cloud target uses this account.
- A protection set that is monitored in HYCU Manager uses this account.

Adding an Azure service principal

Prerequisite

The service principal must be created in Azure and added to HYCU for Azure. For details, see HYCU for Azure documentation.

 **Important** You must always add a dedicated service principal that you have created yourself to HYCU and not use the default one that HYCU for Azure automatically creates for you when you start using the service.

The role that must be assigned to the service principal depends on what data protection tasks you want to perform:

I want to...	Required roles
Migrate data protected with HYCU for Azure from Azure to the on-premises environment.	
Migrate data protected with HYCU from the on-premises environment to Azure.	<ul style="list-style-type: none"> Contributor role assigned at the subscription level Storage Blob Data Contributor role assigned at the subscription, resource group, or storage account level
Perform disaster recovery of data to Azure in the event of a disaster.	
Monitor my HYCU for Azure data protection environment in HYCU Manager.	<ul style="list-style-type: none"> Contributor role assigned at the subscription level

Accessing the Cloud Accounts dialog box

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

Procedure

1. In the Cloud Accounts dialog box, click the **Cloud Provider Accounts** tab, and then click  **New**. The Select Cloud dialog box opens.
2. Select **Add Azure service principal**, and then click **Next**. The Azure Authentication dialog box opens.
3. In the Name field, enter the name for your service principal.
4. In the Tenant ID field, enter your tenant ID.
5. In the Application ID field, enter the HYCU for Azure application registration ID from Microsoft Entra ID.
6. In the Client secret value field, enter the client secret's value.
7. Click **Save**.

You can later edit any of the existing service principals (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**). Keep in mind that you cannot delete a service principal if a protection set that is monitored in HYCU Manager uses this account.

Adding an Azure Government service principal

Prerequisites

- The service principal must be created in Azure Government.
- The service principal must be assigned the Contributor role at the subscription level.

Accessing the Cloud Accounts dialog box

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

Procedure

1. In the Cloud Accounts dialog box, click the **Cloud Provider Accounts** tab, and then click  **New**. The Select Cloud dialog box opens.
2. Select **Add Azure Government service principal**, and then click **Next**. The Azure Government Authentication dialog box opens.
3. In the Name field, enter the name for your service principal.
4. In the Tenant ID field, enter your tenant ID.
5. In the Application ID field, enter the HYCU application registration ID from Microsoft Entra ID.
6. In the Client secret value field, enter the client secret's value.
7. Click **Save**.

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

Adding a HYCU account

If you use HYCU R-Cloud or HYCU for AWS to protect your data, you must add one or more HYCU accounts to HYCU before performing any of the following data protection tasks:

- Migrating data protected with HYCU from the on-premises environment to cloud.
- Migrating data protected with HYCU R-Cloud or HYCU for AWS from cloud to the on-premises environment.
- Performing disaster recovery of data to cloud in the event of a disaster.

Considerations

- *Only if you added an AWS user account with a previous version of HYCU.* After you upgrade HYCU to version 5.0.0, your AWS user account is automatically split into an AWS user account and a HYCU account. Therefore, you do not need to add a HYCU account.
- If you use HYCU for Google Cloud to protect your data, adding a HYCU account is not required.

Accessing the Cloud Accounts dialog box

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

Procedure

1. In the Cloud Accounts dialog box, click the **HYCU Accounts** tab, and then click  **New**.
2. In the Name field, enter a name for your HYCU account.
3. In the HYCU account ID field, enter the HYCU account ID that you received when you subscribed to HYCU R-Cloud or HYCU for AWS.
4. In the Username and Password fields, enter the credentials of the user account that you use for accessing HYCU R-Cloud or HYCU for AWS.
5. Click **Save**.

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

Configuring target encryption

If you enabled target encryption when setting up a target, you can view the information on which algorithm is used, view a list of encrypted targets, export the encryption key to a file, and import the encryption key.

Accessing the Encryption dialog box

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

Exporting an encryption key

Procedure

1. In the Encryption dialog box, click **Export**.
2. Save the exported file to a safe location.

Importing an encryption key

Procedure

1. In the Encryption dialog box, click **Import**.
2. In the Import dialog box, browse for the file containing the encryption key, and then click **Import**.

Integrating HYCU with identity providers

You can integrate HYCU with the Active Directory identity provider and identity providers that support the OpenID Connect authentication protocol, such as Google, Keycloak, Microsoft, Okta, and Active Directory Federation Services to give users the possibility to securely sign in to HYCU by using these identity providers, without the need to maintain dedicated credentials for HYCU.

When integrating HYCU with an identity provider, you must complete the following tasks:

Task	Instructions
1. Add an identity provider to HYCU to be able to authenticate users.	Follow the procedure described in “Adding an identity provider to HYCU” on the next page .
2. Create a user for whom you want to enable signing in by using the identity provider, and then add this user to a user group.	Follow the procedures described in “Creating a user” on page 348 and “Adding a user to a group” on page 352 .

Adding an identity provider to HYCU

Prerequisites

- *Only when adding identity providers that support the OpenID Connect authentication protocol.* HYCU must be registered as a web application within the identity provider that you plan to add to HYCU. When registering HYCU, depending on the identity provider, make sure the following is done:
 - *Active Directory Federation Services:*
 - In Active Directory Federation Services, you must select **Server application accessing a web API** as the client-server application, and **openid** and **allatclaims** when configuring application permissions. Also, make sure that the outgoing claim type to which you map the ObjectGUID attribute is ObjectGUID when configuring rules.
 - *Only if you are adding identity provider groups.* In Active Directory Federation Services, create an issuance transform rule, and then set the groups claim.
 - *Keycloak:*
 - *Only if you are adding identity provider groups.* In Keycloak, when configuring the client scopes, add a mapper of the type Group Membership, with the token claim named `group_mem`. Make sure to disable the **Full group path** switch.
 - *Microsoft:*
 - In Azure, HYCU must be given access permissions to the following Azure API: Microsoft Graph with delegated permissions for `User.Read`.
 - *Only if you are adding identity provider groups.* In Azure, under Token configuration, add a groups claim.
 - *Okta:*
 - In Okta, you must select **Authorization Code** under Client acting on behalf of a user as the grant type.
 - *Only if you are adding identity provider groups.* In Okta, add a new scope named `groups`, and then add a new claim.

For instructions on how to register an application, see the respective identity provider documentation.

- Only if you plan to use LDAPS with Active Directory: LDAPS authentication must be set up. For details, see “[Setting up LDAPS authentication](#)” on [page 448](#).

Considerations

- To increase the security of user accounts further, you can also configure multi-factor authentication within the identity providers. For instructions on how to do this, see the respective identity provider documentation.
- If you use Active Directory as your authentication source in HYCU, you can also enable certificate authentication to allow users to sign in to the HYCU web user interface with a client certificate or a smart card. For instructions, see “[Enabling certificate authentication](#)” on [page 380](#).

Accessing the Identity Providers dialog box

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

Procedure

1. In the Identity Providers dialog box, click  **New**.
2. In the Name field, enter a name for the identity provider.
3. From the Type drop-down menu, select one of the following types of identity providers, and then follow the instructions:

Identity provider type	Instructions
Active Directory	<p>a. In the Domain field, enter the FQDN or domain alias name of the Active Directory. If you plan to use AD groups, it is mandatory to enter the FQDN.</p> <p>For example, if you enter <code>mycompany.com</code> as the FQDN and <code>mc</code> as the alias domain name, the user will be able to sign in to HYCU with <code><Username>@mycompany.com</code> or <code>mc\<Username></code>.</p> <p> Note You can enter more than one FQDN or domain alias name. In this case, press the Spacebar after entering each one.</p>

Identity provider type	Instructions
	<p>b. In the Provider URL field, enter the URL of the corresponding LDAP server in one of the following formats:</p> <ul style="list-style-type: none"> • <code>ldap://<LDAPServerHostnameorIPAddress>:<Port></code> When using the LDAP protocol, the default port is 389. Entering the port is optional if the default value is used. • <i>Only if LDAPS authentication is set up.</i> <code>ldaps://<LDAPServerHostname>:<Port></code> <p>ⓘ Important Make sure that the LDAP server host name matches the DNS entry specified in the Subject Alternative Name (SAN) extension of the LDAP server's certificate. Otherwise, connection to the LDAP server will fail.</p> <p>When using the LDAPS protocol, the default port is 636. Entering the port is optional if the default value is used.</p> <p>ⓘ Note You can enter more than one URL. In this case, press the Spacebar after entering each one.</p> <p>c. <i>Only if you plan to enable certificate authentication.</i> Enable the Use service account switch, and then enter the user name and password of the service account that HYCU will use to sign in to the Active Directory and authorize users.</p>
Active Directory Federation Services	<p>a. In the Client ID field, enter the application ID that is generated by the identity provider.</p> <p>b. In the Client secret field, enter the application secret that is associated with the client ID and generated by the identity provider.</p> <p>c. In the Redirect URI field, enter the URL to which the</p>

Identity provider type	Instructions
	<p>user will be redirected after authentication. The format is as follows:</p> <div data-bbox="568 489 975 523" style="background-color: #f0f0f0; padding: 5px; border-radius: 5px; text-align: center;"><code>https://<ServerName>:8443</code></div> <p>In this instance, <code><ServerName></code> is the fully qualified domain name of the HYCU server.</p> <p>For example:</p> <div data-bbox="568 720 1038 754" style="background-color: #f0f0f0; padding: 5px; border-radius: 5px; text-align: center;"><code>https://hycu.example.com:8443</code></div> <p>d. In the Issuer field, enter the URL of the issuer of the identity provider.</p>
Google	<p>a. In the Client ID field, enter the application ID that is generated by the identity provider.</p> <p>b. In the Client secret field, enter the application secret that is associated with the client ID and generated by the identity provider.</p> <p>c. In the Redirect URI field, enter the URL to which the user will be redirected after authentication. The format is as follows:</p> <div data-bbox="568 1304 975 1338" style="background-color: #f0f0f0; padding: 5px; border-radius: 5px; text-align: center;"><code>https://<ServerName>:8443</code></div>
Keycloak	<p>a. In the Client ID field, enter the application ID that is generated by the identity provider.</p> <p>b. In the Client secret field, enter the application secret that is associated with the client ID and generated by the identity provider.</p>

Identity provider type	Instructions
	<p>c. In the Redirect URI field, enter the URL to which the user will be redirected after authentication. The format is as follows:</p> <div data-bbox="568 534 975 570" style="background-color: #f0f0f0; padding: 5px; text-align: center;"><code>https://<ServerName>:8443</code></div> <p>In this instance, <code><ServerName></code> is the fully qualified domain name of the HYCU server.</p> <p>For example:</p> <div data-bbox="568 770 1038 806" style="background-color: #f0f0f0; padding: 5px; text-align: center;"><code>https://hycu.example.com:8443</code></div> <p>d. In the Issuer field, enter the URL of the issuer of the identity provider.</p>
Microsoft	<p>a. In the Client ID field, enter the application ID that is generated by the identity provider.</p> <p>b. In the Client secret field, enter the application secret that is associated with the client ID and generated by the identity provider.</p> <p>c. In the Redirect URI field, enter the URL to which the user will be redirected after authentication. The format is as follows:</p> <div data-bbox="568 1343 975 1379" style="background-color: #f0f0f0; padding: 5px; text-align: center;"><code>https://<ServerName>:8443</code></div> <p>In this instance, <code><ServerName></code> is the fully qualified domain name of the HYCU server.</p> <p>For example:</p> <div data-bbox="568 1578 1038 1614" style="background-color: #f0f0f0; padding: 5px; text-align: center;"><code>https://hycu.example.com:8443</code></div>
Okta	<p>a. In the Client ID field, enter the application ID that is generated by the identity provider.</p> <p>b. In the Client secret field, enter the application secret that is associated with the client ID and generated by the identity provider.</p>

Identity provider type	Instructions
	<p>c. <i>Optional.</i> In the Authorization server field, enter the ID of the authorization server. If not provided, the default authorization server is used.</p> <p>d. In the Redirect URI field, enter the URL to which the user will be redirected after authentication. The format is as follows:</p> <div data-bbox="568 676 975 709" style="background-color: #f0f0f0; padding: 5px; text-align: center;"><code>https://<ServerName>:8443</code></div> <p>In this instance, <code><ServerName></code> is the fully qualified domain name of the HYCU server.</p> <p>For example:</p> <div data-bbox="568 916 1038 950" style="background-color: #f0f0f0; padding: 5px; text-align: center;"><code>https://hycu.example.com:8443</code></div> <p>e. In the Issuer field, enter the URL of the issuer of the identity provider.</p>
OpenID Connect IdP	<p>a. In the Client ID field, enter the application ID that is generated by the identity provider.</p> <p>b. In the Client secret field, enter the application secret that is associated with the client ID and generated by the identity provider.</p> <p>c. In the Redirect URI field, enter the URL to which the user will be redirected after authentication. The format is as follows:</p> <div data-bbox="568 1500 975 1534" style="background-color: #f0f0f0; padding: 5px; text-align: center;"><code>https://<ServerName>:8443</code></div> <p>In this instance, <code><ServerName></code> is the fully qualified domain name of the HYCU server.</p> <p>For example:</p> <div data-bbox="568 1740 1038 1774" style="background-color: #f0f0f0; padding: 5px; text-align: center;"><code>https://hycu.example.com:8443</code></div> <p>d. In the Issuer field, enter the URL of the issuer of the identity provider.</p>

Identity provider type	Instructions
	<p>e. In the Authorization endpoint field, enter the authorization endpoint of the identity provider.</p> <p>f. In the Token endpoint field, enter the token endpoint of the identity provider.</p> <p>g. In the JWKS endpoint field, enter the JSON web key set endpoint of the identity provider.</p> <p>h. In the UserInfo endpoint field, enter the UserInfo endpoint of the identity provider.</p>

4. Click **Save**.

You can later do the following:

- Edit information about any of the existing identity providers by clicking  **Edit** and making the required modifications.

 **Note** The Redirect URI field shows to which URL the user will be redirected after authentication (for example, <https://hycu.example.com:8443>). The prepopulated host name is the host name of the HYCU backup controller to which you are authenticating user access.

- Delete any of the existing identity providers by clicking  **Delete**.

Enabling certificate authentication

By enabling certificate authentication, you allow Active Directory users to sign in to the HYCU web interface by using a client certificate or a smart card, without having to enter a password.

Prerequisites

- At least one Active Directory with a configured service account must be added to HYCU.
- A CA-signed certificate must be imported to HYCU. For details on how to do this, see “[Importing a custom certificate](#)” on page 435.

Limitation

If certificate authentication is enabled, setting up two-factor authentication for AD users is not supported.

Consideration

When you enable or disable certificate authentication, all affected users that are signed in the HYCU web user interface will lose their connections and will be required to sign in again.

Procedure

1. In the Identity Providers dialog box, use the **Enable certificate authentication** switch if you want to enable certificate authentication.
2. From the CA certificate drop-down menu, select the CA-signed certificate for verifying the client certificate.

Tip By clicking **Manage**, you are automatically directed to the SSL Certificates dialog box where you can manage your certificates.

Managing HYCU instances

All existing HYCU instances in your data protection environment are listed in the HYCU Instances dialog box. Besides viewing all the existing HYCU instances, you can use this dialog box also to create new HYCU instances, view information about each HYCU instance, delete HYCU instances, and apply updates to HYCU instances.

For details on HYCU instances, see “[HYCU instances](#)” on page 63. For instructions on how to apply an update to HYCU instance, see “[Applying an update to a HYCU instance](#)” on page 430.

Accessing the HYCU Instances dialog box

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

Creating a HYCU instance by using the HYCU web user interface

You can create a HYCU instance by using the HYCU web user interface as an alternative to creating it by deploying the HYCU virtual appliance in the HYCU Instance mode.

Prerequisites

- *For creating a HYCU instance on a Nutanix AHV cluster:* The HYCU virtual appliance image must be present on the Nutanix cluster in the following format:
`hycu-<Version>-<Revision>`
For example, `hycu-5.0.0-3634`.
- *For creating a HYCU instance on a Nutanix ESXi cluster or in a vSphere environment:*
 - A user with specific privileges for vCenter Servers must be specified. For details on which privileges must be assigned to a vSphere user, see [“Assigning privileges to a vSphere user” on page 456](#).
 - The HYCU OVF package must be imported to the vCenter Server content library and its format must be as follows:
`hycu-<Version>-<Revision>`
For example, `hycu-5.0.0-3634`.
- The Nutanix cluster, the vCenter Server, or the Google Cloud project where you want to create your HYCU instance must be added as a source to HYCU. For instructions, see [“Adding sources” on page 53](#).

Limitation

For creating a HYCU instance in Google Cloud: Using the web user interface is the only supported method to add a HYCU instance.

Procedure

1. In the HYCU Instances dialog box, click  **New**.
2. In the General section, enter a name for the HYCU instance.
3. In the Network configuration section, do the following:

- a. Enter a host name for the HYCU instance.

! Important Make sure that you enter a unique host name for each HYCU instance that you create and follow these rules:

- The host name contains only letters, numbers, hyphens (-), and periods. The maximum number of characters is 253 and at least one of the characters is a letter.
- The maximum number of characters in each host name segment is 63. A host name segment cannot begin or end with a hyphen.
- The top-level domain cannot begin or end with a number.

- b. Use the **DHCP** switch if you want a dynamic IP address to be assigned to the HYCU instance.

Otherwise, specify the IP address, the netmask, and the gateway. You can click **Auto** for the netmask and the gateway if you want HYCU to automatically set them for you.

4. In the Deployment section, do the following:

- a. From the Destination drop-down menu, select the Nutanix cluster, the vCenter Server, or the Google Cloud project where you want to create your HYCU instance.
- b. *Only if you are creating the HYCU instance on a Nutanix cluster or in a vSphere environment.* From the Storage container drop-down menu, select the storage container where you want to create your HYCU instance.

! Tip If you select **Select automatically**, HYCU will select the storage container with the most available space.

- c. *Only if you are creating the HYCU instance on a Nutanix cluster or in a vSphere environment.* From the Network drop-down menu, select a VLAN.

5. Click **Save**.

Viewing HYCU instance information

You can view the following information about each HYCU instance:

HYCU instance information	Description
VM name	Name of the HYCU instance, if known.
Hostname	Host name of the HYCU instance.

HYCU instance information	Description
Source	Nutanix cluster, the vCenter Server, or the Google Cloud project in which the HYCU instance resides.
Status	Shows if the HYCU instance is up and running, and communicating with the HYCU backup controller.
Version	Version of the HYCU instance (for example, 5.0.0-3634).
IP address	IP address currently assigned to the HYCU instance.

Deleting a HYCU instance

Procedure

1. In the HYCU Instances dialog box, from the list of HYCU instances, select the one that you want to delete, and then click  **Delete**.

! Important The selected HYCU instance will be deleted from both HYCU and the source.
2. In the Delete HYCU Instance dialog box, click **Delete** to confirm that you want to delete the selected HYCU instance.

Setting the iSCSI Initiator secret

During the HYCU deployment, the HYCU iSCSI client, referred to as the iSCSI Initiator, is set up so that HYCU can use iSCSI targets for storing data.

If you want to configure mutual CHAP authentication between the iSCSI Initiator and the iSCSI target, you must specify the iSCSI Initiator secret (the security key). For details on how to enable mutual authentication, see “[Setting up targets](#)” on page 78.

Accessing the iSCSI Initiator dialog box

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

Procedure

1. In the iSCSI Initiator dialog box, in the Secret field, enter the iSCSI Initiator secret.
2. Click **Save**.

Licensing

After you deploy the HYCU virtual appliance, you can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period.

The HYCU license is linked to the HYCU backup controller and you can decide on the license type or a combination of license types that best suits your environment. The following license types are available:

- Standard licenses

- Socket-based licenses

Licenses are based on the number of CPU sockets on the sources (Nutanix clusters, vCenter Servers, Nutanix Files, and servers) that you plan to protect by using HYCU. Socket-based licenses cannot be used to protect AWS GovCloud (US) regions and Azure Government subscriptions.

- VM-based licenses

Licenses are based on the number of virtual machines on all sources and servers that you plan to protect by using HYCU.

- File server licenses

You can use these licenses independently or in combination with standard licenses.

- Socket-based licenses

Licenses are based on the number of CPU sockets on all Nutanix clusters where the Nutanix Files servers that you plan to protect by using HYCU reside.

ⓘ Important This type of licenses is reserved only for Nutanix Files. If you would like to protect Dell PowerScale OneFS, NetApp ONTAP, file shares created by NetApp Cloud Volumes Service for

Google Cloud, S3, or generic file shares, contact your HYCU Sales representative.

- Capacity-based licenses

Licenses are based on the capacity of file server shares, which is calculated automatically as an overall size (in terabytes) of all protected file server shares.

- HYCU R-Cloud license

You can use this license in combination with other licenses to be able to migrate virtual machines across different infrastructures, perform disaster recovery of data to cloud, and monitor HYCU for AWS, HYCU for Google Cloud, and HYCU for Azure data protection environments.

Considerations

- When verifying that your license is valid, HYCU takes into account only the sources containing the entities with the Protected or Protected deleted status.
- The protection of the HYCU backup controller does not require any license.
- When a Managed Service Provider (MSP) license is applied to HYCU, sharing telemetry data with HYCU is enabled by default and cannot be disabled.
- *For Nutanix Community Edition (CE) environment:* No HYCU licenses are required.

Procedure

1. Buy a needed number of HYCU licenses. To discuss the options, contact your Sales representative.
2. Create a license request. For details, see “[Creating a license request](#)” on the [next page](#).
3. Request and obtain licenses from the web licensing portal. For details, see “[Requesting and retrieving licenses](#)” on page 388.
4. Activate the licenses to start using HYCU. For details, see “[Activating licenses](#)” on page 389.

Accessing the Licensing dialog box

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

Creating a license request

To obtain your HYCU licenses, you should submit a request form to the web licensing portal.

Prerequisites

- You bought the required number of HYCU licenses and have an entitlement order number.
- You added sources that you want to protect to the data protection environment. For instructions, see [“Adding sources” on page 53](#).

Procedure

1. In the Licensing dialog box, click **Download Request**.
2. Save the license request file to a temporary location.

Example

license.req file:

```
CN myCompany
ND C0F90A56-3FCC-4437-A49C-EFBA9BD8FC0F
VER V2
PID nutanixbackup
NRP 3
QTY 127
AFS 3
AFSCAP 4
NRPALL 12
QTYALL 167
HYCUVER 4.9.0-66
HSUD
83B770D4D02B9F9D516C9FAD7027F50AEF67C1F85209735165B7C500CCB3BF
DC
NEXT NODE
```

Requesting and retrieving licenses

After you create a license request file, you can obtain the licenses from the licensing portal.

Procedure

1. Connect to the web licensing portal at:

<https://licensing.hycu.com/>

2. If you already have a licensing portal account, click **Login**, enter your user name and password, and then click **Login**. Otherwise, create an account and then sign in with a newly created user account.
3. Click the **Activate licenses** link, and then enter the entitlement order number. Click **Next**.
4. Perform the following:
 - a. Browse for the license request file, and then click **Request License**.
 - b. In the Activate perpetual licenses page, specify the license types and the number of licenses you want to activate. By default, the number of licenses from the license request file is provided. You can specify a different value that may not exceed the number of purchased licenses. Click **Activate Licenses**.

Within a few minutes, you should receive an email with a license file `license.dat` attached.

Example

`license.dat` file:

```

CN myCompany
ND C0F90A56-3FCC-4437-A49C-EFBA9BD8FC0F
VER V2
PID nutanixbackup
EXP 05.04.2024
NRP 3
AFSCAP 5
LK
302C02146B7A48EE010CD1E1212E73B27DD2E58958B6C6ED021426BA2A4CCD
271CC45571A35129B7E8B4E46A75AD
NEXT NODE
  
```

5. Save the license file locally.

Activating licenses

After you submit your license request for the HYCU licenses to the web licensing portal, you get an email with a product license file attached.

Procedure

1. In the Licensing dialog box, click **Upload License**.
2. Browse for the license file that you received by email, and then click **Upload**.

After the licenses are activated, the information related to licensing is updated.

 **Note** You can always add new licenses for your growing environment. Contact your HYCU Sales representative.

On the General and Capacities tabs, you can check the following information related to licensing:

- Status
- License type
- HYCU backup controller ID
- License expiration date
- Managed service provider
- Protected and licensed number of virtual machines and servers
- Protected and licensed number of sockets for Nutanix Files
- Protected and licensed file server capacity
- Discovered, protected, and licensed number of sockets for virtual machines and servers

Setting up logging

You can set up logging to log information at various levels to help you analyze and troubleshoot the entire HYCU operation and diagnose issues with backup and restore performance.

Prerequisite

For sending log files to HYCU Support: Sharing telemetry data with HYCU must be enabled. For instructions, see “[Sharing telemetry data with HYCU](#)” on page 438.

Accessing the Logging dialog box

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

In the Logging dialog box, you can do the following:

- Download and view the existing log files by clicking **Get Logs**.

You download log files with the level that was specified at the time they were recorded. If logging is not set up, the log files are downloaded with the default settings. The changed logging level is applied only to the log files that are recorded after you save new logging settings.

After you extract the `zip` file, check the log files at the following location:

`/opt/grizzly/logs/`

- *Only if Sharing telemetry data with HYCU is enabled.* Send the existing log files to HYCU Support by clicking **Send Logs**.

You send log files with the level that was specified at the time they were recorded. If logging is not set up, the log files are uploaded with the default settings. The changed logging level is applied only to the log files that are recorded after you save new logging settings.

- Set up logging. To do so, follow these steps:
 1. Specify values for the following logging settings:

Logging setting	Description
Maximum log file size (MiB)	The maximum size of a log file. The default log file size is 200 MiB, whereas the maximum log file size is 2047 MiB.
Number of log files	The number of log files. The default number is 9.
Level	<p>The following logging levels are available:</p> <ul style="list-style-type: none"> ◦ Informational (default): Informational messages about the operation of HYCU are recorded to log files. ◦ Detailed: All activity is recorded to log files.
Outbound REST call level <i>(Available only if the Detailed logging level is selected.)</i>	<p>The following levels are available:</p> <ul style="list-style-type: none"> ◦ Off (default): Outbound REST call logs are not recorded to log files. ◦ Informational: Informational messages about the operations related to outbound REST calls are recorded to log files. ◦ Detailed: All activity related to outbound REST calls is recorded to log files.
Inbound REST call level <i>(Available only if the Detailed logging level is selected.)</i>	<p>The following levels are available:</p> <ul style="list-style-type: none"> ◦ Off (default): Inbound REST call logs are not recorded to log files. ◦ Informational: Informational messages about operations related to inbound REST calls are recorded to log files. ◦ Detailed: All activity related to inbound REST calls is recorded to log files.

2. Use the **Keep settings after upgrade** switch if you want the custom logging settings to remain the same after you upgrade HYCU. As you usually set logging for troubleshooting purposes and do not need the

same logging level for regular use of the product, by default, this switch is turned off.

3. Click **Save**.

 **Note** Keep in mind that the changed logging level is applied only to the log files that are recorded after you save new logging settings.

You can later modify the settings by specifying new values and clicking **Save**, or set the default values by clicking **Default**.

Configuring your network

When configuring your network, you can change network settings such as the IP address and the HYCU listening port number, or enable network bandwidth throttling. Depending on what you want to do, see one of the following sections:

- “[Changing network settings](#)” below
- “[Limiting network bandwidth](#)” on page 394

Accessing the Networks dialog box

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

Changing network settings

Changing network settings allows you to configure your network to suit the needs of your environment.

 **Important** After you make any changes to the HYCU network settings, you will be signed out automatically and your session will restart.

Limitations

- Multiple network adapters on the same network are not supported.
- *Only if you plan to make HYCU accessible through an IPv6 connection.*
 - The HYCU web user interface and REST API can be accessed through IPv6. Other IPv6 connections with HYCU are not supported.
 - The communication with HYCU instances uses only the IPv4 connections.

- The IPv6 address cannot be configured during the initial deployment of HYCU.
- Only static IPv6 network configuration is supported.
- The pure IPv6 network configuration without an existing IPv4 configuration is not supported.

Consideration

The network that you specified during the HYCU deployment is set to main and is represented by the  icon. If you later connect your HYCU backup controller to more than one network by using the Nutanix Prism web console or the vSphere (Web) Client, you can use another network as the main network. To do so, make sure that a listening port and an SSL certificate are specified for the preferred network, select this network, and then click  **Set Main**.

Procedure

1. In the Networks dialog box, on the General tab, the host name of your HYCU backup controller and the networks to which it is connected are displayed. Select the network for which you want to change settings, and then click  **Edit**.
2. Configure the IPv4 settings as required: the address, the gateway, and the netmask.
3. *Only if you want to make HYCU accessible through an IPv6 connection.* Use the **Enable IPv6 configuration** switch, and then configure the IPv6 settings as required: the address, the gateway, and the prefix.
4. *Optional.* Configure the domain name and the DNS as required.
5. *Only if your HYCU backup controller is connected to more than one network.* Use the **Enable listening on this port** switch if you want to use this network to access the HYCU web user interface.

 **Note** For the network that you specified during the HYCU deployment, this switch is enabled by default.

6. *Only if the Enable listening on this port switch is enabled.* In the Listening port field, enter the port that you want to use to access the HYCU web user interface (by default, 8443).
-  **Important** If a firewall is configured in your infrastructure, make sure that the port you specified is open.

7. From the SSL certificate drop-down menu, select the SSL certificate that you want to use for this network.
8. Click **Save**.

Limiting network bandwidth

Network bandwidth throttling allows you to limit the bandwidth that is available to HYCU. By defining sites with limited bandwidth, you ensure that enough bandwidth is available for all the network operations in your environment.

Limitation

You can limit network bandwidth only for traffic that is outbound from HYCU.

Considerations

- Network bandwidth throttling is not available in HYCU Manager.
- If the IP address of the storage container to which you plan to restore data is defined in a site for which you want to limit bandwidth, restore performance may be affected.
- Cloud, iSCSI, or SMB targets may utilize multiple IP addresses. Make sure to enter all the utilized IP addresses when defining a site. For details on IP ranges used by public clouds, see respective cloud documentation.
- Throttling network bandwidth for AWS IP addresses also affects telemetry data sharing. Sending log files may take longer.
- *Only if HYCU is used for file share protection.* If you enable network bandwidth throttling, the limit you set applies also to HYCU instances.

Recommendation

It is not recommended to throttle network bandwidth for NFS targets.

Procedure

1. In the Networks dialog box, click the **Throttling** tab, and then click  **New**.
2. Enter a name for the site for which you want to limit bandwidth and, optionally, its description.
3. In the Bandwidth limit field, specify the maximum speed (in KiBps, MiBps, or GiBps) that can be used to transfer data from HYCU to the site.

4. In the IP address/range list field, enter the IP addresses or IP ranges of the sites for which you want to limit bandwidth. You can enter the IP addresses or IP ranges in the following form:
 - Single IPv4 address: 192.0.2.1
 - IPv4 subnet with CIDR prefix: 192.0.2.0/24
 - IPv4 range: 192.0.2.3-192.0.2.100
5. *Only if you want to specify a throttling window.* Select the **Use throttling window** check box, and then, from the Throttling window drop-down menu, select the throttling window that you want to be used for limiting bandwidth. For details on how to create a throttling window, see “[Creating a throttling window](#)” below.

By clicking **Manage**, you are automatically directed to the Throttling Windows dialog box from where you can manage throttling windows.

! Important If you define multiple sites with the same IP addresses, make sure the throttling windows that you assign to these sites do not overlap.

6. Click **Save**.

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

Creating a throttling window

HYCU enables you to define time frames for network bandwidth throttling. If you use a throttling window, network bandwidth is limited only within the specified hours. For example, you can limit network bandwidth during peak production hours when there is more activity on the network.

Procedure

1. In the Networks dialog box, click the **Throttling** tab, and then click **Throttling Windows**.
2. Click  **New**.
3. In the Name field, enter a name for the throttling window.
4. From the Time zone drop-down menu, specify the time zone for the throttling window. You can also click one of the displayed time zones (your local time zone or your HYCU backup controller time zone).

5. Select the weekdays and hours during which you want network bandwidth to be limited.

Tip You can click and drag to quickly select a time frame that includes the days and hours you want to add.

6. Click **Save**.

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

Managing data retention

When data reaches the end of the retention period defined in the policy, it is automatically expired and deleted from targets by the HYCU cleaning process. For details on automatic data expiration, see “[Expiring data automatically](#)” on [page 401](#).

However, if there is a restore point tier (Backup, Copy, Snapshot, and/or Archive) that you do not want to use for restoring data anymore, HYCU enables you to expire it also manually at any time. You can do this by choosing one of the following methods:

Expiration method	Prerequisite	Instructions
Select restore point tiers for expiration from the Retention Management dialog box.	You must be familiar with the information related to expiring data described in this section.	“ Procedure ” on page 399
Select restore point tiers for expiration from the Virtual Machines, Applications, Shares, or Volume Groups panel.		“ Procedure ” on page 403

As part of data retention management, HYCU also allows you to set the preferred expiration date and time for restore points tiers, as well as to specify how long restore point tiers should be retained.

Tip You can check the backup, copy, snapshot, and archive expiration time in the Detail view of the Virtual Machines, Applications, Shares, or Volume Groups panel. For details, see “[Viewing entity details](#)” on [page 313](#).

When managing data retention, you can select among the following options:

Option	Enables you to...
Expire restore point tiers	<p>Expire the selected restore point tiers. All subsequent restore point tiers in the same backup chains will also be expired (except for the file share restore point tiers with incremental forever backups in the last backup chain).</p> <p>For details on incremental forever backups for file shares, see “Configuring file share backup options” on page 264.</p>
Set expiration date	Set the date and time at which the selected restore point tiers will expire.
Set retention period	Specify how long the selected restore point tiers should be retained. The expiration time is calculated based on the restore point creation time.

Prerequisite

You must have an Administrator user role assigned.

Considerations

- Your user role determines the restore point tiers that are visible to you. If you are a member of the infrastructure group, you can view all restore point tiers in the data protection environment. If you are a self-service group administrator, you can view only restore point tiers related to the entities whose owners belong to your self-service group.
- *Only if expiring restore point tiers.*
 - An expire action cannot be undone.
 - If you expire all restore point tiers belonging to a restore point, the backup status is shown as Expired (○). This indicates that the restore point cannot be used for restoring data anymore.
 - Restore point tiers other than Backup that have the Failed backup status are not available for expiration.
 - If you selected Target as the backup target type in your policy, the following applies:
 - *For all entities except file shares with the enabled Incremental forever backup option:*

- If the most recent restore point is expired, the next backup will be a full backup.
- When a restore point is expired, any subsequent incremental backups within the same backup chain will also be expired unless the status of the selected restore point is Failed. In this case, only the selected restore point is expired and not the whole backup chain.
- If you expire the Backup restore point tier, the associated snapshot is also expired, if there is one.
- *For file shares with the enabled Incremental forever backup option:* If you are expiring restore point tiers belonging to the last backup chain, consider the following:
 - When a restore point is expired, any subsequent incremental backups within the same backup chain are not expired and the restore point is merged into the next restore point in the backup chain.
 - If you want to expire the entire backup chain, all restore points in the backup chain must be expired. In this case, the next backup will be a full backup.
 - The most recent restore point is protected and cannot be expired separately, but only as part of expiring the entire backup chain. In addition, merging other restore points into the most recent one is not possible.

For details on the Incremental forever backup option, see [“Configuring file share backup options” on page 264](#).

- *Only if expiring data for a volume group that has been backed up both as part of the virtual machine backup and by assigning a policy directly to it.* Before expiring data, make sure that the data is not being used by any virtual machine to which the volume group might be attached.
- *Only if expiring data for a file share with the enabled Incremental forever backup option.* If your restore point tier data is stored on a target that has WORM protection enabled, removing the data from the target will be performed when the data is no longer WORM protected.
- You cannot set the expiration date or the retention period for restore point tiers whose backup status is Failed, for the Snapshot restore point tiers created for the fast restore, or for the latest Snapshot restore point tiers.

- Only if setting the retention period for restore point tiers with active WORM retention. The retention period of restore point tiers with active WORM retention is not modified if the WORM target retention period is longer than the specified one.

Accessing the Retention Management dialog box

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

Procedure

In the Retention Management dialog box, select the action that you want to perform, click **Continue**, and then follow the instructions:

Option	Instructions
Expire restore point tiers	<ol style="list-style-type: none"> From the list of all restore point tiers, select the ones that you want to expire, and then click Continue. <p> Tip You can use filtering options to display only specific types of restore point tiers. For details, see “Filtering restore point tiers” on the next page.</p> Review the expiration information, and then click Expire. Click Expire to confirm that you want the selected restore point tiers to be expired.
Set expiration date	<ol style="list-style-type: none"> From the list of all restore point tiers, select the ones for which you want to set an expiration date, and then click Continue. <p> Tip You can use filtering options to display only specific types of restore point tiers. For details, see “Filtering restore point tiers” on the next page.</p> In the Expiration date and time field, specify the date and time when you want the selected restore point tiers to be expired. Review the expiration information, and then click Execute.
Set retention period	<ol style="list-style-type: none"> From the list of all restore point tiers, select the ones for which you want to set a retention period, and then click Continue.

Option	Instructions
	<p>Tip You can use filtering options to display only specific types of restore point tiers. For details, see “Filtering restore point tiers” below.</p> <p>2. Depending on whether you want to set a new retention period or to modify the retention period for the selected restore point tiers, select one of the following options, and then do as required:</p> <ul style="list-style-type: none"> Set new retention period: Set a retention period for the data (in months, weeks, days, or hours). Adjust current retention period: Select Increase or Decrease, and then set for how long you want to increase or decrease the retention period of the data (in months, weeks, days, or hours). <p>Important Keep in mind that setting a new retention period or modifying the retention period for the selected restore point tiers can affect the retention period of other restore point tiers in the same backup chains.</p> <p>3. Review retention information, and then click Execute.</p>

The next HYCU cleaning process removes data from the targets.

Filtering restore point tiers

The restore point tier filtering options allow you to filter and focus only on specific types of restore point tiers (for example, the restore point tiers whose data is stored on a particular target). After you apply any of the filters, only restore point tiers that match the filter criteria are displayed and you can easily find what you need.

Tip When you specify any of the filtering options and you want to include all available categories, click **Select all**.

Procedure

1. In the side panel of the Select Restore Point Tiers dialog box that opens after you select the preferred retention management action, select one or more of

the following filtering options:

Filtering option	Filter by...
Tiers	Restore point tiers (Backup, Copy, Snapshot, Archive daily, Archive weekly, Archive monthly, and Archive yearly).
Backup type	Type of backup (Incremental and Full).
Restore point creation date	Time range within which the restore point was created.
Restore point tier expiration date	Time range within which the restore point tier is marked for expiration.
Source	Sources that host the entities with the restore point tiers.
Target	Targets on which restore point tier data is stored.
Type	Entities with the restore point tiers.
Owner	Owners that are assigned to the entities with the restore point tiers.
Backup status	Backup statuses (Success, Failure, and Warning).

2. Click **Apply Filters**.

Expiring data automatically

When any of the restore point tiers reaches its retention period, it is grayed out in the HYCU web user interface. Depending on which backup target type you selected in your policy, tiers are expired as follows:

Backup target type	Conditions for tier expiration
Target	<i>For all entities except file shares with the enabled Incremental forever backup option:</i> A tier is expired when the last tier in the backup chain reaches its retention period. This means that this data is not removed from HYCU or the target until all tiers in the backup chain are expired. However, if there is a

Backup target type	Conditions for tier expiration
	<p>restore point that contains the Archive tier, this restore point is kept although the rest of the backup chain is expired. In addition, if this restore point is an incremental backup, it is changed to full.</p> <p><i>For file shares with the enabled Incremental forever backup option:</i> A tier is expired when it reaches its retention period and is merged into the next restore point. However, if there is a restore point that contains the Archive tier, this restore point is kept.</p>
Snapshot	<p>A tier is expired when the snapshot reaches its retention period. However, if there is a restore point that contains the Archive tier, this restore point is kept although the snapshot is expired.</p>

Considerations

- Changing the retention period in the policy does not affect existing backups.
- HYCU automatically expires the last backup chain of an unprotected entity (the one from which a policy was unassigned or whose policy was deleted), whereas the last backup chain of a protected entity is never expired automatically.

Expiring data from the entity panel

As an alternative to selecting restore point tiers for expiration from the Retention Management dialog box, you can select the preferred restore point tiers also from the Virtual Machines, Applications, Shares, or Volume Groups panel.

Depending on the entity for which you want to expire data, access one of the following panels:

- Accessing the Virtual Machines panel
To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

- Accessing the Applications panel
To access the Applications panel, in the navigation pane, click  **Applications**.
- Accessing the Shares panel
To access the Shares panel, in the navigation pane, click  **Shares**.
- Accessing the Volume Groups panel
To access the Volume Groups panel, in the navigation pane, click  **Volume Groups**.

Procedure

1. In the Virtual Machines, Applications, Shares, or Volume Groups panel, select the entity for which you want to expire data.
2. Select one or more restore points that contain the restore point tiers that you want to expire, and then click  **Expire**.
3. From the list of available restore point tiers, select the ones that you want to expire, and then click **Continue**. You are redirected to the Retention Management > Expire dialog box.
4. Review the expiration information, and then click **Expire**.
5. Click **Expire** to confirm that you want the selected restore point tiers to be expired.

Setting power options

You can set power options for the HYCU backup controller so that its activities are suspended or resumed.

Accessing the Power Options dialog box

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

Power option	Description
Suspend all	Pauses all HYCU backup controller activities. If you want the HYCU backup controller activities to automatically resume after a specified amount of time, in the

Power option	Description
	<p>Auto resume after field, specify the number of hours (1–168) to pass before the activities are resumed.</p> <p>All currently running jobs are allowed to complete normally. All jobs that are in the queue will start when the HYCU backup controller is resumed. While activities are paused, you cannot start any new backup jobs.</p>
Suspend cleanup	<p>Pauses the cleanup of targets and, if enabled, the purge of events and jobs.</p> <p>The snapshot cleanup is not affected.</p>
Resume	Allows HYCU backup controller activities to continue.

Managing secrets

HYCU enables you to securely store, access, and manage your credentials (secrets) by employing the Conjur secrets management solution. After you store your HYCU secrets in Conjur as one or more Conjur configurations (that is, a set of one or more policies where you define your security rules), you can easily manage them and be confident that your resources can be accessed only by authorized parties.

Prerequisites

- You have set up your Conjur environment and stored HYCU secrets as a set of one or more policies. For instructions, see Conjur documentation.
- The SSL certificate of the Conjur server must be imported to HYCU by an infrastructure administrator. For instructions, see “[Configuring SSL certificates](#)” on page 432.

Limitations

- HYCU credentials that you plan to store in Conjur may not start with \${ and end with }.
- HYCU users cannot be managed by using Conjur. For details on HYCU users, see “[Managing users](#)” on page 343.

Considerations

- Mixed mode is possible. This means that you do not have to store all your HYCU secrets in Conjur to be able to benefit from this integration.
- When providing secrets that are stored in Conjur, make sure to use the following syntax in HYCU:

`${<Path to Secret>}`

Recommendation

Only if you plan to change the names of secrets in Conjur. Every time you change the name of a secret in Conjur, it is recommended that you clear the cache in HYCU. To do so, in the Secrets Management dialog box, click  **Clear Cache**. This is done also automatically by HYCU every 24 hours, but for the continuity of your business processes, it is recommended that you do it manually.

 **Tip** All the fields in the HYCU web user interface that support entering values stored in Conjur have the following icon next to them:  for infrastructure group configuration and  for private configuration.

Accessing the Secrets Management dialog box

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

Adding a Conjur configuration

Consideration

For each HYCU data protection environment, one infrastructure group Conjur configuration and one private Conjur configuration for each self-service group can be added.

Procedure

1. In the Secrets Management dialog box, depending on which type of Conjur configuration you want to perform, click one of the following buttons:

Conjur configuration type	Description
Add Infrastructure Group Configuration	<i>Available only if you are an infrastructure group administrator.</i> Enables you to provide secrets stored in Conjur when performing

Conjur configuration type	Description
	all data protection and administrative tasks. For example, when adding sources and targets, adding identity providers, and so on.
Add Private Configuration	<p><i>Available if you are an infrastructure or a self-service group administrator.</i> Enables you to provide secrets stored in Conjur when performing the following tasks:</p> <ul style="list-style-type: none"> • Adding cloud accounts. • Assigning credential groups to virtual machines. • Setting up webhook notifications.

2. In the Appliance URL field, enter the URL of the Conjur server that you are connecting to.
3. In the Account field, enter the name of the account that you specified during the Conjur environment setup.
4. In the Authentication login field, enter the Conjur host user name. For example:

```
host/HycuPolicy/hycuBckupController
```

In this example, host is the type of user, HycuPolicy is the name of the policy to which the user belongs, and hycuBckupController is the user name.

5. In the Authentication API key field, enter the API key that corresponds to the Conjur host user name.
6. *Only if you are an infrastructure group administrator.* When performing one type of Conjur configuration, enable the **Use same values for private configuration** or **Use same values for infrastructure group configuration** switch if you want to use the same values for performing the other type of Conjur configuration.
7. Click **Save**.

Editing a Conjur configuration

Procedure

1. In the Secrets Management dialog box, click  **Edit** next to the Conjur configuration that you want to edit.
2. Edit the selected Conjur configuration as required. For detailed information on Conjur configuration properties, see “[Adding a Conjur configuration](#)” on [page 405](#).
3. Click **Save**.

Deleting a Conjur configuration

Procedure

1. In the Secrets Management dialog box, click  **Delete** next to the Conjur configuration that you want to delete.
2. Click **Yes** to confirm that you want to delete the selected Conjur configuration.

Configuring an SMTP server

Before enabling HYCU to send email notifications, you must configure an SMTP server that HYCU will use.

Prerequisite

For using the STARTTLS or SSL/TLS security mode to secure email traffic: A valid SSL certificate must be imported to HYCU. For instructions, see “[Securing SMTP connections](#)” on [page 451](#).

Limitation

HYCU supports only basic SMTP authentication.

Accessing the SMTP Server Settings dialog box

To access the SMTP Server Settings dialog box, click  **Administration**, and then select **SMTP Server Settings**.

Procedure

1. In the SMTP Server Settings dialog box, provide the following information:

Required information	Description
Username	User name of the account on the SMTP server.
Password	Password of the account on the SMTP server.
Display name	Display name of the email sender.
Hostname or IP address	Host name or IP address of the SMTP server.
Port	Port number to be used (usually set to 25).
Security mode	Protocol used to secure email traffic—can be set to None, STARTTLS, or SSL/TLS.
From email address	Email address from which email notifications will be sent.

2. If you want to verify that the provided SMTP configuration is correct by sending a test email with the SMTP server settings, do the following:
 - a. Enable the **Send test email** switch.
 - b. In the Test email recipient field, enter an email recipient that should receive the test email with the SMTP server settings.
3. Click **Save**.

After you configure an SMTP server, you can continue with configuring HYCU to send email notifications. For details on how to do this, see “[Setting up email notifications](#)” on page 302.

To reconfigure your existing SMTP server, delete the current server settings by clicking **Delete Settings**, and then redefine them.

Upgrading HYCU

You can upgrade HYCU when a new software release version is available.

Prerequisites

- The HYCU backup controller activities must be suspended. For instructions, see “[Setting power options](#)” on page 403.
- Jobs that you do not want to be aborted must be finished (the upgrade process aborts all currently running jobs).
- The HYCU data disk must be larger than the HYCU system disk. For instructions on how to increase disk size, see “[Increasing the size of the HYCU virtual disks](#)” on page 455.

Considerations

- Any users that were signed in to the HYCU web user interface of the HYCU backup controller that is being upgraded should perform a hard reload of the web user interface page in their web browser after the process completes.
- Upgrading removes any previously added update packages from the update folder on the HYCU backup controller.

To upgrade HYCU, see one of the following sections:

Upgrade HYCU...	Instructions
On a Nutanix AHV cluster	“ Upgrading HYCU on a Nutanix AHV cluster ” on the next page
On a Nutanix ESXi cluster	“ Upgrading HYCU on a Nutanix ESXi cluster ” on page 412
In a vSphere environment	“ Upgrading HYCU in a vSphere environment ” on page 417
In an AWS GovCloud (US) environment	“ Upgrading HYCU in an AWS GovCloud (US) environment ” on page 423
In Google Cloud	“ Upgrading HYCU in Google Cloud ” on page 424
In an Azure Government environment	“ Upgrading HYCU in an Azure Government environment ” on page 426

Upgrading HYCU on a Nutanix AHV cluster

Prerequisites

- The Nutanix AHV cluster where the HYCU backup controller and HYCU instances reside must be added to HYCU. For instructions, see [“Adding a Nutanix cluster” on page 54](#).
- The HYCU system disk must be selected as the boot device in the Disks section of the Update VM dialog box in the Nutanix Prism web console.
- The state of the HYCU virtual appliance image that you want to use for an upgrade must be ACTIVE in the Nutanix Prism image service.

For details, see Nutanix documentation.

Considerations

- If you are using HYCU for file share protection, the HYCU instances residing on a Nutanix AHV cluster are upgraded automatically during the HYCU upgrade process if the following is true:
 - The Nutanix cluster where the HYCU instances reside is added to HYCU.
 - The HYCU virtual appliance image is present on the same Nutanix cluster in the following format:
`hycu-<Version>-<Revision>`
 For example, `hycu-5.0.0-3634`.

Otherwise, follow the HYCU upgrade procedure to perform the HYCU instance upgrade.

- After you make sure HYCU was upgraded successfully, you can safely delete the old HYCU backup controller virtual machine from the Nutanix AHV cluster.
- If the HYCU backup controller is part of a Nutanix protection domain (the recommended approach), make sure that the new version of the HYCU backup controller virtual machine is included in this protection domain after the upgrade.

Procedure

1. Sign in to the Nutanix Prism web console, and then upload the HYCU virtual appliance image that you want to use for an upgrade to your Nutanix AHV cluster as follows:

- a. Click  , and then select **Image Configuration**.
- b. In the Image Configuration dialog box, click **Upload Image**.
- c. In the Create Image dialog box, provide the following information:
 - i. Enter a HYCU image name in the format that should correspond to that of the HYCU image file you are uploading.

! Important The HYCU virtual appliance image must be uploaded to the Nutanix AHV cluster in the following format:
`hycu-<Version>-<Revision>`

For example: `hycu-5.0.0-3634`

If you enter the HYCU image name in a different format, you will not be able to use this image for an upgrade.

- ii. *Optional.* Enter an annotation.
- iii. From the Image Type drop-down menu, select **DISK**.
- iv. From the Storage Container drop-down menu, select a storage container for the image to be uploaded.
- v. In the Image Source section, specify the location of the image file.
- vi. Click **Save**.
- vii. Click **Close** after the image is successfully uploaded.

2. Sign in to the HYCU web user interface, and then do as follows:
 - a. Click  **Administration**, and then select **Software Upgrade**.
 - b. In the Software Upgrade dialog box, on the Release tab, check the current version of HYCU and all available versions. You can also check whether any newer version is available on the HYCU Support portal by clicking the **Check for new version** link.
 - c. Select the version to which you want to upgrade HYCU.
 - d. Click **Software Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.
3. *Only if HYCU is used for file share protection.* If the Nutanix cluster where the HYCU instances reside is not added to HYCU or the appropriate HYCU virtual appliance image is not present on the same Nutanix cluster, upgrade the HYCU instances as follows:
 - a. Remove the existing HYCU instances. For details on how to do this, see “[Deleting a HYCU instance](#)” on page 384.

- b. Create new HYCU instances with the latest HYCU version. For details on how to do this, see “[Creating a HYCU instance by using the HYCU web user interface](#)” on page 382.

You will be signed out of HYCU and you can track the upgrade progress in the Nutanix Prism web console as follows:

- The old HYCU backup controller virtual machine will remain on the Nutanix AHV cluster and will be renamed to `<HYCUBackupControllerName>_version_<OldHYCUVersion>`.
- The new upgraded HYCU backup controller virtual machine will replace the old one.
- The upgraded HYCU backup controller virtual machine will be powered on automatically.

After the upgrade process completes, you can sign in to the HYCU web user interface.

① Important Before you sign in to the HYCU web user interface again, make sure to perform a hard reload of its webpage in your web browser.

Upgrading HYCU on a Nutanix ESXi cluster

To upgrade HYCU on a Nutanix ESXi cluster, you can choose one of the following approaches:

Upgrade approach ^a	Instructions
By importing the HYCU OVF package to a content library.	“ Upgrading HYCU by importing the HYCU OVF package to a content library ” on the next page
By deploying the HYCU OVF package to a vCenter Server inventory.	“ Upgrading HYCU by deploying the HYCU OVF package to a vCenter Server inventory ” on page 415

^aIf HYCU is used for file share protection, you must upgrade HYCU by importing the HYCU OVF package to a content library so that the HYCU instances connected to your HYCU backup controller can be upgraded. For instructions, see “[Upgrading HYCU instances on a Nutanix ESXi cluster](#)” on page 417.

Prerequisites

- The Nutanix ESXi cluster where the HYCU backup controller and HYCU instances reside must be added to HYCU. For instructions, see “[Adding a Nutanix cluster](#)” on page 54.
- As a vSphere user, you must have the required upgrade privileges. For details on upgrade privileges, see “[Assigning privileges to a vSphere user](#)” on page 456.
- A snapshot of the HYCU backup controller must be created by using the Nutanix protection domain. For details, see Nutanix documentation.
- Any HYCU snapshots created by using VMware vSphere must be removed.

Considerations

- After you make sure HYCU was upgraded successfully, you can safely delete the old HYCU backup controller virtual machine from the Nutanix ESXi cluster.
- If the HYCU backup controller is part of a Nutanix protection domain (the recommended approach), make sure that the new version of the HYCU backup controller virtual machine is included in this protection domain after the upgrade.
- After you upgrade HYCU or HYCU instances, on some Nutanix ESXi clusters you might get an error message that there is a MAC address conflict. You can safely ignore this message.

Upgrading HYCU by importing the HYCU OVF package to a content library

① Important You can use either the vSphere Web Client or the vSphere Client as the interface for performing the procedure described in this section. As an example, you are guided through the steps that you must perform if you are using the vSphere Web Client.

Procedure

1. Sign in to the vSphere Web Client, and then do as follows:
 - a. Navigate to the content library to which you want to import the HYCU OVF package.
 - b. Right-click your content library, and then select **Import Item**. The Import Library Item dialog box opens.

c. In the Source section, specify the location of the OVF package:

URL	Specify a URL to the HYCU OVF package.
	Browse your file system for the HYCU OVF package.
Local file	<p>! Important When you are browsing your file system, make sure to select both the .ovf file and the .vmdk file related to the OVF package.</p>

Click **OK**.

d. In the Destination section, enter a name and description for the item, and then click **OK**.

! Important Make sure the item name you enter matches the HYCU OVF package name. For example, `hycu-5.0.0-3634`.

2. Sign in to the HYCU web user interface, and then do as follows:

- Click  **Administration**, and then select **Software Upgrade**.
- In the Software Upgrade dialog box, on the Release tab, check the current version of HYCU and all available versions. You can also check whether any newer version is available on the HYCU Support portal by clicking the **Check for new version** link.
- Select the version to which you want to upgrade HYCU.

! Tip The icon next to each version shows the location of the HYCU upgrade image, **CL** (a content library) or **VC** (a vCenter Server inventory).

- Click **Software Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

You will be signed out of HYCU and you can track the upgrade progress in the Nutanix Prism web console as follows:

- The old HYCU backup controller virtual machine will remain on the Nutanix ESXi cluster and will be renamed to `<HYCUBackupControllerName>_version_<OldHYCUVersion>`.
- The new upgraded HYCU backup controller virtual machine will replace the old one.
- The upgraded HYCU backup controller virtual machine will be powered on automatically.

After the upgrade process completes, you can sign in to the HYCU web user interface.

① Important Before you sign in to the HYCU web user interface again, make sure to perform a hard reload of its webpage in your web browser.

Upgrading HYCU by deploying the HYCU OVF package to a vCenter Server inventory

① Important You can use either the vSphere Web Client or the vSphere Client as the interface for performing the procedure described in this section. As an example, you are guided through the steps that you must perform if you are using the vSphere Web Client.

Procedure

1. Sign in to the vSphere Web Client, and then do as follows:
 - a. Right-click your vCenter Server, and then select **Deploy OVF Template....** The Deploy OVF Template dialog box opens.
 - b. In the Select template section, specify the location of the OVF package:

URL	Specify a URL to the HYCU OVF package.
	Browse your file system for the HYCU OVF package.
Local file	<p>① Important When you are browsing your file system, make sure to select both the .ovf file and the .vmdk file related to the OVF package.</p>

Click **Next**.

- a. In the Select name and location section, enter a name for the HYCU backup controller virtual machine and specify a location where you want to deploy it, and then click **Next**.

① Important Make sure the virtual machine name you enter matches the HYCU OVF package name. For example, `hycu-5.0.0-3634`.
- d. In the Select a resource section, select where to run the deployed package, and then click **Next**.
- e. In the Review details section, verify the package details, and then click **Next**.

- f. In the Select storage section, select where to store the files for the deployed package, and then click **Next**.
- g. In the Select networks section, select a destination network, and then click **Next**.
- h. In the Customize template section, enter the values for the following:
 - *Optional.* Host name for the virtual machine

Note The default host name is generated automatically during the HYCU virtual appliance deployment. The host name should begin with a letter and may contain only letters, numbers, and hyphens (-).

 - IPv4 address (for example, 10.1.100.1)
 - Subnet mask (for example, 255.0.0.0)
 - Default gateway (for example, 10.1.1.1)
 - *Optional.* DNS server (for example, 10.1.1.5)
 - *Optional.* Search domain (for example, domain.com)

Note The domain name should begin with a letter and contain one or more periods. It may also contain only letters, numbers, and hyphens (-).

Click **Next**.

- i. In the Ready to complete section, review data, and then click **Finish**.
2. Sign in to the HYCU web user interface, and then do as follows:
 - a. Click  **Administration**, and then select **Software Upgrade**.
 - b. In the Software Upgrade dialog box, on the Release tab, check the current version of HYCU and all available versions. You can also check whether any newer version is available on the HYCU Support portal by clicking the **Check for new version** link.
 - c. Select the version to which you want to upgrade HYCU.

Tip The icon next to each version shows the location of the HYCU upgrade image,  (a content library) or  (a vCenter Server inventory).

- d. Click **Software Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

You will be signed out of HYCU and you can track the upgrade progress in the Nutanix Prism web console as follows:

- The old HYCU backup controller virtual machine will remain on the Nutanix ESXi cluster and will be renamed to `<HYCUBackupControllerName>_version_<OldHYCUVersion>`.
- The new upgraded HYCU backup controller virtual machine will replace the old one.
- The upgraded HYCU backup controller virtual machine will be powered on automatically.

After the upgrade process completes, you can sign in to the HYCU web user interface.

① Important Before you sign in to the HYCU web user interface again, make sure to perform a hard reload of its webpage in your web browser.

Upgrading HYCU instances on a Nutanix ESXi cluster

An upgrade of the HYCU instances residing on a Nutanix ESXi cluster starts automatically after the HYCU upgrade if the HYCU OVF package is imported to the vCenter Server content library and its format is as follows:

`hycu-<Version>-<Revision>`

For example, `hycu-5.0.0-3634`.

Otherwise, you must upgrade the HYCU instances manually as follows:

1. Remove the existing HYCU instances. For instructions, see “[Deleting a HYCU instance](#)” on page 384.
2. Create new HYCU instances by using the latest HYCU version. For instructions, see “[Creating a HYCU instance by using the HYCU web user interface](#)” on page 382.

>Note If you made any changes to the default user credentials, you can use only the default operating system user credentials (hycu, hycu/4u) after the HYCU instance upgrade. You can later make modifications to meet the needs of your environment.

Upgrading HYCU in a vSphere environment

To upgrade HYCU in a vSphere environment, you can choose one of the following approaches:

Upgrade approach ^a	Instructions
By importing the HYCU OVF package to a content library.	“Upgrading HYCU by importing the HYCU OVF package to a content library” on the next page
By deploying the HYCU OVF package to a vCenter Server inventory.	“Upgrading HYCU by deploying the HYCU OVF package to a vCenter Server inventory” on page 420

^aIf HYCU is used for file share protection, you must upgrade HYCU by importing the HYCU OVF package to a content library so that the HYCU instances connected to your HYCU backup controller can be upgraded. For instructions, see “Upgrading HYCU instances in a vSphere environment” on page 422.

Prerequisites

- The vCenter Server that manages the ESXi hosts where the HYCU backup controller and HYCU instances reside must be added to HYCU. For instructions, see “Adding a vCenter Server” on page 57.
- As a vSphere user, you must have the required upgrade privileges. For details on upgrade privileges, see “Assigning privileges to a vSphere user” on page 456.
- *For importing the HYCU OVF package to a content library:* A content library must be created in the vSphere (Web) Client.

Considerations

- *For upgrading HYCU if the HYCU backup controller is connected to a distributed switch:* After the upgrade, the port configured on the upgraded HYCU backup controller is different from the distributed switch port configured on the old HYCU backup controller. If you need your upgraded HYCU backup controller to use the same port as before, delete the port on the old HYCU backup controller, and then modify the port number in the new HYCU backup controller settings. For details on how to do this, see VMware documentation.
- It is not recommended that the HYCU backup controller is deployed on a VMware Virtual SAN (vSAN) datastore. However, if this is your case, before upgrading HYCU, contact HYCU Support.
- After you make sure HYCU was upgraded successfully, you can safely delete the old HYCU backup controller virtual machine from the vSphere environment.

- After you upgrade HYCU or HYCU instances, in some vSphere environments you might get an error message that there is a MAC address conflict. You can safely ignore this message.

Upgrading HYCU by importing the HYCU OVF package to a content library

ⓘ Important You can use either the vSphere Web Client or the vSphere Client as the interface for performing the procedure described in this section. As an example, you are guided through the steps that you must perform if you are using the vSphere Web Client.

Procedure

1. Sign in to the vSphere Web Client, and then do as follows:
 - a. Navigate to the content library to which you want to import the HYCU OVF package.
 - b. Right-click your content library, and then select **Import Item**. The Import Library Item dialog box opens.
 - c. In the Source section, specify the location of the OVF package:

URL	Specify a URL to the HYCU OVF package.
	Browse your file system for the HYCU OVF package.
Local file	ⓘ Important When you are browsing your file system, make sure to select both the .ovf file and the .vmdk file related to the OVF package.

Click **OK**.

2. In the Destination section, enter a name and description for the item, and then click **OK**.

ⓘ Important Make sure the item name you enter matches the HYCU OVF package name. For example, `hycu-5.0.0-3634`.
2. Sign in to the HYCU web user interface, and then do as follows:
 - a. Click  **Administration**, and then select **Software Upgrade**.
 - b. In the Software Upgrade dialog box, on the Release tab, check the current version of HYCU and all available versions. You can also check whether any newer version is available on the HYCU Support portal by clicking the **Check for new version** link.

- c. Select the version to which you want to upgrade HYCU.

 **Tip** The icon next to each version shows the location of the HYCU upgrade image, **CL** (a content library) or **VC** (a vCenter Server inventory).

- d. Click **Software Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

You will be signed out of HYCU and you can track the upgrade progress in the vSphere (Web) Client as follows:

- The old HYCU backup controller virtual machine will remain in the vSphere environment and will be renamed to `<HYCUBackupControllerName>_version_<OldHYCUVVersion>`.
- The new upgraded HYCU backup controller virtual machine will replace the old one.
- The upgraded HYCU backup controller virtual machine will be powered on automatically.

After the upgrade process completes, you can sign in to the HYCU web user interface.

 **Important** Before you sign in to the HYCU web user interface again, make sure to perform a hard reload of its webpage in your web browser.

Upgrading HYCU by deploying the HYCU OVF package to a vCenter Server inventory

 **Important** You can use either the vSphere Web Client or the vSphere Client as the interface for performing the procedure described in this section. As an example, you are guided through the steps that you must perform if you are using the vSphere Web Client.

Procedure

1. Sign in to the vSphere Web Client, and then do as follows:
 - a. Right-click your vCenter Server, and then select **Deploy OVF Template....**. The Deploy OVF Template dialog box opens.
 - b. In the Select template section, specify the location of the OVF package:

URL

Specify a URL to the HYCU OVF package.

	Browse your file system for the HYCU OVF package.
Local file	<p>Important When you are browsing your file system, make sure to select both the .ovf file and the .vmdk file related to the OVF package.</p>

Click **Next**.

- c. In the Select name and location section, enter a name for the HYCU backup controller virtual machine and specify a location where you want to deploy it, and then click **Next**.

Important Make sure the virtual machine name you enter matches the HYCU OVF package name. For example, `hycu-5.0.0-3634`.

- d. In the Select a resource section, select where to run the deployed package, and then click **Next**.
- e. In the Review details section, verify the package details, and then click **Next**.
- f. In the Select storage section, select where to store the files for the deployed package, and then click **Next**.
- g. In the Select networks section, select a destination network, and then click **Next**.

Important Make sure not to select a vSphere distributed switch (dvSwitch) for the virtual NIC option.

- h. In the Customize template section, enter the values for the following:
 - *Optional.* Host name for the virtual machine

Note The default host name is generated automatically during the HYCU virtual appliance deployment. The host name should begin with a letter and may contain only letters, numbers, and hyphens (-).

- IPv4 address (for example, `10.1.100.1`)
- Subnet mask (for example, `255.0.0.0`)
- Default gateway (for example, `10.1.1.1`)
- *Optional.* DNS server (for example, `10.1.1.5`)
- *Optional.* Search domain (for example, `domain.com`)

Note The domain name should begin with a letter and contain one or more periods. It may also contain only letters, numbers, and hyphens (-).

Click **Next**.

- i. In the Ready to complete section, review data, and then click **Finish**.
2. Sign in to the HYCU web user interface, and then do as follows:
 - a. Click  **Administration**, and then select **Software Upgrade**.
 - b. In the Software Upgrade dialog box, on the Release tab, check the current version of HYCU and all available versions. You can also check whether any newer version is available on the HYCU Support portal by clicking the **Check for new version** link.
 - c. Select the version to which you want to upgrade HYCU.

Tip The icon next to each version shows the location of the HYCU upgrade image, **CL** (a content library) or **VC** (a vCenter Server inventory).

- d. Click **Software Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

You will be signed out of HYCU and you can track the upgrade progress in the vSphere (Web) Client as follows:

- The old HYCU backup controller virtual machine will remain in the vSphere environment and will be renamed to `<HYCUBackupControllerName>_version_<OldHYCUVVersion>`.
- The new upgraded HYCU backup controller virtual machine will replace the old one.
- The upgraded HYCU backup controller virtual machine will be powered on automatically.

After the upgrade process completes, you can sign in to the HYCU web user interface.

Important Before you sign in to the HYCU web user interface again, make sure to perform a hard reload of its webpage in your web browser.

Upgrading HYCU instances in a vSphere environment

An upgrade of the HYCU instances residing in a vSphere environment starts automatically after the HYCU upgrade if the HYCU OVF package is imported to the vCenter Server content library and its format is as follows:

`hycu-<Version>-<Revision>`

For example, `hycu-5.0.0-3634`.

Otherwise, you must upgrade the HYCU instances manually as follows:

1. Remove the existing HYCU instances. For instructions, see [“Deleting a HYCU instance” on page 384](#).
2. Create new HYCU instances by using the latest HYCU version. For instructions, see [“Creating a HYCU instance by using the HYCU web user interface” on page 382](#).

 **Note** If you made any changes to the default user credentials, you can use only the default operating system user credentials (`hycu`, `hycu/4u`) after the HYCU instance upgrade. You can later make modifications to meet the needs of your environment.

Upgrading HYCU in an AWS GovCloud (US) environment

Prerequisites

- The AWS GovCloud (US) region that contains the HYCU backup controller must be added to HYCU. For instructions, see [“Adding an AWS GovCloud \(US\) region” on page 59](#).
- The HYCU virtual appliance image must be uploaded to an Amazon S3 bucket. To obtain the image, contact [HYCU Support](#).

Considerations

- The upgraded HYCU backup controller is renamed to `<HYCUBackupControllerName>-<NewHYCUVersion>-<Revision>` (for example, `HYCU-5.0.0-3634`) and a new IP address is assigned to it.
- After upgrading HYCU, a warning message stating that the virtual machine agent status is not ready may be displayed. You can safely ignore this message.

Procedure

1. Sign in to the AWS GovCloud (US) console.
2. Import the HYCU virtual appliance image as a snapshot to AWS GovCloud (US), and then create an Amazon Machine Image (AMI) from the snapshot:

```
aws ec2 register-image --name "<HYCUVirtualApplianceImage>"  
--block-device-mappings DeviceName="/dev/sda1",Ebs=  
{SnapshotId=<YourSnapshotID>} --root-device-name "/dev/sda1"
```

① Important Make sure that you use the following format for the HYCU virtual appliance image:

`hycu-<Version>-<Revision>`

For example: `hycu-5.0.0-363`.

If you enter the HYCU image name in a different format, you will not be able to use this image for the upgrade.

For details on how to import the HYCU virtual appliance image as a snapshot to AWS GovCloud (US) and create an AMI from the snapshot, see AWS documentation.

3. Sign in to the HYCU web user interface.
4. Click  **Administration**, and then select **Software Upgrade**.
5. In the Software Upgrade dialog box, on the Release tab, check the current version of HYCU and all available versions. You can also check whether any newer version is available on the HYCU Support portal by clicking the **Check for new version** link.
6. Select the version to which you want to upgrade HYCU.
7. Click **Software Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

① Important If upgrading HYCU fails for any reason, manually power on the original HYCU backup controller and associate it with the original IP address, and then repeat the HYCU upgrade steps from step 3.

Upgrading HYCU in Google Cloud

Prerequisites

- The Google Cloud project where the HYCU backup controller and HYCU instances reside must be added to HYCU. For instructions, see “[Adding a Google Cloud project](#)” on page 60.
- The HYCU virtual appliance image must be uploaded to a Google Cloud bucket. To obtain the image, contact [HYCU Support](#).

- Private Google Access must be enabled in the subnet where the HYCU backup controller resides.
- *Only if the HYCU backup controller uses a public IP address.* The public IP address must be static.

Considerations

- The upgraded HYCU backup controller is renamed to <HYCUBackupControllerName>-<NewHYCUVersion> (for example, HYCU-4-9-0-3634).
- An upgrade of the HYCU instances residing in Google Cloud starts automatically after the HYCU upgrade.
- A new private IP address is assigned to the upgraded HYCU backup controller.
- After you make sure that HYCU was upgraded successfully, you can safely delete the old HYCU backup controller from Google Cloud.

Procedure

1. Sign in to the Google Cloud console.
2. Create a custom image from the HYCU virtual appliance image. To do so, follow these steps:
 - a. Go to the Create an image page.
 - b. Under Name, enter a HYCU image name in the following format:
hycu-<Version>-<Revision>
| **Important** The delimiting characters in <Version> must be dashes (-). For example, hycu-4-9-0-3634.
 - c. Under Source, select **Cloud Storage file**.
 - d. In the Cloud Storage file field, click **Browse** to select the HYCU virtual appliance image that is stored in the bucket.
 - e. Under Location, select a region for your custom image.
 - f. Click **Create** at the bottom of the page.
3. Sign in to the HYCU web user interface.
4. Click  **Administration**, and then select **Software Upgrade**.
5. In the Software Upgrade dialog box, on the Release tab, check the current version of HYCU and all available versions. You can also check whether any newer version is available on the HYCU Support portal by clicking the

Check for new version link.

6. Select the version to which you want to upgrade HYCU.
7. Click **Software Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

① Important If upgrading HYCU fails for any reason, shut down the new HYCU backup controller, unassign the public IP address from the new HYCU backup controller and assign it to the original HYCU backup controller (if not using the private IP address), turn on the old HYCU backup controller, and then repeat the HYCU upgrade steps from step 3.

Upgrading HYCU in an Azure Government environment

Prerequisites

- The Azure Government subscription that contains the HYCU backup controller must be added to HYCU. For instructions, see “[Adding an Azure Government subscription](#)” on page 61.
- The HYCU virtual appliance image must be uploaded to a storage container in Azure Government. To obtain the image, contact [HYCU Support](#).

Considerations

- The upgraded HYCU backup controller is renamed to <HYCUBackupControllerName>-<NewHYCUVersion> (for example, HYCU-5.0.0-3634).
- After upgrading HYCU, a warning message stating that the virtual machine agent status is not ready may be displayed. You can safely ignore this message.

Recommendation

It is recommended that you use Azure Storage Explorer to upload the HYCU virtual appliance image to Azure Government. For details, see Azure documentation.

Procedure

1. Sign in to Azure Government.
2. Create a managed image from the HYCU virtual appliance image:

- a. In the Images navigation pane, click **Create**. In the Create an image menu that opens, make sure you specify the following:
 - In the Instance details section, in the Name field, enter the name of the HYCU virtual appliance image in the following format:
`hycu-<Version>-<Revision>`
 For example, `hycu-5.0.0-3634`.
 - In the OS disk section, select the following:
 - OS type: **Linux**
 - VM Generation: **Gen 1**

You can leave the default values for the remaining options, or adjust them to your needs.
- b. Click **Review + Create** to review the information, and then click **Create** to create the managed image.
3. Sign in to the HYCU web user interface.
4. Click  **Administration**, and then select **Software Upgrade**.
5. In the Software Upgrade dialog box, on the Release tab, check the current version of HYCU and all available versions. You can also check whether any newer version is available on the HYCU Support portal by clicking the **Check for new version** link.
6. Select the version to which you want to upgrade HYCU.
7. Click **Software Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

 **Important** If upgrading HYCU fails for any reason, manually power on the original HYCU backup controller and associate it with the original IP address, and then repeat the HYCU upgrade steps from step 3.

Applying HYCU updates

After you receive a HYCU update from HYCU Support, you can apply it to your current product version. An update can be applied only to an installed compatible product version. For example, an update labeled 1.2.3-4567 can be applied to the product version 1.2.3 whereas an update labeled 1.2.4-5678 cannot.

 **Note** Each HYCU update addresses a cumulative set of issues.

Prerequisites

- *For applying an update to a HYCU backup controller:* The HYCU backup controller activities must be suspended. For instructions on how to do this, see “[Setting power options](#)” on page 403.
- Jobs that you do not want to be aborted must be finished (the update application process aborts all currently running jobs). You can check this by filtering the Jobs list by the Executing job status. For instructions, see “[Filtering and sorting data](#)” on page 317.
- *For applying an update to a HYCU instance:* The same update must be applied to the corresponding HYCU backup controller.
- *For applying an update by using the shell script:* You must know credentials of an operating system user account that has administrative user rights on the HYCU virtual machine where you plan to apply the update.

① Important Unless instructed otherwise by HYCU Support, you must apply the same updates to all your HYCU virtual machines: HYCU backup controllers, HYCU instances, and HYCU Managers.

Considerations

- The update that you apply to the HYCU backup controller is not automatically applied to HYCU instances or HYCU Managers, if there are any in your data protection environment.
- *For applying an update to a HYCU backup controller or a HYCU Manager:* Any users that were signed in to the HYCU web user interface of the HYCU virtual machine where the update is being applied should perform a hard reload of the web user interface page in their web browser after the process completes.

Recommendation

Before applying an update to a HYCU backup controller, back up the HYCU backup controller. For instructions, see “[Backing up virtual machines](#)” on page 153.

You can apply HYCU updates:

- From the HYCU web user interface
Use this method if you want to apply an update to a HYCU backup controller, a HYCU instance, or a HYCU Manager. For instructions, see “[Applying an update by using the HYCU web user interface](#)” below.

- By using the shell script

Use this method if you are unable to sign in to the HYCU web user interface.

For instructions, see “[Applying an update by using the shell script](#)” on [page 431](#).

Applying an update by using the HYCU web user interface

From the HYCU web user interface, you can apply an update to any kind of HYCU virtual machine by using the following procedures:

- “[Applying an update to a HYCU backup controller or a HYCU Manager](#)” below
- “[Applying an update to a HYCU instance](#)” on the next page

Applying an update to a HYCU backup controller or a HYCU Manager

Procedure

1. Sign in to the HYCU web user interface.
2. Click  **Administration**, and then select **Software Upgrade**.
3. In the Software Upgrade dialog box, click the **Updates** tab.
4. In the Update label column, check if the package of the preferred update has already been added to HYCU or HYCU Manager, and then do one of the following:
 - If the update label is not present, follow these steps:
 - a. Click  **Add**.
 - b. Click **Browse**, and then browse for the update package (in the ZIP format).
 - c. Click **Add Update**.
 - d. Select the update label.
 - If the update label is present, select it.
5. Click **Apply Update**.
6. Verify that the displayed digital fingerprint matches the one that you were given by HYCU Support, and then click **Yes** to start the update process.

 **Tip** Click  **Info** to review the list of issues that the update resolves.

You are automatically signed out of the HYCU web user interface, and can track the progress of applying the update on the web user interface sign-in page.

7. When the process completes, perform a hard reload of the HYCU web user interface page in your web browser.
8. *Only if you applied an update to a HYCU backup controller.* Do the following:
 - a. Sign in to the HYCU web user interface.
 - b. Resume activities of the HYCU backup controller. For instructions on how to do this, see “[Setting power options](#)” on page 403.

You can later delete any of the update packages that you do not need anymore by selecting it and clicking  **Delete**.

Applying an update to a HYCU instance

Procedure

1. Sign in to the HYCU web user interface.
2. Click  **Administration**, and then select **HYCU Instances**.
3. In the HYCU Instances dialog box, select the HYCU instance to which you want to apply an update, and then click  **Updates**.
4. In the Update label column, check if the package of the preferred update has already been added to HYCU, and then do one of the following:
 - If the update label is not present, follow these steps:
 - a. Click  **Add**.
 - b. Click **Browse**, and then browse for the update package (in the ZIP format).
 - c. Click **Add Update**.
 - If the update label is present, select it.
5. Click **Apply Update**.
6. Verify that the displayed digital fingerprint matches the one that you were given by HYCU Support, and then click **Yes** to start the update process.

 **Note** Each update that is applied to a HYCU instance is first uploaded to the corresponding HYCU backup controller.

d. Select the update label.

• If the update label is present, select it.

 **Tip** You can click  **Info** to review the list of issues that the update resolves.

The HYCU instance status icon in the HYCU Instances dialog box turns gray to indicate the ongoing process. You can track the progress of the process by checking the status of the corresponding job in the Jobs panel. When the update is applied, the HYCU instance status icon turns green.

You can later delete any of the update packages that you do not need anymore by selecting it, and then clicking  **Delete**.

Applying an update by using the shell script

Procedure

1. Sign in to the web user interface that you are using to manage your virtualization environment, and connect to the HYCU virtual machine where you plan to apply the update.
2. Sign in to the operating system with a user account that has administrative user rights.
3. Open a command shell, and then run the following command:

```
cd /opt/grizzly/bin/
```

4. Run the following command to retrieve the list of update packages that are already added to the HYCU virtual machine:

```
sudo ./HycuPatch.sh -list_patches
```

5. If the label of the preferred update is not present on the list, follow these steps:
 - a. Extract the contents of the update package (in the ZIP format). The package contains the main update file, installation instructions, and digital fingerprints.
 - b. Use the /usr/bin/cksum and /usr/bin/md5sum commands to verify that the digital fingerprint of the main update file matches the one that you were given by HYCU Support.
 - c. Copy the main update file in the archived TAR (.tar.gz) format to the following directory on the HYCU virtual machine:

```
/hycudata/opt/grizzly/updates
```

 **Tip** Run the following command to review the list of issues that the update resolves:

```
sudo ./HycuPatch.sh -patch_info <UpdateLabel>
```

- Run the following command to apply the update to the HYCU virtual machine:

```
sudo ./HycuPatch.sh -apply_patch <UpdateLabel>
```

- Only if you applied an update to a HYCU backup controller. Do the following:
 - Sign in to the HYCU web user interface.
 - Resume activities of the HYCU backup controller. For instructions on how to do this, see [“Setting power options” on page 403](#).

Configuring SSL certificates

To establish trusted and secure communication in your data protection environment, you must configure SSL certificates.

Accessing the SSL Certificates dialog box

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

In the SSL Certificates dialog box that opens, you can view the information about your SSL certificate—the certificate name, the certificate common name, the certificate expiry date, the certificate key type, and the generated certificate signing request (CSR).

Consideration

After you create or import an SSL certificate, make sure to update also the HYCU network settings by specifying this certificate. For details on how to do this, see [“Configuring your network” on page 392](#).

Recommendation

It is recommended to replace the self-signed certificate that is generated automatically during HYCU deployment with a CA-signed certificate.

Depending on what you want to do, see one of the following:

I want to...	Procedure
Create a self-signed certificate.	“Creating a self-signed certificate”

I want to...	Procedure
	below
Create a certificate signing request.	“Creating a certificate signing request” on the next page
Import a custom certificate to HYCU.	“Importing a custom certificate” on page 435

Creating a self-signed certificate

Procedure

1. In the SSL Certificates dialog box, click **⊕ Generate**.
2. Select **Generate self-signed certificate**, and then click **Next**.
3. Provide the following certificate-related information:
 - Name
 - Common name
 - Organization
 - Organization unit
 - Location
 - Country
 - *Optional.* State
 - Key algorithm
 - Key size
4. Click **Generate**.

The self-signed certificate is added to the list of SSL certificates. Keep in mind that each SSL certificate that is generated through HYCU is valid for 824 days and that you must maintain the validity of the certificate.

Creating a certificate signing request

Procedure

1. In the SSL Certificates dialog box, click **⊕ Generate**.
2. Select **Generate certificate signing request**, and then click **Next**.

3. Provide the following certificate-related information:

- Name
- Common name
- Organization
- Organization unit
- Location
- Country
- *Optional.* State
- Key algorithm
- Key size
- *Optional.* Subject alternative name (SAN)

The SAN is a list of possible names of the HYCU backup controller. Each name can be one of the following:

- FQDN (for example, `hycu-bc.hycu.local`)
- Wildcard prefix "`*`." followed by a domain name (for example, `*.hycu.local`)
- IPv4 address (for example, `10.1.100.1`)
- IPv6 address (for example, `fe80::1234:5678:9abc:def0`)

! Important The maximum number of characters in each field is 64.

4. Click **Generate**.

An SSL private key is added to the list of SSL certificates and the  icon in the CSR column indicates that the certificate signing request has been created.

Click  **Download Generated CSR** to download the generated CSR.

After you create and download the CSR, you can send it to a certificate authority to create a certificate. The certificate that is created from the CSR by the certificate authority must be added to the SSL private key to complete the SSL key pair in HYCU. For instructions, see “[Completing the SSL key pair](#)” on the [next page](#).

Completing the SSL key pair

Procedure

1. In the SSL Certificates dialog box, select the SSL private key, and then click  **Edit**.
2. Browse for the following files:

- Certificate: The file with the certificate that was created from the CSR.
- *Optional.* CA certificate/chain: The file with the CA-signed certificate or trust chain certificates.

3. Click **Update**.

Importing a custom certificate

HYCU enables you to import an SSL key pair, or a CA-signed certificate or trust chain certificates.

Prerequisites

- *For importing an SSL key pair:* The private key and the certificate must be available.
- *For importing an SSL key pair from PEM files:*
 - All certificate files must be unencrypted.
 - The certificate must be compliant with the PKCS#7 standard and encoded in the PEM format.
- *For importing a CA-signed certificate or trust chain certificates from a file:* The CA-signed certificate or trust chain certificates must be available.

Considerations

- If the certificate uses a wildcard for the Common Name (CN), make sure that the Certificate Subject Alt Name field includes all possible host names or FQDNs, and their corresponding IP addresses. Otherwise, the certificate may be recognized as invalid by your web browser or hyCLI.
- If you are importing an SSL key pair from a PFX file, consider the following:
 - The PFX file must contain the entire trust chain to the root CA certificate.
 - The PFX must contain a single private key along with its associated certificate. If multiple keys exist in the PFX file, the import may fail or only one private key may be imported.

Procedure

1. In the SSL Certificates dialog box, click  **Import**.
2. Depending on whether you want to import an SSL key pair, a CA-signed certificate or trust chain certificates from a file, or a CA-signed certificate or trust chain certificates from a host, select one of the following options, and

then click **Next** and follow the instructions:

Option	Instructions
Import SSL keypair from PEM files	<p>a. Enter a name for your certificate.</p> <p>b. Browse for the following files:</p> <ul style="list-style-type: none"> • <i>Optional.</i> CA certificate/chain: The file with the CA-signed certificate or trust chain certificates. • Certificate: The file with the certificate corresponding to the private key that you are importing. • Private key: The file with the private key that is associated with the certificate that you are importing. <p>The private key should be created with the RSA or ECDSA algorithm in the PKCS#1 or PKCS#8 format. The minimum key size for private keys created with the RSA algorithm is 2048 bits.</p> <p> Note If you use Conjur for managing your HYCU secrets, you can enable the Retrieve values from secrets manager switch if you want to provide the secret instead of browsing for the file. For details on managing secrets, see “Managing secrets” on page 404.</p>
Import SSL keypair from PFX file	<p>a. Enter a name for your certificate.</p> <p>b. Browse for the PFX file that contains the required SSL key pair.</p> <p>c. <i>Optional.</i> Enter the passphrase of the PFX file.</p>
Import CA certificate/chain from file	<p>a. Enter a name for your certificate.</p> <p>b. Browse for the file with the CA-signed certificate or trust chain certificates.</p>
Import CA certificate/chain	<p>a. Enter the server host name or IP address, and the port.</p>

Option	Instructions
from host	<p>The following examples show which host names and ports to use in common configuration scenarios:</p> <ul style="list-style-type: none"> • If HYCU is configured to use HTTPS for WinRM connections to virtual machines, enter the host name or IP address of the virtual machine for which you want to establish an HTTPS connection, and the HTTPS port (usually 5986). • If HYCU is configured to use LDAP over SSL (LDAPS), enter the LDAPS server host name or IP address, and the LDAPS port (usually 636). • If you are using STARTTLS or SSL/TLS for SMTP connections, enter the SMTP server host name or IP address, and the port for authenticated SMTP connections (465 for the SSL/TLS security mode, and 587 or 25 for the STARTTLS security mode). <p>b. From the Security mode drop-down menu, select the preferred security mode.</p> <p>c. Click Retrieve.</p> <p>d. Review the certificates and select the one that you want to import. If you select the CA-signed certificate for the import, its trust chain certificates will be trusted as well.</p>

3. Click **Import**.

You can also change the name of any self-signed or custom certificate (click  **Edit** and make the required modification) or delete the ones that you do not need anymore (click  **Delete**).

Sharing telemetry data with HYCU

You can configure HYCU to collect telemetry data. This data helps HYCU to provide proactive support and improved performance to better meet your data protection environment needs.

Sharing diagnostic data through telemetry enables proactive, contextualized support for HYCU as follows:

1. Collects detailed data on your data protection environment that includes the syslog files, HYCU internal data base (PostgreSQL) logs, system activity information (sar), data distribution statistics for file share backups, HYCU license information, and other detailed information on your specific infrastructure, and then sends this data to HYCU Support.

 **Important** HYCU does not collect any sensitive information from your data protection environment.

2. Analyzes collected data, generates internal reports, and identifies eventual problems or unfavorable trends considerably reducing issue resolution time.
3. Provides you with feedback on your HYCU environment that addresses eventual issues and instructs you on how to adjust your environment and to improve infrastructure and performance.

 **Note** You need to enable telemetry data sharing for each HYCU backup controller that you want to include in the advanced troubleshooting.

Prerequisite

You must have a valid HYCU Support user account.

Consideration

When a Managed Service Provider (MSP) license is applied to HYCU, sharing telemetry data with HYCU is enabled by default and cannot be disabled.

Accessing the Telemetry dialog box

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

Procedure

In the Telemetry dialog box, use the **Share telemetry data with HYCU, Inc.** switch to allow HYCU to collect your telemetry data, and then click **Save**.

HYCU starts collecting data and sends it to HYCU Support. Later, the telemetry diagnostic data is sent to HYCU Support once a day. You can view the collection job status in the Jobs panel.

If you later decide that you no longer want to share your telemetry data with HYCU, disable the **Share telemetry data with HYCU, Inc.** option for each configured HYCU backup controller.

 **Note** When the **Share telemetry data with HYCU, Inc.** option is enabled, you can send the log files to HYCU Support. For more information, see “Setting up logging” on page 389.

Removing HYCU

When you remove HYCU from your environment, you also need to perform additional cleanup tasks.

To remove HYCU, follow these steps:

1. Sign in to HYCU, and then unassign policies from all entities as follows:
 - To unassign policies from virtual machines:
 - a. In the navigation pane, click  **Virtual Machines**.
 - b. Select all virtual machines, and then click  **Set Policy**.
 - c. Click **Unassign**.
 - d. Click **Yes** to confirm that you want to unassign the policies from the selected virtual machines.
 - To unassign policies from applications:
 - a. In the navigation pane, click  **Applications**.
 - b. Select all discovered applications, and then click  **Set Policy**.
 - c. Click **Unassign**.
 - d. Click **Yes** to confirm that you want to unassign the policies from the selected applications.
 - To unassign policies from file shares:
 - a. In the navigation pane, click  **Shares**.
 - b. Select all file shares, and then click  **Set Policy**.
 - c. Click **Unassign**.

- d. Click **Yes** to confirm that you want to unassign the policies from the selected file shares.
- To unassign policies from volume groups:
 - a. In the navigation pane, click  **Volume Groups**.
 - b. Select all volume groups, and then click  **Set Policy**.
 - c. Click **Unassign**.
 - d. Click **Yes** to confirm that you want to unassign the policies from the selected volume groups.

2. *Only if HYCU was used for file share protection.* Do the following:

- a. Remove the existing HYCU instances. For instructions, see “[Deleting a HYCU instance](#)” on page 384.
- b. Remove the file server snapshots created by HYCU. To do so, on the HYCU backup controller, run the `/opt/grizzly/bin/HycuCleanup.pl` script as follows:

```
sudo perl HycuCleanup.pl -c <FileServer> -u <Username> -p
<Password> -dnfs -all
```

In this instance, `<FileServer>` is the name of the file server in the following format: `https://<ServerName>:<Port>`.

 **Important** By running this command, you will also remove all file server snapshots whose names start with `hycu-` (case insensitive).

3. *For Nutanix clusters:* On the HYCU backup controller, run the `/opt/grizzly/bin/HycuCleanup.pl` script as follows:

- To remove virtual machine and volume group snapshots created by HYCU:

```
sudo perl HycuCleanup.pl -c <NutanixCluster> -u <Username> -
p <Password> -dvms -all
```

```
sudo perl HycuCleanup.pl -c <NutanixCluster> -u <Username> -
p <Password> -dvgs -all
```

In these instances, `<NutanixCluster>` is the name of the Nutanix cluster in the following format: `https://<ServerName>:<Port>`.

 **Important** By running these commands, you will also remove all third-party snapshots created by using Nutanix REST API v3 whose names start with the IP address.

- To remove volume groups created by HYCU:

```
sudo perl HycuCleanup.pl -c <NutanixCluster> -u <Username> -  
p <Password> -dvg -all
```

In this instance, *<NutanixCluster>* is the name of the Nutanix cluster in the following format: <https://<ServerName>:<Port>>.

⚠ Important By running this command, you will also remove all volume groups created by using Nutanix REST API v3 whose names start with HYCU- (case insensitive).

4. Remove data from targets. To do so, on each target, delete the `bkpctrl` folder.
5. Sign in to the Nutanix Prism web console, the vSphere (Web) Client, the AWS GovCloud (US) console, the Azure Government portal, or the Google Cloud console, and then delete the HYCU backup controller virtual machine. For instructions, see the relevant documentation.

Chapter 12

Tuning your data protection environment

Administration tasks that you perform through the ☰ **Administration** menu to customize HYCU for your data protection environment are usually sufficient to successfully manage it. However, sometimes the needs of your organization require additional administration tasks to be performed for optimal performance, a higher security level, or interaction with external applications, as well as for taking advantage of a broader spectrum of HYCU options.

I want to...	Procedure
Access the HYCU backup controller virtual machine by using SSH.	“Accessing the HYCU backup controller virtual machine by using SSH” on the next page
Enable HTTPS for WinRM connections.	“Enabling HTTPS for WinRM connections” on page 446
Configure FIPS-compliant mode for HYCU.	“Configuring FIPS mode for HYCU” on page 446
Set up LDAPS authentication.	“Setting up LDAPS authentication” on page 448
Set up two-factor authentication.	“Setting up two-factor authentication” on page 448
Manage API keys.	“Managing API keys” on page 449
Manage FIDO authenticators.	“Managing FIDO authenticators” on page 450
Secure SMTP connections.	“Securing SMTP connections” on page 451
Set up HYCU to use multiple	“Setting up HYCU to use multiple

I want to...	Procedure
networks.	“networks” on page 452
Increase the size of the HYCU virtual disks.	“Increasing the size of the HYCU virtual disks” on page 455
Assign required privileges to a vSphere user.	“Assigning privileges to a vSphere user” on page 456
Use the HYCU REST API to automate tasks.	“Using the HYCU REST API Explorer” on page 460
Use hyCLI.	“Using the command-line interface” on page 460
Use the pre and post scripts to perform necessary actions before and after the backup and the restore are performed.	“Using the pre and post scripts” on page 461

Accessing the HYCU backup controller virtual machine by using SSH

You can perform most administrative tasks of the HYCU backup controller by using the HYCU web user interface or command-line user interface (hyCLI). The only two exceptions for which you should use SSH are restarting the HYCU application server (the Grizzly server) or the entire appliance.

① Important Using SSH to perform any tasks other than restarting the HYCU application server or the entire appliance is not recommended.

After you deploy the HYCU virtual appliance, you can use the following default credentials to access the HYCU backup controller virtual machine by using SSH:

User name: **hycu**

Password: **hycu/4u**

Changing the default SSH password

For security purposes, it is highly recommended that you change the default SSH password. To do so, follow these steps:

1. Open a remote session to the HYCU backup controller virtual machine:

```
ssh hycu@<HYCUBackupControllerIPAddress>
```

When requested, enter the default password.

2. Change the password for the hycu user:

```
passwd
```

When requested, enter the default password again, and then enter and verify your new password.

Configuring SSH public key authentication

Adding an SSH public key to HYCU and using it to access the HYCU backup controller enables you to add an additional layer of security to your data protection environment by providing a more secure alternative to SSH password authentication. If you are using HYCU for file share protection and you configure SSH public key authentication for accessing the HYCU backup controller, you can use the same SSH public key also to access your HYCU instances. For added security, you can choose to disable SSH password authentication.

Limitation

The supported SSH key types are RSA, ECDSA, and Ed25519.

Accessing the SSH Authentication dialog box

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

Procedure

1. In the SSH Authentication dialog box, click  **Add Public Key**.
2. Enter a name for the SSH public key and the SSH public key.
3. Click **Save**.

The SSH public key is added to HYCU. For each added key, the name, creation date, and key fingerprint are displayed.

You can also delete any of the existing SSH public keys by selecting the key, and then clicking  **Delete**.

If after configuring SSH public key authentication you want to disable SSH password authentication, you can do so by disabling the **Allow password authentication** switch, and then clicking **Save**.

Disabling SSH access

You can disable SSH access at any time. To do so, follow these steps:

1. Open a remote session to the HYCU backup controller virtual machine:

```
ssh hycu@<HYCUBackupControllerIPAddress>
```

When requested, enter the password for the hycu user.

2. Shut down the SSH service:

```
sudo systemctl stop sshd.service
```

When requested, enter the password for the hycu user.

3. Disable the SSH service:

```
sudo systemctl disable sshd.service
```

If requested, enter the password for the hycu user.

After performing this procedure, your SSH connection will be disabled. To re-enable SSH, you need to connect to the HYCU backup controller virtual machine through the console of the respective source.

Managing the HYCU application server

To manage the HYCU application server, follow these steps:

1. Open a remote session to the HYCU backup controller virtual machine:

```
ssh hycu@<HYCUBackupControllerIPAddress>
```

When requested, enter the password for the hycu user.

2. Perform the preferred operation on the HYCU application server:

```
sudo service grizzly {start | stop | restart}
```

When requested, enter the password for the hycu user.

① Important If you plan to restart the PostgreSQL server, make sure the HYCU application server is stopped before and started after restarting the PostgreSQL server.

Enabling HTTPS for WinRM connections

If you want to add an additional layer of security, you can configure HYCU to use HTTPS for WinRM connections to virtual machines.

Procedure

For each virtual machine for which you want to enable HTTPS for WinRM connections, do the following:

1. Configure WinRM for HTTPS. For details on how to do this, see Microsoft documentation.
2. *Only if WinRM is configured with a certificate that was signed by a private certificate authority or with a self-signed certificate.* Import the CA-signed certificate or trust chain certificates to HYCU. For instructions, see “Importing a custom certificate” on page 435.

Configuring FIPS mode for HYCU

HYCU can be configured to operate to be compliant with the Federal Information Processing Standards (FIPS) 140-2 that establish security requirements for cryptography modules (which encryption algorithms and methods for generating encryption keys can be used).

Depending on the nature of your business, you can either enable or disable FIPS mode for HYCU. To check whether FIPS mode is enabled (disabled by default), open a remote session to the HYCU backup controller, and then as the root user or by using sudo, run the following command:

```
/opt/grizzly/bin/enable_fips.sh --status
```

Limitations

When FIPS mode is enabled, the following limitations apply:

- SMB targets cannot be used for storing data.
- Applications cannot be discovered and therefore protected.
- Individual files cannot be restored.
- Windows servers cannot be protected.
- SMB file shares on Nutanix Files servers cannot be protected.

Considerations

- *Only if HYCU is used for file share protection.* You must enable FIPS mode for each HYCU instance separately (independent of the HYCU backup controller).
- After you upgrade HYCU, FIPS mode will be disabled. If required, make sure to re-enable it.

Enabling FIPS mode for HYCU

Procedure

Open a remote session to the HYCU backup controller, and then as the root user or by using sudo, do the following:

1. Stop the HYCU application server:

```
systemctl stop grizzly.service
```

2. Enable FIPS-compliant mode:

```
/opt/grizzly/bin/enable_fips.sh
```

3. Reboot the HYCU backup controller:

```
reboot
```

Disabling FIPS mode for HYCU

Procedure

Open a remote session to the HYCU backup controller, and then as the root user or by using sudo, do the following:

1. Stop the HYCU application server:

```
systemctl stop grizzly.service
```

2. Disable FIPS-compliant mode:

```
/opt/grizzly/bin/enable_fips.sh -d
```

3. Reboot the HYCU backup controller:

```
reboot
```

Setting up LDAPS authentication

If you want to add an extra layer of protection and ensure the confidentiality of data, you can configure HYCU to use LDAP over SSL (LDAPS) for secure user authentication. For this authentication to work, HYCU must trust the LDAPS server certificate. Depending on the type of the LDAPS server certificate, trust is established in one of the following ways:

- If the LDAPS server certificate was signed by a public certificate authority, HYCU will trust it automatically.
- If the LDAPS server certificate was signed by a private certificate authority, or if it is self-signed, you must import the CA-signed certificate or trust chain certificates to HYCU. For instructions, see [“Importing a custom certificate” on page 435](#).

Setting up two-factor authentication

You can set up two-factor authentication to add an extra layer of security when signing in to HYCU. The following authentication methods are supported:

- Time-based one-time passwords (OTP) generated by an OTP application.
- Authenticators compliant with the FIDO protocol (FIDO authenticators), such as security keys and fingerprint reader.

When setting up two-factor authentication for HYCU, you must complete the following tasks:

Task	Instructions
1. Perform the necessary preparation steps for the selected authentication method.	<ul style="list-style-type: none"> • <i>For OTP:</i> Provide instructions to users and make sure that they have access to

Task	Instructions
	<p>an OTP application.</p> <ul style="list-style-type: none"> • <i>For FIDO authenticators:</i> Make sure that the following conditions are met: <ul style="list-style-type: none"> ◦ Authenticators are set up correctly. For instructions, see the authenticator documentation. ◦ The DNS is configured properly. ◦ The host name is resolved properly.
2. Create or edit a user for whom you want to enable two-factor authentication, and then add this user to a user group.	Follow the procedures described in “ Creating a user ” on page 348 and “ Adding a user to a group ” on page 352 .

After you perform these tasks, users must authenticate their sign-ins by signing in to HYCU as described in “[Signing in to HYCU](#)” on page 49.

Managing API keys

API keys are needed in the following scenarios:

- If you enable two-factor authentication for using the REST API or the HYCU command-line user interface (hyCLI).
- If you enable API key authentication when adding a Hybrid Cloud Edition controller to HYCU Manager.

You can generate or revoke your API keys by using the API keys option.

Consideration

As a user with the Administrator role assigned, you can edit other users' information through the Self-Service panel. For details, see “[Creating a user](#)” on [page 348](#).

Accessing the API Keys dialog box

To access the API Keys dialog box, click  at the upper right of the screen, and then select **API Keys**.

Generating an API key

Procedure

1. In the API Keys dialog box, click  **New**.
2. Enter a name for the API key.
3. *Optional.* Set the expiry date. If you do not set the expiry date, the API key does not expire.
4. Click **Generate**. The API key is displayed.
5. Write the API key down and store it safely.

 **Important** For security reasons, the API key will never be displayed again, so make sure to write it down and keep it safe. You can copy the API key to the clipboard by clicking  **Copy to Clipboard**.

Your API key can be used to access your data, therefore, treat it like a password.

6. Click **Finish**.

Revoking an API key

Procedure

1. In the API Keys dialog box, select the API key that you want to revoke, and then click  **Revoke**.
2. Click **Revoke** to confirm that you want to revoke the API key. The API key is immediately revoked.

Managing FIDO authenticators

If the FIDO two-factor authentication method is enabled for your account, you need to set up a FIDO authenticator. You can add or revoke your FIDO authenticators by using the FIDO Authenticators option.

Adding a new FIDO authenticator

Considerations

- As a user in the Infrastructure group with the Administrator role assigned, you can edit other users' information through the Self-Service panel. For details, see “[Creating a user](#)” on page 348.
- Make sure that you use a fully qualified domain name when signing in to HYCU and that DNS is correctly configured. Otherwise, authentication may fail.

Accessing the FIDO Authenticators dialog box

To access the FIDO Authenticators dialog box, click  at the upper right of the screen, and then select **FIDO Authenticators**.

Procedure

- In the FIDO Authenticators dialog box, click  **New**. The Security Setup wizard opens.
- Follow the wizard instructions to create the FIDO authenticator. The process depends on the type of authenticator you select and the operating system version.
- In the Name field, enter a name for the FIDO authenticator.
- Click **Register**.

Revoking a FIDO authenticator

Procedure

- In the FIDO Authenticators dialog box, select the authenticator that you want revoke, and then click  **Revoke**.
- Click **Revoke** to confirm that you want to revoke the FIDO authenticator. The FIDO authenticator is immediately revoked.

Securing SMTP connections

If you want to add an extra layer of protection and ensure the confidentiality of data, you can configure HYCU to use SMTP over SSL/TLS or STARTTLS for secure user authentication. For this authentication to work, HYCU must trust

the SMTP server certificate. Depending on the type of the SMTP server certificate, trust is established in one of the following ways:

- If the SMTP server certificate was signed by a public certificate authority, HYCU will trust it automatically.
- If the SMTP server certificate was signed by a private certificate authority, or if it is self-signed, you must import the CA-signed certificate or trust chain certificates to HYCU. For instructions, see “[Importing a custom certificate](#)” on page 435.

Setting up HYCU to use multiple networks

You can set up HYCU to operate in a multi-network environment, allowing it to have two network adapters assigned to different VLANs or network segments. This is especially useful if you have dedicated storage used for backups in a different network than HYCU. For example:

- HYCU could be located on the 10.0.0.0/16 VLAN and a storage box could be located on the 192.168.0.0/24 VLAN.
- You need to access the HYCU web user interface from a network other than the virtual machine network. In this case, it is recommended to have a dedicated NIC for data transfer that must be on the same VLAN as the Nutanix Controller virtual machines, in addition to the NIC for the web user access.

 **Note** *For Nutanix clusters:* While the bulk of data traffic during a backup takes place over the additional network, part of it is still done through the management network. This is because HYCU uses the Nutanix data services IP address to consume data through Nutanix Volumes, which must be in the same subnet as the management network of the CVMs.

For details on this limitation, see Nutanix documentation.

Limitation

You cannot set up HYCU to use multiple networks in AWS GovCloud (US), Google Cloud, and Azure Government environments.

File server environment considerations

- The main network must correspond to a network segment where both the HYCU backup controller and the additional HYCU instances can see and establish a connection to each other.
- Both virtual machines (the HYCU backup controller and one or more connected HYCU instances) must be able to connect to the file server.
- Each network adapter must be on a different subnet.
- *Only if the DNS servers are specified.* The DNS servers on all subnets must return the same results.
- *For Nutanix ESXi clusters:* When upgrading HYCU, network settings on all additional network adapters will be set to the default values. Make sure to reconfigure the HYCU instance after the upgrade.

Depending on the environment in which you want to set up HYCU to use multiple networks, perform one of the following procedures:

- “Setting up HYCU to use multiple networks on a Nutanix AHV or Nutanix ESXi cluster” below
- “Setting up HYCU to use multiple networks in a vSphere environment” on the next page

Setting up HYCU to use multiple networks on a Nutanix AHV or Nutanix ESXi cluster

Procedure

1. Sign in to the Nutanix Prism web console, and then add an additional network adapter:
 - a. In the menu bar, click **Home**, and then select **VM**.
 - b. Click the **Table** tab to display the VM Table view, and then, from the list of virtual machines, select your HYCU virtual machine.
 - c. Click **Update**, and then navigate to the Network Adapters (NIC) section.
 - d. Click **Add New NIC**, and then select the required VLAN and click **Add**.
 - e. Click **Save**.

For details, see Nutanix documentation.

2. Configure the network. To do so, depending on how the VLAN is set up, select one of the following approaches:

- VLAN has IP address (DHCP) management enabled

Assign the IP address directly from the Nutanix Prism web console.

- VLAN does not have IP address (DHCP) management enabled

Depending on whether you are configuring an additional network for a HYCU instance, a HYCU backup controller, or a HYCU Manager, do one of the following:

- *For a HYCU instance:* Configure the network manually:
 - Open a remote session to the HYCU backup controller virtual machine:

```
ssh hycu@<HYCUBackupControllerIPAddress>
```

 - Open the `ifcfg-mainnetwork.template` file located at `/opt/grizzly/misc/`, and then follow the instructions provided in this template. Make sure to run the specified commands as the root user or by using sudo.
 - For a HYCU backup controller or a HYCU Manager:* Configure the network from the HYCU backup controller or HYCU Manager user interface. For details, see “[Configuring your network](#)” on page 392.

After the new network adapter is properly configured, you can add a target located on another VLAN to HYCU.

Setting up HYCU to use multiple networks in a vSphere environment

① Important You can use either the vSphere Web Client or the vSphere Client as the interface for performing the procedure described in this section. As an example, you are guided through the steps that you must perform if you are using the vSphere Web Client.

Procedure

1. Sign in to the vSphere Web Client, and then add an additional network adapter:
 - a. Click the **VMs** tab, and then navigate to your HYCU backup controller.
 - b. Right-click the HYCU backup controller, and then select **Edit Settings**.

- c. From the New device drop-down menu, select **Network**, and then click **Add**.
- d. From the New Network drop-down menu, select the required network.

! Important Make sure not to select a vSphere distributed switch (dvSwitch) for the virtual NIC option.

- e. Click **OK**.

For details, see VMware documentation.

2. Depending on whether you are configuring an additional network for a HYCU instance, a HYCU backup controller, or a HYCU Manager, do one of the following:
 - *For a HYCU instance:* Configure the network manually:
 - a. Open a remote session to the HYCU backup controller virtual machine:

```
ssh hycu@<HYCUBackupControllerIPAddress>
```

 - b. Open the `ifcfg-mainnetwork.template` file located at `/opt/grizzly/misc/`, and then follow the instructions provided in this template. Make sure to run the specified commands as the root user or by using sudo.
- *For a HYCU backup controller or a HYCU Manager:* Configure the network from the HYCU backup controller or HYCU Manager user interface. For details, see “[Configuring your network](#)” on page 392.

After the new network adapter is properly configured, you can add a target located on another network to HYCU.

Increasing the size of the HYCU virtual disks

If you are running out of disk space on your HYCU backup controller, you can increase the size of the HYCU virtual disks as needed.

Procedure

1. Depending on the environment in which the HYCU backup controller resides, sign in to the Nutanix Prism web console, the vSphere Web Client

or the vSphere Client, the AWS GovCloud (US) console, the Google Cloud console, or the Azure Government portal.

2. Shut down the HYCU backup controller.
3. Increase the size of the HYCU disk as required.
4. Turn on the HYCU backup controller.

For instructions on how to perform these steps, see Nutanix, VMware, AWS, Google Cloud, or Azure documentation.

Assigning privileges to a vSphere user

You can assign required privileges to a vSphere user by using the vSphere (Web) Client.

① Important You can use either the vSphere Web Client or the vSphere Client as the interface for performing the procedure described in this section. As an example, you are guided through the steps that you must perform if you are using the vSphere Web Client.

Procedure

1. Sign in to the vSphere Web Client as an administrator.
2. Click **Administration > Roles**.
3. Add a new role, and then type its name (for example, **HYCU**).
4. Depending on your data protection environment, select the required privileges for the role:

Nutanix ESXi cluster

Privilege category	Backup privileges	Restore privileges	Upgrade and HYCU instance creation privileges
Cryptographic operations	Direct Access	Direct Access	Not applicable
Datastore	Browse datastore	Browse datastore	• Allocate space

Privilege category	Backup privileges	Restore privileges	Upgrade and HYCU instance creation privileges
			<ul style="list-style-type: none"> • Low level file operations
Network	Not applicable	Not applicable	Assign network
vApp	Not applicable	Not applicable	Import
Virtual Machine > Configuration	Not applicable	Not applicable	<ul style="list-style-type: none"> • Add existing disk • Add new disk • Change Settings • Remove disk
Virtual Machine > Interaction	Not applicable	Not applicable	Power On
Virtual Machine > Inventory	Not applicable	Not applicable	<ul style="list-style-type: none"> • Create from existing • Remove
Virtual Machine > Provisioning	Not applicable	Not applicable	Clone virtual machine
vSphere Tagging	Assign or Unassign vSphere Tag	Assign or Unassign vSphere Tag	Not applicable

vSphere environment

Privilege category	Backup privileges	Restore privileges	Upgrade privileges
Cryptographic operations	Direct Access	Not applicable	Not applicable

Privilege category	Backup privileges	Restore privileges	Upgrade privileges
Datastore	<ul style="list-style-type: none"> • Browse datastore • Low level file operations 	<ul style="list-style-type: none"> • Allocate space • Low level file operations 	<ul style="list-style-type: none"> • Allocate space • Low level file operations
Global	<ul style="list-style-type: none"> • Disable methods • Enable methods 	Not applicable	Not applicable
Host > Local operations	Not applicable	<ul style="list-style-type: none"> • Create virtual machine • Delete virtual machine • Reconfigure virtual machine 	Not applicable
Network	Not applicable	<ul style="list-style-type: none"> • Assign network • Configure 	Assign network
Resource	Not applicable	Assign virtual machine to resource pool	Not applicable
vApp	Not applicable	Add virtual machine	Import
Virtual Machine > Configuration	<ul style="list-style-type: none"> • Toggle disk change tracking • Change Settings 	All privileges	<ul style="list-style-type: none"> • Add existing disk • Add new disk • Add or remove device • Change Settings • Remove disk • Rename
Virtual	Power On	• Answer	Power On

Privilege category	Backup privileges	Restore privileges	Upgrade privileges
Machine > Interaction		question • Connect devices • Power Off • Power On	
Virtual Machine > Inventory	Not applicable	• Create new • Register • Remove • Unregister	• Create from existing • Remove
Virtual Machine > Provisioning	• Allow read-only disk access • Allow virtual machine download • <i>For backing up a template:</i> Mark as template • <i>For backing up a template:</i> Mark as virtual machine	Allow disk access	Clone virtual machine
Virtual Machine > Snapshot management	• Create snapshot • Remove snapshot	Revert to snapshot	Not applicable
vSphere Tagging	Assign or Unassign vSphere Tag	Assign or Unassign vSphere Tag	Not applicable

5. Assign the created role to the vSphere user.

For details, see VMware documentation.

Using the HYCU REST API Explorer

HYCU provides a REST API that can be used by external applications to interact with the HYCU backup controller, retrieve information from it, and automate tasks. All functionality exposed through the HYCU user interface is also available through the HYCU REST API. You can use the HYCU REST API Explorer to interact with the API and view the expected input and output formats for each endpoint.

To access the HYCU REST API Explorer, follow these steps:

1. Click **?** **Help** at the upper right of the screen, and then select **REST API Explorer**. The HYCU REST API Explorer opens.
2. In the list of functionality groups, you can expand the preferred group by clicking **List Operations**. A list of API endpoints is displayed.
3. Click any of the endpoints to show the description, the parameters, and the output format. You can fill in the fields, and then click **Try it out!** to call an API and get output data.

Using the command-line interface

You can manage your data protection environment also by using the HYCU command-line user interface (hyCLI). hyCLI provides the functionality comparable to the HYCU web user interface and enables you to implement scripts for automating certain tasks.

To enable the usage of hyCLI, follow these steps:

1. Download the `hycli.zip` package. To do so, click **?** **Help** at the upper right of the screen, and then select **Download hyCLI Beta**.
2. Save and extract the `hycli.zip` file to any location on your system.
3. Add the folder containing the extracted files to the PATH environment variable.
4. *Only if two-factor authentication is enabled for your account.* Generate an API key. You will need to provide this key each time you run a hyCLI command. For details, see [“Managing API keys” on page 449](#).

 **Note** hyCLI log files are located at `.Hycu/log` in the user's home folder. You can change logging settings for hyCLI in the `logging.properties` files located in the folder containing the extracted files.

For detailed information about hyCLI, see the `README.txt` file that you can find in the folder containing the extracted files.

For more information on the hyCLI structure, commands, and usage, run the `hycli help` command.

Using the pre and post scripts

If you want to use the pre/post scripts to perform necessary actions before and after the backup and the restore are performed, these scripts should return an exit code of 0 for success and any other value for failure. In the latter case, the data protection operation is also affected as follows:

- An exit code is greater than 0: The status of the job (and the backup in the case of the backup operation) will be set to Warning and the job will continue.
- An exit code is less than 0: The status of the job (and the backup in the case of the backup operation) will be set to Failed.

During the execution of the scripts, the following environment variables are exported:

Environment variable	Description
<code>HYCU_BKPCTRL_URL</code>	HYCU backup controller URL
<code>HYCU_BKPCTRL_UUID</code>	HYCU backup controller UUID
<code>HYCU_VM_UUID</code>	Virtual machine UUID
<code>HYCU_BACKUP_UUID</code>	Backup UUID
<code>HYCU_JOB_UUID</code>	Job UUID
<code>HYCU_TARGET_UUID</code>	Target UUID
<code>HYCU_VM_NAME</code>	Virtual machine name ^a
<code>HYCU_TARGET_NAME</code>	Target name ^a
<code>HYCU_TARGET_PATH</code>	Path to the data on the target
<code>HYCU_SUCCESS</code>	<i>Available only for post scripts.</i> Success of the data protection operation.
<code>HYCU_PREEXEC_RETURN_CODE</code>	<i>Available only for post scripts.</i> Exit code

Environment variable	Description
	of the pre script.

^a If the name contains the space character or any of the following characters: " ' , ; & % € () < > { } | ^ ` ^ , these characters are replaced with an underscore before the export.

For details on how to specify pre and post scripts, see the following sections:

- “[Specifying pre/post-backup and pre/post-snapshot scripts](#)” on page 150
- “[Restoring individual files](#)” on page 202

Chapter 13

Monitoring data protection environments

HYCU Manager is designed to provide you with the visibility you need to proactively monitor all your data protection environments, allowing you to view their overall status from a single console. With HYCU Manager, data protection information received from all registered HYCU controllers is consolidated in one place and you can easily access this information for the on-premises (HYCU) and the following cloud data protection environments:

- HYCU Data Protection as a Service for AWS (HYCU for AWS)
- HYCU Data Protection as a Service for Google Cloud (HYCU for Google Cloud)
- HYCU Data Protection as a Service for Azure (HYCU for Azure)
- HYCU for Microsoft 365 and Google Workspace

Before you can start monitoring your data protection environments, you must complete the following tasks:

Task	Instructions
1. Deploy the HYCU virtual appliance in the HYCU Manager mode.	“Deploying the HYCU virtual appliance” on page 20
2. Add a HYCU controller to HYCU Manager.	“Adding a HYCU controller” on the next page

Managing HYCU controllers

You can use the HYCU Controllers panel to add, edit, and delete the HYCU controllers, as well as to view the information about each of them.

Accessing the HYCU Controllers panel

To access the HYCU Controllers panel, in the navigation pane, click  **HYCU Controllers**.

 **Tip** You can update data related to the data protection environments by clicking  **Refresh**.

Adding a HYCU controller

I want to monitor...	HYCU controller to add	Instructions
On-premises (HYCU) data protection environment	Hybrid Cloud Edition controller	“Adding a Hybrid Cloud Edition controller” below
HYCU for AWS data protection environment	AWS controller	“Adding an AWS controller” on the next page
HYCU for Google Cloud data protection environment	Google Cloud controller	“Adding a Google Cloud controller” on page 466
HYCU for Azure data protection environment	Azure controller	“Adding an Azure controller” on page 466
HYCU for Microsoft 365 and Google Workspace data protection environment	Microsoft 365 and Google Workspace controller	“Adding a Microsoft 365 and Google Workspace controller” on page 467

Adding a Hybrid Cloud Edition controller

Procedure

1. In the HYCU Controllers panel, click  **Add**.
2. Select **Hybrid Cloud Edition controller**, and then click **Next**.
3. Enter the name of the HYCU backup controller.
4. Enter the URL of the HYCU backup controller.
5. Depending on the type of authentication you want to use, do one of the following:

- *Basic authentication:* Make sure the **Use API key authentication** switch is disabled, and then enter the user name and password of an infrastructure group administrator.
- *API key authentication:* Enable the **Use API key authentication** switch, and then enter your API key. For details on how to generate and revoke an API key, see “[Managing API keys](#)” on page 449.

6. Click **Add**.

After you add a Hybrid Cloud Edition controller, you can view data protection information received from this HYCU controller for each related virtual machine, application, file share, server, and volume group. For details, see “[Viewing entity data](#)” on page 472.

Adding an AWS controller

Prerequisites

- You must own a HYCU R-Cloud license. For more information, see “[Licensing](#)” on page 385.
- You must have an active subscription for HYCU for AWS. For details, see HYCU for AWS documentation.

Procedure

1. In the HYCU Controllers panel, click **Add**. The Add Controller dialog box opens.
2. Select **AWS controller**, and then click **Next**. The Add Controller > AWS Controller dialog box opens.
3. Enter the credentials of your user account for accessing HYCU for AWS.
4. Enter the HYCU account ID that you received when you subscribed to HYCU for AWS, and then click **Next**.
5. Select the HYCU for AWS protection sets that you want to monitor. You can also search for a protection set by entering its name in the Search field.
 **Tip** You can see all AWS accounts that are included in each available protection set by clicking .
6. Click **Add**.

After you add an AWS controller, you can select it and click **Details** to see its protection sets and all included accounts.

Adding a Google Cloud controller

Prerequisites

- You must own a HYCU R-Cloud license. For more information, see “[Licensing](#)” on page 385.
- You must have an active subscription for HYCU for Google Cloud. For details, see HYCU for Google Cloud documentation.
- A cloud account must be added to HYCU. For instructions, see “[Adding a Google Cloud service account](#)” on page 367.
- The projects included in the protection set that you plan to monitor must be linked to the Google Cloud billing account that was selected when subscribing to HYCU for Google Cloud. For details, see HYCU for Google Cloud documentation.

Consideration

If the required cloud account is not added to HYCU, the option for adding the Google Cloud controller is grayed out in the HYCU Manager console.

Procedure

1. In the HYCU Controllers panel, click  **Add**. The Add Controller dialog box opens.
2. Select **Google Cloud controller**, and then click **Next**. The Add Controller > Google Cloud Controller dialog box opens.
3. Select the HYCU for Google Cloud protection sets that you want to monitor. You can also search for a protection set by entering its name in the Search field.

 **Tip** You can see all Google Cloud projects that are included in each available protection set by clicking .
4. Click **Add**.

After you add a Google Cloud controller, you can select it and click  **Details** to see its protection sets and all included projects.

Adding an Azure controller

Prerequisites

- You must own a HYCU R-Cloud license. For more information, see “[Licensing](#)” on page 385.

- You must have an active subscription for HYCU for Azure. For details, see HYCU for Azure documentation.
- A cloud account must be added to HYCU. For instructions, see “[Adding an Azure service principal](#)” on page 369.

Consideration

If the required cloud account is not added to HYCU, the option for adding the Azure controller is grayed out in the HYCU Manager console.

Procedure

1. In the HYCU Controllers panel, click **+** **Add**. The Add Controller dialog box opens.
2. Select **Azure controller**, and then click **Next**. The Add Controller > Azure Controller dialog box opens.
3. Select the HYCU for Azure protection sets that you want to monitor. You can also search for a protection set by entering its name in the Search field.
Tip You can see all Azure resource groups that are included in each available protection set by clicking **>**.
4. Click **Add**.

After you add an Azure controller, you can select it and click **Details** to see its protection sets and all included resource groups.

Adding a Microsoft 365 and Google Workspace controller

Prerequisite

You must have an active subscription for HYCU for Microsoft 365 and Google Workspace. For details, see the [HYCU for Microsoft 365 and Google Workspace Quick Start Guide](#).

Procedure

1. In the HYCU Controllers panel, click **+** **Add**. The Add Controller dialog box opens.
2. Select **Microsoft 365 and Google Workspace controller**, and then click **Next**. The Add Controller > Microsoft 365 and Google Workspace Controller dialog box opens.
3. Enter the name of the Microsoft 365 and Google Workspace controller.

4. Enter the URL of your HYCU for Microsoft 365 and Google Workspace web user interface.
5. Enter the access and reseller tokens that you received when you subscribed to HYCU for Microsoft 365 and Google Workspace.
6. Click **Add**.

Viewing information about HYCU controllers

You can view specific information about each added HYCU controller.

However, keep in mind that not all information might be applicable to your data protection scenario.

HYCU controller information	Description
Name	<p>Name of the HYCU controller.</p> <p>A Hybrid Cloud Edition controller is represented by the  icon and the name of the HYCU backup controller. If you use HYCU Manager to monitor also the cloud data protection environments, you can view cloud controllers. A cloud controller is represented by:</p> <ul style="list-style-type: none"> • <i>HYCU for AWS</i>: The  icon and the name of the HYCU for AWS protection set. • <i>HYCU for Google Cloud</i>: The  icon and the name of the Google Cloud service account and the HYCU for Google Cloud protection set. • <i>HYCU for Azure</i>: The  icon and the name of the Azure service principal and the HYCU for Azure protection set. • <i>HYCU for Microsoft 365 and Google Workspace</i>: The  icon and the name of the Microsoft 365 and Google Workspace controller. <p>Note If you click the name of the HYCU controller, you are directed to the relevant web user interface.</p> <p><i>For all HYCU controllers except Microsoft 365 and Google Workspace controllers</i>: You can click any of the icons representing the information about the HYCU controller</p>

HYCU controller information	Description
	and you are automatically directed to the specific panel listing all the corresponding items. If your HYCU controller is a Hybrid Cloud Edition controller, these items are also filtered according to your selection. For example, if you click an icon representing the percentage of the protected virtual machines, you are directed to the Virtual Machines panel listing only all the protected virtual machines.
Version	HYCU software release version on the HYCU backup controller.
Status	Status of the HYCU controller (Active, Suspended, or Unavailable).
Backups	Exact number and percentage of successful and failed backups (also per user groups).
DR-ready VMs	Number of DR-ready virtual machines and servers. A virtual machine or a server is DR-ready if all backups in the current backup chain are stored on one of the cloud targets and a successful platform readiness check is performed during its latest backup.
VM protection	Exact number and percentage of protected and unprotected virtual machines (also per user groups).
App protection	Exact number and percentage of protected and unprotected applications (also per user groups).
Share protection	Exact number and percentage of protected and unprotected file shares (also per user groups).
Policy compliance	Percentage of compliant and non-compliant policies.
Target utilization	Percentage of used storage space on targets.

You can export data that you view in the HYCU Controllers panel to a file in JSON or CSV format. For details on how to do this, see [“Exporting the contents of the panel” on page 324](#).

Editing a HYCU controller

Limitation

You can edit only the Hybrid Cloud Edition controller and the Microsoft 365 and Google Workspace controller.

Procedure

1. In the HYCU Controllers panel, select the HYCU controller that you want to edit, and then click  **Edit**. The Edit Controller dialog box opens.
2. Edit the selected HYCU controller as required.
3. Click **Save**.

Deleting a HYCU controller

Procedure

1. In the HYCU Controllers panel, select the HYCU controller that you want to delete from HYCU Manager, and then click  **Delete**.
2. Click **Delete** to confirm that you want to delete the selected HYCU controller.

Using the HYCU Manager console

The HYCU Manager console provides you with an at-a-glance overview of the data collected from all the data protection environments for which you are responsible.

Accessing the Console panel

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

Within each widget in the HYCU Manager console, you can find information related to your data protection environments. However, keep in mind that not all widgets might be applicable to your data protection scenario.

Console widget	Description
Virtual Machines	<p>Shows the percentage of protected and unprotected virtual machines in your data protection environments, and the percentage of compliant and non-compliant virtual machines. A virtual machine is considered:</p> <ul style="list-style-type: none"> Protected: If it has at least one valid backup available and does not have the Exclude policy assigned. Compliant: If the time since the last successful backup is lower than its RPO and the estimated time to recover is lower than its RTO.
Applications	<p>Shows the percentage of protected and unprotected applications in your data protection environments, and the percentage of compliant and non-compliant applications. An application is considered:</p> <ul style="list-style-type: none"> Protected: If it has at least one valid backup available and does not have the Exclude policy assigned. Compliant: If the time since the last successful backup is lower than its RPO and the estimated time to recover is lower than its RTO.
HYCU Controllers	<p>Shows the number of available and unavailable HYCU controllers in your data protection environments.</p>
Backups	<p>Shows the percentage of successful and unsuccessful backups in your data protection environments for the last 7 days, the number of successful backups, and the number of virtual machines with the DR-ready status. You can safely ignore the DR-ready label if you do not plan to use the SpinUp functionality.</p>
Shares	<p>Shows the percentage of protected and unprotected file shares in your data protection environments. A file share is considered protected if it has at least one valid backup available and does not have the Exclude policy assigned.</p>
Targets	<p>Shows the list of all targets in your data protection environments, and the information on how much space is used and available for storing data on each target and on all targets in your data protection environments</p>

Console widget	Description
	combined.
Policies	Shows the number of entities in your data protection environments, the number of entities that have no policy assigned, and the number of entities that are compliant and non-compliant with the RPO and RTO set in their assigned policy. The number of compliant and non-compliant entities for specific policies is also shown.
M365 / G Workspace	Shows the overview of protected users, Microsoft 365 SharePoint sites, Groups and Teams, and Google Workspace Shared Drive files in your data protection environments. For users, the total number of protected emails, files, contacts, calendar items, and tasks is also displayed.

Viewing entity data

If you have one or more Hybrid Cloud Edition controllers added to HYCU Manager, you can view data protection information received from these HYCU controllers for each related virtual machine, application, file share, server, and volume group.

Depending on what kind of data protection information you want to view, access one of the following panels:

- Accessing the Applications panel
To access the Applications panel, in the navigation pane, click  **Applications**.
- Accessing the Virtual Machines panel
To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.
- Accessing the Shares panel
To access the Shares panel, in the navigation pane, click  **Shares**.

- Accessing the Volume Groups panel

To access the Volume Groups panel, in the navigation pane, click  **Volume Groups**.

 **Note** HYCU Manager performs the automatic synchronization of entities every five minutes. However, you can at any time update the list of entities also manually by clicking  **Refresh** in the corresponding panel.

Within each panel, besides viewing data protection information related to your entities, you can also do the following:

I want to...	Instructions
<p>Navigate directly from HYCU Manager to the web user interface of the HYCU backup controller containing the entity whose restore point I select.</p> <p> Note By navigating to the specific web user interface, you will be able not only to perform restore operations related to the selected restore point, but also all other operations that are available in your data protection environment.</p>	<p>To navigate to the web user interface of the HYCU backup controller containing the preferred restore point, follow these steps:</p> <ol style="list-style-type: none"> 1. In the Applications, Virtual Machines, Shares, or Volume Groups panel, click the entity to whose restore point you want to navigate. The Detail view appears at the bottom of the screen. 2. In the Detail view, select the preferred restore point. 3. Click  Navigate to Restore Point. <p>You are automatically directed to the corresponding HYCU web user interface. For instructions on how to sign in to HYCU, see “Signing in to HYCU” on page 49.</p>
<p>Apply different types of filters to entities, including filtering them by the HYCU controller to which they belong.</p>	<p>To filter entities by the HYCU controller, follow these steps:</p> <ol style="list-style-type: none"> 1. In the selected panel, click  Filters. 2. In the side panel that opens, from the Backup controller name drop-down menu, select the preferred Hybrid Cloud Edition controller. 3. Click Apply Filters. <p>For more details on filtering, see “Filtering</p>

I want to...	Instructions
	and sorting data” on page 317.
Export data to a file in JSON or CSV format.	For details on how to do this, see “ Exporting the contents of the panel ” on page 324.

Viewing events

You can use the Events panel to view all events that occurred on your HYCU Manager and check details about the selected event, list events that match the specified filter, configure HYCU to send notifications when events occur, export the contents of the panel to a file in JSON or CSV format, and enable the purge of events.

Accessing the Events panel

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

I want to...	Procedure
View events and check details about the selected event.	“ Managing HYCU events ” on page 301
Apply filters to events.	“ Filtering and sorting data ” on page 317
Configure HYCU to send notifications when events occur.	“ Configuring event notifications ” on page 302
Export event data.	“ Exporting the contents of the panel ” on page 324
Enable the purge of events.	“ Enabling the purge of events and jobs ” on page 305

Performing administration tasks

After you deploy the HYCU virtual appliance in HYCU Manager mode, you can perform various administration tasks through the  **Administration** menu.

Note The procedures for administering HYCU deployed in the HYCU Manager mode are the same as for HYCU deployed in the HYCU Backup Controller mode. Therefore, in most cases, you can follow the same instructions.

Keep in mind that a varied set of administration tasks is available depending on the selected deployment mode.

I want to...	Procedure
Add Azure or Google Cloud accounts to be able to monitor cloud data protection environments.	“Adding a cloud account” on page 363
Integrate HYCU Manager with identity providers.	“Integrating HYCU with identity providers” on page 373
Configure log file settings to troubleshoot problems if HYCU does not perform as expected.	“Setting up logging” on page 389
Change network settings.	“Changing network settings” on page 392
Securely store, access, and manage my credentials (secrets) by employing the Conjur secrets management solution.	“Managing secrets” on page 404
Configure an SMTP server.	“Configuring an SMTP server” on page 407
Upgrade HYCU to a new available version.	“Upgrading HYCU” on page 408
Important Before upgrading, make sure you have added the source where your HYCU Manager virtual machine resides as described in “Adding sources” on page 53 .	
Access the HYCU Manager virtual machine by using SSH.	“Accessing the HYCU backup controller virtual machine by using SSH” on page 443
Configure the SSL certificate.	“Configuring SSL certificates”

I want to...	Procedure
	on page 432
Manage HYCU Manager users.	“Managing users” below

In addition, you can do the following:

- Use hyCLI. For details, see “Using the command-line interface” on page 460.
- Use the HYCU REST API Explorer. For details, see “Using the HYCU REST API Explorer” on page 460.

Managing users

You can use the Manage Users dialog box to give the specified users access to HYCU Manager. Managing users includes creating, editing, deleting, and activating or deactivating users.

Accessing the User Management dialog box

To access the User Management dialog box, from the  **Administration** menu, select **User Management**.

Creating a new user

Procedure

1. In the User Management dialog box, click  **New**.
2. Depending on what kind of user you are adding, enter one of the following:
 - *For a HYCU user, an AD user, an OIDC user, or an OIDC group:* User name
 - *For an AD group:* Common name

 **Important** When entering a name, make sure it complies with the SAM account name limitations—name length may not exceed 20 characters and contain any of the following characters: "/ \ [] : ; | = , + * ? < >. In addition, HYCU does not allow the at sign (@) in the name. If your environment requires it, these limitations can be overridden by editing the `ad.username.filter.regex` configuration setting. However, this is not supported and could cause authentication issues. For details on how to customize HYCU configuration settings, see “Customizing HYCU configuration settings” on page 523.

3. From the Authentication Type drop-down menu, select one of the following authentication types, and then follow the instructions:

Authentication type	Instructions
HYCU	<ol style="list-style-type: none"> a. From the Language drop-down menu, select the preferred language for the user. b. In the Name field, enter a display name for the user. c. <i>Optional.</i> In the Email field, enter the email address of the user. d. In the Password field, enter the user password. <p>Note The minimum password length is six characters.</p>
OIDC User	<ol style="list-style-type: none"> a. From the Language drop-down menu, select the preferred language for the user. b. From the Identity Provider drop-down menu, select the identity provider. c. In the Identity Provider User ID field, enter the ID of the identity provider user. <p>Note Depending on your identity provider, the user ID corresponds to the following:</p> <ul style="list-style-type: none"> • <i>Active Directory Federation Services:</i> Object GUID • <i>Google:</i> User email address • <i>Keycloak:</i> User ID • <i>Microsoft:</i> Object ID • <i>Okta:</i> Part of the URL when you navigate to the user's profile <p>For details, see the respective identity provider documentation.</p>
OIDC Group	<ol style="list-style-type: none"> a. From the Language drop-down menu, select the preferred language for the group. b. From the Identity Provider drop-down menu, select the identity provider. c. In the Identity Provider Group ID field, enter the ID of the identity provider group.

Authentication type	Instructions
	<p>Note Depending on your identity provider, the group ID corresponds to the following:</p> <ul style="list-style-type: none"> • <i>Active Directory Federation Services</i>: Object GUID • <i>Keycloak</i>: Group ID • <i>Microsoft</i>: Group Object IDs • <i>Okta</i>: Group name <p>For details, see the respective identity provider documentation.</p>
AD User	<ol style="list-style-type: none"> From the Language drop-down menu, select the preferred language for the user. From the Identity Provider drop-down menu, select the Active Directory the AD user belongs to.
AD Group	<ol style="list-style-type: none"> From the Language drop-down menu, select the preferred language for the user. From the Identity Provider drop-down menu, select the Active Directory the AD group belongs to.

4. *Only if you are adding a HYCU user, an AD user, or an AD group.* Use the **Two-factor authentication** switch if you want to enable two-factor authentication for the user, and then select one of the following two-factor authentication methods:

- **Time-based one-time password**

This option enables the use of a time-based one-time password (OTP) generated by an OTP application. The user needs to set up an OTP during the first sign-in after two-factor authentication is enabled.

- **FIDO**

This option enables the use of an authenticator complying with FIDO protocols (FIDO authenticator). The user needs to register a FIDO authenticator. For details, see “[Managing FIDO authenticators](#)” on [page 450](#).

5. *Only if you enabled two-factor authentication.* To prevent the user from disabling two-factor authentication, make sure the **User cannot disable two-factor authentication** check box is selected. If you clear the check box, the user can disable two-factor authentication. An infrastructure group administrator can disable two-factor authentication even if this option is enabled.

 **Note** If a user disables two-factor authentication, the administrator is notified with a security warning.

6. Click **Save**.

The user is added to the list of all users.

You can later do the following:

- Edit any of the existing HYCU or identity provider users by clicking  **Edit** and making the required modifications. Keep in mind that the built-in user, AD users, and AD groups cannot be edited.
- Enable or disable specific users from signing in to HYCU. For details, see [“Performing administration tasks” on page 474](#).
- Delete any of the existing users by clicking  **Delete**. Keep in mind that the built-in user cannot be deleted.

 **Important** *For creating a user by using hyCLI:* As opposed to creating a new user through the HYCU Manager console where this is done automatically, if using hyCLI, you must also add the created user to the infrastructure group and assign this user the Administrator role.

Chapter 14

Employing Nutanix Mine with HYCU

Nutanix Mine with HYCU is the only hyperconverged backup and recovery solution that provides backup and recovery as a native service of the Nutanix platform and eliminates the need for an isolated infrastructure. It allows you to preserve hyperconverged infrastructure simplicity while ensuring all of your data is fully protected.

The Nutanix Mine with HYCU solution allows you to use a single pane of glass to manage both production and backup infrastructures. You can optimize your data protection environment by introducing Nutanix Mine storage as a target, which will increase your Nutanix Mine cluster's effective storage capacity, and improve backup and restore performance.

Task	Instructions
1. Register HYCU as a service of the Nutanix Mine platform.	“Registering HYCU with Nutanix Prism” below
2. Add Nutanix Mine storage as a target for storing protected data.	“Setting up a Nutanix target” on page 85
3. Use a single pane of glass to manage both production and backup infrastructures.	“Accessing HYCU from the Nutanix Prism web console” on the next page

Registering HYCU with Nutanix Prism

Prerequisites

- You acquired a Nutanix Mine appliance.
- The HYCU backup controller must reside on a Nutanix Mine cluster and this cluster must be added to HYCU as a source. For details, see [“Deploying](#)

HYCU to a Nutanix AHV cluster” on page 35 and “Adding a Nutanix cluster” on page 54.

- *For repeating the registration procedure:* Currently running jobs that you do not want to be aborted must be finished.

Considerations

- All instructions that apply to the Nutanix AHV cluster apply also to the Nutanix Mine cluster.
- If you receive a warning message indicating that there have been changes on the Nutanix Mine cluster, you must register HYCU with Nutanix Prism again. You receive such a message in the following cases:
 - The IP address/host name or port of the HYCU backup controller was changed.
 - The AOS of the Nutanix Mine cluster was upgraded to a new version.
 - A new HYCU backup controller was added to the Nutanix Mine cluster.

Accessing the Sources dialog box

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

Procedure

1. In the Sources dialog box, on the Hypervisor tab, from the list of all sources, select the Nutanix Mine cluster.
2. Click  **Register with Prism**.
3. Click **Yes** to confirm that you want to proceed.

 **Important** Registering HYCU with Nutanix Prism may take some time. The Nutanix Prism web console will not be available during this time.

You can at any time unregister HYCU from Nutanix Prism. To do so, select the respective Nutanix Mine cluster, and then click  **Unregister from Prism**.

Accessing HYCU from the Nutanix Prism web console

After you enable register HYCU with Nutanix Prism, you can view the Nutanix Mine with HYCU dashboard and also launch the HYCU web user interface

directly from the Nutanix Prism web console.

Procedure

1. Sign in to the Nutanix Prism web console.
2. From the drop-down menu on the left, select **HYCU**. The Nutanix Mine with HYCU dashboard appears.
3. Click **Launch HYCU**. The HYCU user web interface opens in another tab, allowing you to manage your data protection environment.

Viewing the Nutanix Mine with HYCU dashboard

The Nutanix Mine with HYCU dashboard provides you with an at-a-glance overview of the data protection status in your environment. This intuitive dashboard enables you to monitor all data protection activity and to quickly identify areas that need your attention. You can use this dashboard as a starting point for your everyday tasks related to data protection because it enables you to easily access the area of interest by simply clicking the corresponding links.

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

Dashboard widget	Description
VM Protection Status	Percentage of virtual machines that are protected and the number of protected and unprotected virtual machines in the data protection environment. A virtual machine is considered protected if it has at least one valid backup available and does not have the Exclude policy assigned. For details on protecting virtual machines, see “Protecting virtual machines” on page 128 .
App Protection Status	Percentage of applications that are protected and the number of protected and unprotected applications in the data protection environment. An application is considered protected if it has at least one valid backup available and does not have the Exclude policy assigned. For details on protecting applications, see “Protecting applications” on page 209 .

Dashboard widget	Description
Compliance	<p>Percentage of policies that are compliant and the number of compliant and non-compliant policies in the data protection environment. A policy is considered compliant if all entities that have this policy assigned are compliant with the RPO and RTO requirements. For details on policies, see “Defining your backup strategy” on page 107.</p>
Backups	<p>Backup success rates for the last seven days.</p>
Mine Storage	<ul style="list-style-type: none"> • List of Nutanix targets, and the information on how much space is used and available for storing data, the data compression ratio, and the data deduplication ratio. • List of Nutanix Objects and S3 compatible targets, and the information on how much space is used and available for storing data. <p>For details on targets, see “Setting up targets” on page 78.</p>
Target Summary	<p>List of all targets in the data protection environment, not including the Nutanix, Nutanix Objects, and S3 compatible targets, and the information on how much space is used and available for storing data. For details on targets, see “Setting up targets” on page 78.</p>
HYCU Controller	<p>Information on whether the HYCU backup controller is protected and its license is valid, as well as the resource information about the HYCU backup controller (storage, memory, and vCPU). For details on what to do if any of the resource values reaches a critical value, see “Adjusting the HYCU backup controller resources” on page 341.</p>
Events	<p>Number of events in the data protection environment in the last 56 hours according to their status (Success, Warning, and Failed). For details on events, see “Managing HYCU events” on page 301.</p>
Jobs	<p>Number of jobs in the data protection environment in the last 56 hours according to their status (Success, Warning, Failed, In progress, and Queued). For details on jobs, see “Managing HYCU jobs” on page 300.</p>

 **Tip** You can rearrange the dashboard widgets by dragging and dropping them so that you have the most important data you want to view at the top of your dashboard.

Chapter 15

Protecting data across on-premises and cloud environments

The SpinUp functionality ensures business continuity of your data protection environment across different infrastructures. You can ensure data resilience by migrating virtual machines across the on-premises and cloud (AWS, Google Cloud, global Azure, or Azure Government) infrastructures. In the event of a disaster in your on-premises environment, the SpinUp functionality provides disaster recovery of data to cloud.

Depending on your cloud environment, see one of the following sections:

- “Protecting data across on-premises and AWS environments” below
- “Protecting data across on-premises and Google Cloud environments” on page 496
- “Protecting data across on-premises and Azure environments” on page 505
- “Protecting data across on-premises and Azure Government environments” on page 515

Protecting data across on-premises and AWS environments

You can use the SpinUp functionality to migrate protected data across the on-premises and Amazon Web Services (AWS) environments. In the event of a disaster, it provides disaster recovery of data to AWS.

Depending on what you want to do, see one of the following:

I want to...	Instructions
Migrate protected data across the on-	“Migrating virtual machines across

I want to...	Instructions
on-premises and AWS environments.	different environments ” on the next page
Perform disaster recovery of data to AWS.	“Performing disaster recovery of data to AWS” on page 493

Prerequisites

- You must have an active subscription for HYCU R-Cloud or HYCU for AWS. For instructions, see HYCU R-Cloud or HYCU for AWS documentation.
- An AWS user account must be added to HYCU. For instructions, see [“Adding an AWS user account”](#) on page 365.
- A HYCU account must be added to HYCU. For instructions, see [“Adding a HYCU account”](#) on page 371.
- You must own a HYCU R-Cloud license. For details, see [“Licensing”](#) on page 385.

Migrating virtual machines across different environments

You can migrate protected data across the on-premises and AWS environments as follows:

- [“Migrating data to cloud”](#) below
- [“Migrating data from cloud”](#) on page 491

Migrating data to cloud

You can migrate virtual machines, servers, and applications running on virtual machines and servers to cloud by using the SpinUp functionality. Keep in mind that when you migrate an application, the whole virtual machine or server on which this application is running is migrated to cloud.

 **Note** The instructions for protecting virtual machine data apply also to servers except where specifically stated otherwise.

Prerequisites

- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected, and must have

a successful platform readiness check during the backup. For more information, see [“SpinUp specifics” on page 142](#).

- In the HYCU R-Cloud or HYCU for AWS web user interface, make sure that the AWS user account is granted the Administrator role in the Subscription context.
- *For Windows virtual machines:* In AWS, the required permissions must be specified in your IAM policy for VM Import/Export. Make sure to use `hycu-tmp` instead of `mys3bucket`, `disk-image-file-bucket`, and `export-bucket` in the sample policies. For instruction on how to specify required permissions for VM Import/Export, see AWS documentation.

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for migrating data.
- *For Nutanix clusters:* You cannot migrate volume groups.
- *For vSphere environments:*
 - You cannot migrate virtual machine templates.
 - Migrating data from snapshots is not supported.

Consideration

If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for migrating data.

Depending on whether you want to migrate virtual machine or application data to cloud, access one of the following panels:

- Accessing the Virtual Machines panel
To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.
- Accessing the Applications panel
To access the Applications panel, in the navigation pane, click  **Applications**.

Procedure

1. In the Virtual Machines or Applications panel, select the entity that you want to migrate.
2. In the Detail view that appears at the bottom of the screen, select the virtual machine or application restore point that you want to use for the migration.

Note The Detail view appears only if you click an entity. Selecting the check box before the name of the entity will not open the Detail view.
3. Click  **SpinUp to Cloud**.
4. Select **SpinUp to AWS**, and then click **Next**.
5. From the AWS user account drop-down menu, select the AWS user account.

Note By default, the AWS account to which the selected AWS user account belongs and to which the entity will be migrated is displayed.
6. From the HYCU account drop-down menu, select the HYCU account.
7. From the AWS account ID, select the AWS account ID.
8. From the Region drop-down menu, select the AWS region to which you want to migrate the entity.
9. From the Availability zone drop-down menu, select the preferred Availability Zone within the selected AWS region, and then click **Next**. The VM Settings dialog box opens.
10. From the SpinUp from drop-down menu, select which tier you want to use for the migration. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest migration of data to cloud.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
11. In the New VM name field, enter a name for the migrated virtual machine instance.

Important Make sure the migrated virtual machine instance name is unique.
12. In the vCPU threads field, enter the number of CPUs for the migrated virtual machine instance multiplied by the number of cores per CPU and the

number of threads per core.

13. In the Memory field, set the amount of memory (in GiB) for the migrated virtual machine instance. The default value is the amount of memory in GiB of the original virtual machine.
14. From the Virtual machine type drop-down menu, select the machine type for the migrated virtual machine instance.

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

15. *Only if virtual disks have been excluded from the backup (manually or automatically)*: Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the migrated virtual machine.
16. Under Network interfaces, do the following:
 - a. Click **Add Network Interface**. The Add Network Interface dialog box opens.
 - b. From the Subnets drop-down menu, select the subnet.
 - c. From the Security groups drop-down menu, select the AWS security group.
 - d. In the Public address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	The network interface does not use a public IP address.
Auto-assign	 Note Auto-assign will not work if the Auto-assign public IPv4 address on a subnet option is set to No or if more than one network interface is specified.
Elastic IP (Reserved)	The network interface uses an elastic public IP address that was reserved in Amazon EC2 in advance.
Elastic IP	The network interface uses a new elastic public

(New)	<p>IP address.</p> <p>Note Allocation of the IP address in Amazon EC2 is performed at the very beginning of the migration. If the allocation fails, the migration task is terminated without being logged.</p>
-------	---

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

Option	Description
Auto-assign	The network interface uses an automatically allocated private IP address.
Custom	<p>The network interface uses a private IP address that you define.</p> <p>Important Using this option might result in IP address conflicts.</p>

f. Click **Add**.

You can also edit any of the existing network interfaces (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**). Keep in mind that you cannot migrate the virtual machine without a network interface.

17. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:

- **Linux**
- **Windows**

18. Under Operating system license, select one of the following options:

OS license option	Select this option if you want to...
Keep existing license	<p>Keep the existing OS license on the migrated virtual machine instance.</p> <p>Important Make sure that the existing license is applicable also in AWS.</p>

OS license option	Select this option if you want to...
<i>Available only for the Windows Server OS.</i> Replace existing license with AWS license	Replace the existing OS license with an AWS license on the migrated virtual machine instance.

19. Click **SpinUp**.

The Migration to cloud job starts. When it finishes successfully, you can check the migrated virtual machine instance in the Instances panel in HYCU R-Cloud or HYCU for AWS. For details, see HYCU R-Cloud or HYCU for AWS documentation.

After migrating data to cloud

- *For Windows virtual machines:* If you decided to keep the existing OS license on the migrated virtual machine instance, reactivate the Windows license.
- Enable the protection of the migrated virtual machine instances by using HYCU R-Cloud or HYCU for AWS. For details, see HYCU R-Cloud or HYCU for AWS documentation.

Migrating data from cloud

You can migrate virtual machine instances from cloud by using the SpinUp functionality.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click  **SpinUp from Cloud**.
2. Select **SpinUp from AWS**, and then click **Next**.
3. From the AWS user account drop-down menu, select the AWS user account.

Note By default, the AWS account to which the selected AWS user account belongs and from which the virtual machine instance will be migrated is displayed.
4. From the HYCU account drop-down menu, select the HYCU account.
5. From the Virtual machine drop-down menu, select the virtual machine instance that you want to migrate.

6. From the Checkpoint drop-down menu, select the checkpoint from which you want to migrate virtual machine instance data.
7. Click **Next**. The VM Settings dialog box opens.
8. From the Storage container drop-down menu, select where you want to migrate the virtual machine instance.
9. In the New VM name field, enter a name for the migrated virtual machine.
10. *Only if the virtual machine instance that you are migrating was created in the on-premises environment, migrated to cloud, and now you are migrating it back to the on-premises environment.* If you want the migrated virtual machine to have the same virtual machine settings as it had in the on-premises environment, enable the **Keep original on-premises settings** option, and then continue with step 13.

Otherwise, leave the Keep original on-premises settings option disabled and continue with the next step.

11. Specify the following values for the migrated virtual machine:

Value	Description
vCPU(s)	Number of virtual CPUs.
Cores per vCPU	Number of cores per virtual CPU.
Memory	Amount of memory (in GiB).

 **Note** The default values are the ones that the virtual machine had in the environment in which it was created, either in the on-premises or cloud one.

12. Under Network adapters, depending on your data protection needs, do one of the following:
 - Add one or more network adapters:
 - a. Click **Add Network Adapter**. The New Network Adapter dialog box opens.
 - b. From the Network drop-down menu, select the network.
 - c. Click **Save**.
 - Edit any of the existing network adapters to connect the virtual machine to a different network. To do so, select a network adapter, click  **Edit**, and make the required modification.

- Delete any of the existing network adapters by selecting it, and then clicking  **Delete**. If you delete all the existing network adapters, your virtual machine will be migrated without network connectivity.

13. Use the **Power virtual machine on** switch if you want to turn the migrated virtual machine on after the migration.
14. Click **SpinUp**.

The Migration from cloud job starts. When it finishes successfully, you can view the migrated virtual machine in the Virtual Machines panel.

After migrating data from cloud

- *For virtual machines on a Nutanix AHV cluster:* Make sure that the latest version of NGT is installed on the virtual machine. For instructions, see Nutanix documentation.
- *For virtual machines on a Nutanix ESXi cluster:* Make sure that the latest versions of VMware Tools and NGT are installed on the virtual machine. For instructions, see Nutanix and VMware documentation.
- *For virtual machines in a vSphere environment:* Make sure that the latest version of VMware Tools is installed on the virtual machine. For instructions, see VMware documentation.
- *For Linux virtual machines:* If a virtual machine on a Nutanix ESXi cluster or in a vSphere environment does not boot, change the controller type from SCSI to SATA, and then install the necessary SCSI drivers to switch back to SCSI.
- *For Windows virtual machines:* Reactivate the Windows licenses.
- *Only if you migrated virtual machines without network connectivity.* Make sure to configure the network settings on the virtual machine.
- Enable protection of the migrated data. For details on how to do this, see “[Protecting virtual machines](#)” on page 128 and “[Protecting applications](#)” on page 209.

Performing disaster recovery of data to AWS

You can perform disaster recovery of data from the on-premises environment to AWS in the event of a disaster.

Prerequisites

- You must have the HYCU virtual appliance image for AWS. To obtain the image, on AWS Marketplace, browse the AMI Catalog for the following Amazon Machine Image (AMI):

	The image name is represented in the following format:
Image name	hycu-<Version>-<Revision> For example: hycu-5.0.0-3634
Owner	The owner is represented by the following AWS account ID: 787223699828

For instructions, see AWS documentation.

- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected and must have the DR-ready status. For more information, see “[SpinUp specifics](#)” on [page 142](#).

Considerations

- When the HYCU backup controller is deployed in AWS, changing network settings is prevented in HYCU.
- Make sure the imported target is in the region to which you plan to migrate your virtual machines. This ensures the disaster recovery process is as fast and as cost-effective as possible.
- After you deploy the HYCU backup controller and use it to perform disaster recovery, you can keep the HYCU backup controller to stay prepared for disaster recovery in the future. However, every time you upgrade HYCU, you must deploy a new HYCU backup controller to be able to perform disaster recovery of data to cloud.

Procedure

1. Deploy a HYCU backup controller. To do so, select the HYCU virtual appliance image in the AWS AMI Catalog, and then click **Launch Instance with AMI**. For instructions, see AWS documentation.
2. In AWS, create a new firewall rule to allow ingress network traffic through TCP port 8443 from the entire subnetwork to which the HYCU backup controller belongs. For instructions, see AWS documentation.

3. Sign in to the HYCU web user interface by specifying the following URL:

```
https://<IPAddress>:8443
```

In this instance, *<IPAddress>* is the external IP address of the newly deployed HYCU backup controller.

ⓘ Important The credentials you provided in AWS during virtual machine instance creation cannot be used to sign in to HYCU and perform disaster recovery of data to AWS. For details on what credentials you can use to sign in to HYCU or to access the HYCU backup controller by using SSH, see “[Signing in to HYCU](#)” on page 49 or “[Accessing the HYCU backup controller virtual machine by using SSH](#)” on page 443.

4. Add an AWS user account. For instructions, see “[Adding an AWS user account](#)” on page 365.
5. Import the Amazon S3 target on which your backup data is stored to HYCU:
 - a. In the Targets panel, click  **Import**. The Import Target dialog box opens.
 - b. From the Type drop-down menu, select **Amazon S3 / S3 Compatible**.
 - c. In the Service endpoint field, enter the service endpoint URL.
 - d. In the Bucket name field, enter the Amazon S3 bucket name as it was specified in the original target configuration.
 - e. In the Access key ID field, enter the access key ID of your AWS user account.
 - f. In the Secret access key, enter the secret access key of your AWS user account.
 - g. Enable the **Path style access** switch if you want HYCU to use a path-style URL (<https://s3.amazonaws.com/<BucketName>>) to access the bucket. HYCU by default uses a virtual-hosted-style URL (<https://<BucketName>.s3.amazonaws.com>).
 - h. Click **Next**. The Import Backup Catalog dialog box opens.
 - i. Select the HYCU backup controller whose backup data you want to import, and then click **Next**.
 - j. In the Multiple Targets dialog box, one or more targets that store backup data are displayed. If any additional targets are found, select them one by one and specify the values so that they match the original target

configuration. For each target, click **Validate** to check the configuration.

- k. After you validate all the targets, click **Import**.
6. Migrate your virtual machines or applications to cloud. For instructions, see “[Migrating data to cloud](#)” on page 486.

Protecting data across on-premises and Google Cloud environments

You can use the SpinUp functionality to migrate protected data across the on-premises and Google Cloud environments. In the event of a disaster, it provides disaster recovery of data to Google Cloud.

Depending on what you want to do, see one of the following:

I want to...	Instructions
Migrate protected data across the on-premises and Google Cloud environments.	“ Migrating virtual machines across different environments ” on the next page
Perform disaster recovery of data to Google Cloud.	“ Performing disaster recovery of data to Google Cloud ” on page 503

Prerequisites

- You have an active subscription for HYCU R-Cloud or HYCU for Google Cloud. For instructions, see HYCU R-Cloud or HYCU for Google Cloud documentation.
- A Google Cloud service account is added to HYCU. For instructions, see “[Adding a Google Cloud service account](#)” on page 367.
- *Only if your cloud data is protected with HYCU R-Cloud.* A HYCU account must be added to HYCU. For instructions, see “[Adding a HYCU account](#)” on page 371.
- You must own a HYCU R-Cloud license. For details, see “[Licensing](#)” on page 385.

Migrating virtual machines across different environments

You can migrate protected data across the on-premises and Google Cloud environments:

- “Migrating data to cloud” below
- “Migrating data from cloud” on page 501

Migrating data to cloud

You can migrate virtual machines, servers, and applications running on virtual machines and servers to cloud by using the SpinUp functionality. Keep in mind that when you migrate an application, the whole virtual machine or server on which this application is running is migrated to cloud.

 **Note** The instructions for protecting virtual machine data apply also to servers except where specifically stated otherwise.

Prerequisite

The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected, and must have a successful platform readiness check during the backup. For more information, see “SpinUp specifics” on page 142.

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for migrating data.
- *For Nutanix clusters:* You cannot migrate volume groups.
- *For vSphere environments:*
 - You cannot migrate virtual machine templates.
 - Migrating data from snapshots is not supported.

Consideration

If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for migrating data.

Depending on whether you want to migrate virtual machine or application data to cloud, access one of the following panels:

- Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

- Accessing the Applications panel

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

Procedure

1. In the Virtual Machines or Applications panel, select the entity that you want to migrate.
2. In the Detail view that appears at the bottom of the screen, select the virtual machine or application restore point that you want to use for the migration.

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

3. Click  **SpinUp to Cloud**.
4. Select **SpinUp to Google Cloud**, and then click **Next**.
5. From the Cloud account drop-down menu, select the Google Cloud service account to which the project where you want to migrate the virtual machine is linked.
6. From the HYCU account drop-down menu, select the HYCU account.

 **Important** If you use HYCU for Google Cloud to protect your data, select **No HYCU account**.

7. From the Project, Target region, and Target zone drop-down menus, select the required values, and then click **Next**. The VM Settings dialog box opens.
8. From the SpinUp from drop-down menu, select which tier you want to use for the migration. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest migration of data to cloud.
 - **Backup**
 - **Copy**

- **Archive**
- **Snapshot**

9. In the New VM name field, enter a name for the migrated virtual machine instance.

! **Important** Make sure the migrated virtual machine instance name is unique.

10. In the vCPU cores field, enter the number of virtual CPUs for the migrated virtual machine multiplied by the number of cores per virtual CPU.

11. In the Memory field, set the amount of memory (in GiB) for the migrated virtual machine instance. The default value is the amount of memory in GiB of the original virtual machine.

12. From the Virtual machine type drop-down menu, select the machine type for the migrated virtual machine instance.

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

13. *Only if virtual disks have been excluded from the backup (manually or automatically)*: Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the migrated virtual machine.

14. Under Network interfaces, the default network interface is displayed and you can check to which network it is assigned (based on the selected project and region). If required, you can also modify network settings.

Modifying network settings

Depending on your data protection needs, you can leave the default network interface or do one of the following:

- Add a new network interface:
 - Click **Add Network Interface**.

! **Note** The maximum number of network interfaces that you can add depends on the selected virtual machine type.

- From the Target networks drop-down menu, select a network to which you want to add the migrated virtual machine instance. You

can choose among the networks configured in the selected project and other networks that your cloud account has access to.

- c. Select the external address type for the network interface and, if required, the name of the preferred external IP address resource. For details, see HYCU R-Cloud or HYCU for Google Cloud documentation.
- d. Select the internal address type for the network interface and, if required, depending on the address type, do one of the following:
 - In the Internal address field, enter the preferred IP address.
 - From the Internal address drop-down menu, select the name of the preferred internal IP address resource.For details, see HYCU R-Cloud or HYCU for Google Cloud documentation.
- e. Click **Save**.

- Select another network for the existing network interface by selecting it, clicking  **Edit** and making the required modifications.
- Delete the existing network interface by selecting it, and then clicking  **Delete**.

15. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
 - **Linux**
 - **Windows**
16. Click **SpinUp**.

The Migration to cloud job starts. When it finishes successfully, you can check the migrated virtual machine instance in the Instances panel in HYCU R-Cloud or HYCU for Google Cloud. For details, see HYCU R-Cloud or HYCU for Google Cloud documentation.

After migrating data to cloud

- Install the Google Compute Engine guest environment on the virtual machine.
- *For Windows virtual machines:* Reactivate the Windows licenses.
- Enable protection of the migrated virtual machines by using HYCU R-Cloud or HYCU for Google Cloud. For details, see HYCU R-Cloud or HYCU for Google Cloud documentation.

Migrating data from cloud

You can migrate virtual machine instances from cloud by using the SpinUp functionality.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click  **SpinUp from Cloud**.
2. Select **SpinUp from Google Cloud**, and then click **Next**.
3. From the Cloud account drop-down menu, select the Google Cloud service account to which the project containing the virtual machine instance that you want to migrate is linked.
4. From the HYCU account drop-down menu, select the HYCU account.

 **Important** If you use HYCU for Google Cloud to protect your data, select **No HYCU account**.

5. From the Project drop-down menu, select the Google Cloud project to which the virtual machine instance that you want to migrate belongs.
6. From the Virtual machine drop-down menu, select the virtual machine instance that you want to migrate.
7. From the Checkpoint drop-down menu, select the checkpoint from which you want to migrate virtual machine instance data.
8. Click **Next**. The VM Settings dialog box opens.
9. From the Storage container drop-down menu, select where you want to migrate the virtual machine instance.
10. In the New VM name field, enter a name for the migrated virtual machine.
11. *Only if the virtual machine instance that you are migrating was created in the on-premises environment, migrated to cloud, and now you are migrating it back to the on-premises environment.* If you want the migrated virtual machine to have the same virtual machine settings as it had in the on-premises environment, enable the **Keep original on-premises settings** option, and then continue with step 13.

Otherwise, leave the Keep original on-premises settings option disabled and continue with the next step.

12. Specify the following values for the migrated virtual machine:

- The number of virtual CPUs.
- The number of cores per virtual CPU.
- The amount of memory (in GiB).

 **Note** The default values are the ones that the virtual machine had in the environment in which it was created, either in the on-premises or cloud one.

13. Under Network adapters, depending on your data protection needs, do one of the following:
 - Add one or more network adapters:
 - a. Click **Add network adapter**. The New Network Adapter dialog box opens.
 - b. From the Networks drop-down menu, select the network for the virtual adapter.
 - c. Click **Save**.
 - Edit any of the existing network adapters to connect the virtual machine to a different network. To do so, select a network adapter, click  **Edit**, and make the required modification.
 - Delete any of the existing network adapters by selecting it, and then clicking  **Delete**. If you delete all the existing network adapters, your virtual machine will be migrated without network connectivity.
14. Use the **Power virtual machine on** switch if you want to turn the migrated virtual machine on after the migration.
15. Click **SpinUp**.

The Migration from cloud job starts. When it finishes successfully, you can view the migrated virtual machine in the Virtual Machines panel.

After migrating data from cloud

- Remove the Google Compute Engine guest environment from the virtual machine.
- *For virtual machines on a Nutanix AHV cluster:* Make sure that the latest version of NGT is installed on the virtual machine. For details on how to do this, see Nutanix documentation.
- *For virtual machines on a Nutanix ESXi cluster:* Make sure that the latest versions of VMware Tools and NGT are installed on the virtual machine. For details on how to do this, see Nutanix and VMware documentation.

- *For virtual machines in a vSphere environment:* Make sure that the latest version of VMware Tools is installed on the virtual machine. For details on how to do this, see VMware documentation.
- *For Linux virtual machines:* If a virtual machine on a Nutanix ESXi cluster or in a vSphere environment does not boot, change the controller type from SCSI to SATA, and then install the necessary SCSI drivers to switch back to SCSI.
- *For Windows virtual machines:* Reactivate the Windows licenses.
- *Only if you migrated virtual machines without network connectivity.* Make sure to configure the network settings on the virtual machine.
- Enable protection of the migrated data. For details on how to do this, see “[Protecting virtual machines](#)” on page 128 and “[Protecting applications](#)” on page 209.

Performing disaster recovery of data to Google Cloud

You can perform disaster recovery of data from the on-premises environment to Google Cloud in the event of a disaster.

Prerequisites

- You have a Google Account with the following permissions:
 - To access Google Cloud Storage buckets in the Google Cloud project where you want to deploy your new HYCU backup controller.
 - To deploy Google Compute Engine virtual machine instances to the Google Cloud project where you want to deploy your new HYCU backup controller.
 - To set up a firewall rule in the Google Cloud network where you plan to deploy your new HYCU backup controller.
- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected and must have the DR-ready status. For more information, see “[SpinUp specifics](#)” on page 142.

Considerations

- When the HYCU backup controller is deployed in Google Cloud, changing network settings is prevented in HYCU.

- Make sure the imported target is in the region to which you plan to migrate your virtual machines. This ensures the disaster recovery process is as fast and as cost-effective as possible.
- After you deploy the HYCU backup controller and use it to perform disaster recovery, you can keep the HYCU backup controller to stay prepared for disaster recovery in the future. However, every time you upgrade HYCU, you must deploy a new HYCU backup controller to be able to perform disaster recovery of data to cloud.

Procedure

1. Deploy a HYCU backup controller by using the HYCU R-Cloud or HYCU for Google Cloud web user interface. For details on how to do this, see HYCU R-Cloud or HYCU for Google Cloud documentation.
2. In Google Cloud, in the VPC network pane, in the Firewall rules context, create a new firewall rule to allow ingress network traffic through the TCP port 8443 from the entire subnetwork which the HYCU backup controller belongs to. For details on how to do this, see Google Cloud documentation.
3. Sign in to the HYCU web user interface by specifying the following URL:

```
https://<IPAddress>:8443
```

In this instance, *<IPAddress>* is the external IP address of the newly deployed HYCU backup controller.

4. Add a Google Cloud service account with permissions to access the Google Cloud Storage buckets where backup data of the protected virtual machines is stored. For details on how to do this, see [“Adding a Google Cloud service account” on page 367](#).
5. Import the Google Cloud target with your backup data:
 - a. In the Targets panel, click  **Import**. The Import Target dialog box opens.
 - b. In the Bucket Name field, enter the name as it was specified in the original target configuration.
 - c. From the Cloud Account drop-down list, select an imported Google Cloud service account, and then click **Next**.
 - d. Click the target name to confirm your selection, and then click **Next**.
 - e. In the Multiple Targets dialog box, one or more targets that store backup data are displayed. If any additional targets are found, select them one

by one and specify the values so that they match the original target configuration. For each target, click **Validate** to check the configuration.

- f. After you validated all the targets required for your restore, click **Import**.
6. Migrate your virtual machines or applications to cloud. For instructions, see “[Migrating data to cloud](#)” on page 497.

Protecting data across on-premises and Azure environments

You can use the SpinUp functionality to migrate protected data across the on-premises and Azure environments. In the event of a disaster in the on-premises environment, it provides disaster recovery of data to Azure.

Prerequisites

- You have an active subscription for HYCU for Azure. For details, see HYCU for Azure documentation.
- An Azure service principal is added to HYCU. For instructions, see “[Adding an Azure service principal](#)” on page 369.
- You must own a HYCU R-Cloud license. For details, see “[Licensing](#)” on page 385.
- A storage account that is dedicated exclusively to migration operations must be created in Azure. This storage account must be in the same region and resource group as the virtual machine that you plan to migrate, must have public network access enabled, and its type must be Standard general-purpose v2 or Premium block blobs.

Depending on what you want to do, see one of the following:

I want to...	Instructions
Migrate protected data across the on-premises and Azure environments.	“ Migrating virtual machines across different environments ” on the next page
Perform disaster recovery of data to Azure.	“ Performing disaster recovery of data to Azure ” on page 513

Migrating virtual machines across different environments

You can migrate protected data across the on-premises and Azure environments:

- “Migrating data to cloud” below
- “Migrating data from cloud” on page 510

Migrating data to cloud

You can migrate virtual machines, servers, and applications running on virtual machines and servers to Azure by using the SpinUp functionality. Keep in mind that when you migrate an application, the whole virtual machine or server on which this application is running is migrated to cloud.

 **Note** The instructions for protecting virtual machine data apply also to servers except where specifically stated otherwise.

Prerequisite

The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected, and must have a successful platform readiness check during the backup. For more information, see “SpinUp specifics” on page 142.

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for migrating data.
- *For Nutanix clusters:* You cannot migrate volume groups.
- *For vSphere environments:*
 - You cannot migrate virtual machine templates.
 - Migrating data from snapshots is not supported.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for migrating data.

- After you migrate data to cloud, an Azure temporary disk is automatically assigned to the migrated virtual machine. This disk is not a managed disk and it is used only for short-term data storage.
- *For virtual machines with secure boot enabled:* Because Azure does not currently support the secure boot feature for virtual machines, after you migrate such a virtual machine to cloud, secure boot cannot be enabled for it.

Depending on whether you want to migrate virtual machine or application data to cloud, access one of the following panels:

- Accessing the Virtual Machines panel
To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.
- Accessing the Applications panel
To access the Applications panel, in the navigation pane, click  **Applications**.

Procedure

1. In the Virtual Machines or Applications panel, select the entity that you want to migrate.
2. In the Detail view that appears at the bottom of the screen, select the virtual machine or application restore point that you want to use for the migration.
 **Note** The Detail view appears only if you click an entity. Selecting the check box before the name of the entity will not open the Detail view.
3. Click  **SpinUp to Cloud**.
4. Select **SpinUp to Azure**, and then click **Next**.
5. From the Service principal drop-down menu, select the service principal that has access to the required resources.
6. From the Subscription drop-down menu, select the appropriate subscription for the migrated virtual machine.
7. From the Resource group drop-down menu, select the resource group for the migrated virtual machine.
8. From the Location drop-down menu, select the geographic region for the migrated virtual machine.

9. From the Availability zone drop-down menu, select the zone for the migrated virtual machine.

Note The selected geographic region and the size of the virtual machine determine to which zones you can migrate data. If you do not want to migrate data to any zone, select **None**.
10. From the Storage account drop-down menu, select the storage account that is dedicated exclusively to migration operations.
11. Click **Next**. The VM Settings dialog box opens.
12. From the SpinUp from drop-down menu, select which tier you want to use for the migration. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: Ensures the fastest migration of data to cloud.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
13. In the New VM name field, enter a name for the migrated virtual machine.
14. In the vCPU cores field, enter the number of virtual CPUs to be assigned to the migrated virtual machine multiplied by the number of cores per virtual CPU.
15. In the Memory field, enter the amount of memory (in GiB) to be assigned to the migrated virtual machine.
16. From the Virtual machine type drop-down menu, select the virtual machine type.

Note The list shows virtual machine types that match the specified number of virtual CPUs and amount of memory, and the boot type of the virtual machine you are migrating to cloud (BIOS or UEFI). If no virtual machine type exactly corresponds to the specified values, the closest matches are shown.
17. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the migrated virtual machine.
18. Under Network interfaces, you can view the network interface that will be added to the migrated virtual machine. By default, this is the first network interface from the subscription that you selected for the migrated virtual

machine. If required, you can also modify network settings.

Modifying network settings

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

Note When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow these steps:
 - Only if you are adding a network interface.* From the Network dropdown menu, select the network for the network interface.
 - Note** The list of available networks includes only the ones within the region you selected for the migrated virtual machine.
 - Select the subnet to which the network interface should be assigned.
 - In the Public IP address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the migrated virtual machine.
Dynamic	A dynamic IP address will be assigned to the network interface on the migrated virtual machine.
Static	A static IP address will be assigned to the network interface on the migrated virtual machine.
Existing	A preferred public IP address resource that you have created in Azure Government will be assigned to the network interface on the migrated virtual machine.

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

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the migrated virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the migrated virtual machine.

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

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

19. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:

- **Linux**
- **Windows**

20. Click **SpinUp**.

The Migration to cloud job starts. When it finishes successfully, you can view the migrated virtual machine in the Virtual Machines panel in HYCU for Azure. For details, see HYCU for Azure documentation.

After migrating data to cloud

- *For Windows virtual machines:* Reactivate the Windows licenses.
- *For Linux virtual machines:* Install the Linux Integration Services for Hyper-V and Azure on the virtual machine. For details, see Microsoft documentation.
- Enable protection of the migrated virtual machines by using HYCU for Azure. For details on how to do this, see HYCU for Azure documentation.

Migrating data from cloud

You can migrate virtual machines from Azure by using the SpinUp functionality.

Limitations

- Migrating virtual machines with unmanaged disks is not supported.
- *For Nutanix clusters:* You can migrate Azure Generation 2 virtual machines only to clusters that support UEFI virtual machines.

Consideration

After you migrate data from cloud, the migrated virtual machine does not contain the temporary disk that was automatically assigned to it in Azure.

Accessing the Virtual Machines panel

To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.

Procedure

1. In the Virtual Machines panel, click  **SpinUp from Cloud**.
2. Select **SpinUp from Azure**, and then click **Next**.
3. From the Service principal drop-down menu, select the service principal that has access to the required resources.
4. From the Subscription drop-down menu, select the HYCU for Azure subscription to which the virtual machine that you want to migrate belongs.
5. From the Resource group drop-down menu, select the resource group to which the virtual machine that you want to migrate belongs.
6. From the Virtual machine drop-down menu, select the virtual machine that you want to migrate.
7. From the Checkpoint drop-down menu, select the checkpoint from which you want to migrate virtual machine data.
8. From the Storage account drop-down menu, select the storage account that is dedicated exclusively to migration operations.
9. Click **Next**. The VM Settings dialog box opens.
10. From the Storage container drop-down menu, select where you want to migrate the virtual machine.
11. In the New VM name field, enter a name for the migrated virtual machine.
12. *Only if the virtual machine that you are migrating was created in the on-premises environment, migrated to cloud, and now you are migrating it back to the on-premises environment.* If you want the virtual machine to have the same virtual machine settings as it had in the on-premises environment, enable

the **Keep original on-premises settings** option, and then continue with step 15.

Otherwise, leave the Keep original on-premises settings option disabled and continue with the next step.

13. Specify the following values for the migrated virtual machine:
 - The number of virtual CPUs.
 - The number of cores to be assigned to each virtual CPU.
 - The amount of memory (in GiB).

 **Note** The default values are the ones that the virtual machine had in the environment in which it was created, either in the on-premises or cloud one.

14. Under Network adapters, depending on your data protection needs, do one of the following:
 - Add one or more network adapters:
 - a. Click **Add Network Adapter**. The Network dialog box opens.
 - b. From the Network drop-down menu, select the virtual network for the network adapter.
 - c. Click **Add**.
 - Edit any of the existing network adapters to connect the virtual machine to a different network. To do so, select a network adapter, click  **Edit**, and make the required modification.
 - Delete any of the existing network adapters by selecting it, and then clicking  **Delete**. If you delete all the existing network adapters, your virtual machine will be migrated without network connectivity.
15. Use the **Power virtual machine on** switch if you want to turn the migrated virtual machine on after the migration.
16. Click **SpinUp**.

The Migration from cloud job starts. When it finishes successfully, you can view the migrated virtual machine in the Virtual Machines panel.

After migrating data from cloud

- *For virtual machines on a Nutanix AHV cluster:* Make sure that the latest version of NGT is installed on the virtual machine. For details, see Nutanix documentation.

- *For virtual machines on a Nutanix ESXi cluster:* Make sure that the latest versions of VMware Tools and NGT are installed on the virtual machine. For details, see Nutanix and VMware documentation.
- *For virtual machines in a vSphere environment:* Make sure that the latest version of VMware Tools is installed on the virtual machine. For details, see VMware documentation.
- *For Windows virtual machines:* Reactivate the Windows licenses.
- *For Linux virtual machines:* If a virtual machine on a Nutanix ESXi cluster or in a vSphere environment does not boot, change the disk controller from SCSI to IDE, and then install the latest version of VMware Tools on the virtual machine. You can later set the disk controller back to SCSI.
- *Only if you migrated virtual machines without network connectivity.* Make sure to configure the network settings on the virtual machine.
- Enable protection of the migrated data. For details, see “[Protecting virtual machines](#)” on page 128 and “[Protecting applications](#)” on page 209.

Performing disaster recovery of data to Azure

You can perform disaster recovery of data from the on-premises environment to Azure in the event of a disaster.

Prerequisite

The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected and must have the DR-ready status. For more information, see “[SpinUp specifics](#)” on page 142.

Considerations

- When the HYCU backup controller is deployed in Azure, changing network settings is prevented in HYCU.
- Make sure the imported target is in the region to which you plan to migrate your virtual machines. This ensures the disaster recovery process is as fast and as cost-effective as possible.
- After you deploy the HYCU backup controller and use it to perform disaster recovery, you can keep the HYCU backup controller to stay prepared for disaster recovery in the future. For instructions on how to upgrade the HYCU backup controller when a new software release version is available, contact [HYCU Support](#).

Procedure

1. Deploy a HYCU backup controller by using the HYCU for Azure web user interface. For details on how to do this, see HYCU for Azure documentation.
2. In Azure, create a new firewall rule to allow ingress network traffic on TCP port 8443 from the entire subnetwork to which the HYCU backup controller belongs. For details, see Azure documentation.
3. Sign in to the HYCU web user interface by specifying the following URL:

```
https://<IPAddress>:8443
```

In this instance, `<IPAddress>` is the external IP address of the newly deployed HYCU backup controller.

⚠ Important The credentials you provided in Azure during virtual machine creation cannot be used to sign in to HYCU and perform disaster recovery of data to Azure. For details on what credentials you can use to sign in to HYCU or to access the HYCU backup controller by using SSH, see “[Signing in to HYCU](#)” on page 49 or “[Accessing the HYCU backup controller virtual machine by using SSH](#)” on page 443.

4. Import the Azure target on which your backup data is stored to HYCU:
 - a. In the Targets panel, click  **Import**. The Import Target dialog box opens.
 - b. From the Type drop-down menu, select **AZURE**.
 - c. In the Storage account name field, enter the Azure storage account name as it was specified in the original target configuration.
 - d. In the Secret access key field, enter the secret access key for your Azure account.
 - e. In the Storage container name, enter the name of the storage container that is associated with the target and where the backup data is stored.
 - f. Click **Next**. The Import Backup Catalog dialog box opens.
 - g. Select the HYCU backup controller whose backup data you want to import, and then click **Next**.
 - h. In the Multiple Targets dialog box, do one of the following:
 - *If backup data is stored on one target:*
Click **Import**.
 - *If backup data is stored on more than one target:*

- Select each target one by one and specify the values so that they match the original target configuration.
- For each target, click **Validate** to check the configuration.
- Click **Import**.

5. Migrate your virtual machines or applications to cloud. For instructions, see “[Migrating data to cloud](#)” on page 506.

Protecting data across on-premises and Azure Government environments

You can use the SpinUp functionality to migrate protected data from your on-premises environment to Azure Government. In the event of a disaster in the on-premises environment, it provides disaster recovery of data to Azure Government.

Prerequisites

- An Azure Government service principal must be added to HYCU. For instructions, see “[Adding an Azure Government service principal](#)” on page 371.
- You must own a HYCU R-Cloud license. For details, see “[Licensing](#)” on page 385.

Depending on what you want to do, see one of the following:

I want to...	Instructions
Migrate protected data from the on-premises environment to Azure Government.	“ Migrating virtual machines to cloud ” below
Perform disaster recovery of data to Azure Government.	“ Performing disaster recovery of data to Azure Government ” on page 520

Migrating virtual machines to cloud

You can migrate virtual machines, servers, and applications running on virtual machines and servers to Azure Government by using the SpinUp functionality.

Keep in mind that when you migrate an application, the whole virtual machine or the server on which this application is running is migrated to cloud.

 **Note** The instructions for protecting virtual machine data apply also to servers except where specifically stated otherwise.

Prerequisite

The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected, and must have a successful platform readiness check during the backup. For more information, see [“SpinUp specifics” on page 142](#).

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for migrating data.
- *For Nutanix clusters:* You cannot migrate volume groups.
- *For vSphere environments:*
 - You cannot migrate virtual machine templates.
 - Migrating data from snapshots is not supported.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for migrating data.
- After you migrate data to cloud, an Azure temporary disk is automatically assigned to the migrated virtual machine. This disk is not a managed disk and it is used only for short-term data storage.
- *For virtual machines with secure boot enabled:* Because Azure does not currently support the secure boot feature for virtual machines, after you migrate such a virtual machine to cloud, secure boot cannot be enabled for it.

Depending on whether you want to migrate virtual machine or application data to cloud, access one of the following panels:

- Accessing the Virtual Machines panel
To access the Virtual Machines panel, in the navigation pane, click  **Virtual Machines**.
- Accessing the Applications panel
To access the Applications panel, in the navigation pane, click  **Applications**.

Procedure

1. In the Virtual Machines or Applications panel, select the entity that you want to migrate.
2. In the Detail view that appears at the bottom of the screen, select the virtual machine or application restore point that you want to use for the migration.

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

3. Click  **SpinUp to Cloud**.
4. Select **SpinUp to Azure Government**, and then click **Next**.
5. From the Service principal drop-down menu, select the service principal that has access to the required resources.
6. From the Subscription drop-down menu, select the appropriate subscription for the migrated virtual machine.
7. From the Resource group drop-down menu, select the resource group for the migrated virtual machine.
8. From the Location drop-down menu, select the geographic region for the migrated virtual machine.
9. From the Availability zone drop-down menu, select the zone for the migrated virtual machine.

 **Note** The selected geographic region and the size of the virtual machine determine to which zones you can migrate data. If you do not want to migrate data to any zone, select **None**.

10. Click **Next**. The VM Settings dialog box opens.
11. From the SpinUp from drop-down menu, select which tier you want to use for the migration. Your restore point can contain one or more tiers among which you can select:

- **Automatic:** Ensures the fastest migration of data to cloud.
- **Backup**
- **Copy**
- **Archive**
- **Snapshot**

12. In the New VM name field, enter a name for the migrated virtual machine.
13. In the vCPU cores field, enter the number of virtual CPUs to be assigned to the migrated virtual machine multiplied by the number of cores per virtual CPU.
14. In the Memory field, enter the amount of memory (in GiB) to be assigned to the migrated virtual machine.
15. From the Virtual machine type drop-down menu, select the virtual machine type.

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

16. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the migrated virtual machine.
17. Under Network interfaces, you can view the network interface that will be added to the migrated virtual machine. By default, this is the first network interface from the subscription that you selected for the migrated virtual machine. If required, you can also modify network settings.

Modifying network settings

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

Note When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:

- Only if you are adding a network interface.* From the Network drop-down menu, select the network for the network interface.

 **Note** The list of available networks includes only the ones within the region you selected for the migrated virtual machine.

- Select the subnet to which the network interface should be assigned.
- In the Public IP address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the migrated virtual machine.
Dynamic	A dynamic IP address will be assigned to the network interface on the migrated virtual machine.
Static	A static IP address will be assigned to the network interface on the migrated virtual machine.
Existing	A preferred public IP address resource that you have created in Azure Government will be assigned to the network interface on the migrated virtual machine.

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

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the migrated virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the migrated virtual machine.

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

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

18. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
 - **Linux**
 - **Windows**
19. Click **SpinUp**.

The Migration to cloud job starts.

After migrating data to cloud

- *For Windows virtual machines:* Reactivate the Windows licenses.
- *For Linux virtual machines:* Install the Linux Integration Services for Hyper-V and Azure on the virtual machine. For details, see Microsoft documentation.

Performing disaster recovery of data to Azure Government

You can perform disaster recovery of data from the on-premises environment to Azure Government in the event of a disaster.

Prerequisites

- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected and must have the DR-ready status. For more information, see “[SpinUp specifics](#)” on [page 142](#).
- You have the HYCU virtual appliance image for Azure Government. To obtain the image and further instructions, contact [HYCU Support](#).

Considerations

- When the HYCU backup controller is deployed in Azure Government, changing network settings is prevented in HYCU.
- Make sure the imported target is in the region to which you plan to migrate your virtual machines. This ensures the disaster recovery process is as fast and as cost-effective as possible.

Procedure

1. Deploy a HYCU backup controller:
 - a. In Azure Government, create a managed image from the HYCU virtual appliance image.
 - b. Create a virtual machine from the managed image. Make sure the virtual machine is configured with a public IP address and an additional disk of 32 GiB in size.

For details, see Azure documentation.
2. In Azure Government, create a new firewall rule to allow ingress network traffic on TCP port 8443 from the entire subnetwork to which the HYCU backup controller belongs. For details, see Azure documentation.
3. Sign in to the HYCU web user interface by specifying the following URL:

```
https://<IPAddress>:8443
```

In this instance, <IPAddress> is the external IP address of the newly deployed HYCU backup controller.

① Important The credentials you provided in Azure Government during virtual machine creation cannot be used to sign in to HYCU and perform disaster recovery of data to Azure Government. For details on what credentials you can use to sign in to HYCU or to access the HYCU backup controller by using SSH, see “[Signing in to HYCU](#)” on page 49 or “[Accessing the HYCU backup controller virtual machine by using SSH](#)” on page 443.

4. Import the Azure Government target on which your backup data is stored to HYCU:
 - a. In the Targets panel, click  **Import**. The Import Target dialog box opens.
 - b. From the Type drop-down menu, select **AZURE Government**.
 - c. In the Storage account name field, enter the Azure Government storage account name as it was specified in the original target configuration.
 - d. In the Secret access key field, enter the secret access key for your Azure Government account.
 - e. In the Storage container name, enter the name of the storage container that is associated with the target and where the backup data is stored.
 - f. Click **Next**. The Import Backup Catalog dialog box opens.

- g. Select the HYCU backup controller whose backup data you want to import, and then click **Next**.
- h. In the Multiple Targets dialog box, do one of the following:
 - *If backup data is stored on one target:*
Click **Import**.
 - *If backup data is stored on more than one target:*
 - Select each target one by one and specify the values so that they match the original target configuration.
 - For each target, click **Validate** to check the configuration.
 - Click **Import**.

5. Migrate your virtual machines or applications to cloud. For instructions, see “[Migrating virtual machines to cloud](#)” on page 515.

Appendix A

Customizing HYCU configuration settings

You can find all HYCU configuration settings in the `config.properties.template` file in the `/opt/grizzly` folder on your HYCU backup controller. This file contains a list of all available configuration settings and their default values. If you want to adjust any of these configuration settings to meet your specific data protection environment needs and provide optimal performance, create a new `config.properties` file in the `/hycudata/opt/grizzly` folder, and then specify the preferred configuration settings and their new values.

 **Note** When you upgrade HYCU, the `config.properties` file will be kept. However, you may want to check the updated `config.properties.template` file for new configuration settings that you can use with the new HYCU version.

Depending on which configuration settings you want to customize, see one of the following sections:

- “Snapshot settings” on page 525
- “Utilization threshold settings” on page 525
- “Display settings” on page 526
- “SQL Server application settings” on page 526
- “Settings for aborting jobs” on page 526
- “File server settings” on page 527
- “Data rehydration settings” on page 528
- “Disaster recovery settings” on page 528
- “User management settings” on page 529

Procedure

1. Open a remote session to the HYCU backup controller virtual machine:

```
ssh hycu@<HYCUBackupControllerIPAddress>
```

When requested, enter the password for the `hycu` user.

For detailed information about accessing the HYCU backup controller virtual machine by using SSH, see [“Accessing the HYCU backup controller virtual machine by using SSH” on page 443](#).

2. Access and open the `config.properties` file by using one of the following text editors:

- Vim:

```
sudo vi /hycudata/opt/grizzly/config.properties
```

- Nano:

```
sudo nano /hycudata/opt/grizzly/config.properties
```

3. Edit any of the existing configuration settings as required.

4. Save and exit the `config.properties` file.

Changes to the configuration settings are applied based on their `ReloadClass` annotation in the `config.properties.template` file:

Annotation	Description
Job	The changes are applied when a new job is started.
Mount	The changes are applied in the following scenarios: <ul style="list-style-type: none"> • When any new target is added to HYCU. • When an existing NFS, SMB, Nutanix, iSCSI, or tape target is activated again after being deactivated with the Detach storage option enabled.
Operation	The changes are applied when a new operation that does not create a job is executed (for example, when using the HYCU web user interface, REST API, SSH, or WinRM).
Service	The changes are applied when the HYCU application server (the Grizzly server) is restarted.

If a configuration setting has no annotation, it is recommended to restart the HYCU application server (the Grizzly server). To do so, run the following command:

```
sudo service grizzly restart
```

Snapshot settings

You can use the following settings to configure the snapshot retention threshold at which an event is triggered:

Settings/descriptions

max.snapshots.per.vm

If the number of snapshots that are retained per virtual machine exceeds the specified value, a warning event is triggered. The default value is 24.

max.snapshots.per.cluster

If the number of snapshots that are retained per Nutanix cluster exceeds the specified value, a warning event is triggered. The default value is 2400.

Utilization threshold settings

You can use the following settings to configure the system and data disks as well as target utilization thresholds:

Settings/descriptions

controller.disk.full.warning.threshold.fraction

If the HYCU backup controller utilization of the system or data disk exceeds the specified value, an event is triggered. The default value is 0.90.

target.utilization.threshold.red.fraction

If the HYCU backup controller utilization of the target exceeds the specified value, its health status indicator becomes red. The default value is 0.95.

target.utilization.threshold.yellow.fraction
--

If the HYCU backup controller utilization of the target exceeds the specified value, its health status indicator becomes yellow. The default value is 0.90.

For detailed information about the health status of the target, see “[Viewing target information](#)” on page 325.

Display settings

You can use the following setting to customize the maximum number of displayed items:

Setting/description

<code>items.per.directory.in.flr</code>

Maximum number of files that are displayed for each directory when restoring individual files. The default value is 1000.

SQL Server application settings

You can use the following setting to customize the backup of SQL Server applications:

Setting/description

<code>sql.translog.compress</code>

During the backup of an SQL Server application, transaction log compression is enabled by default (the default value is <code>true</code>). If you want to disable it, make sure to set the value for this setting to <code>false</code> .

Settings for aborting jobs

You can use the following settings to configure when a job that has the Executing status will be aborted automatically:

Settings/descriptions

<code>jobs.abort.deadline.minutes</code>
--

Time (in minutes) within which a job must be completed. The default value is 1440.
--

<code>jobs.abort.interval.minutes</code>
--

Time interval (in minutes) at which all jobs that have the Executing status are retrieved and stopped if they have been in this status longer than specified in the <code>jobs.abort.deadline.minutes</code> setting. The default value is 15.
--

File server settings

You can use the following settings to configure file share backups:

Settings/descriptions

`afs.reindex.interval.count`

Number of incremental file share backups after which a full reindex is performed, which increases the responsiveness of the file restore process. The default value is 5.

`afs.partial.success.threshold.count`

and

`afs.partial.success.max.fail.fraction`

Maximum number and fraction of failed file backups up to which the backup status of the corresponding file share is marked as Completed with errors (and not as Failed). The default values are the following:

- *For the maximum number of failed file backups*
(`afs.partial.success.threshold.count`): 10000
- *For the maximum fraction of failed file backups*
(`afs.partial.success.max.fail.fraction`): 0.01

! Important Both these values must be exceeded for the backup status of the file share to be marked as Failed.

`afs.instance.afs.cluster.priority`

HYCU uses an internal algorithm to distribute the load among multiple HYCU instances. It prioritizes the HYCU instances that are running on the same Nutanix cluster as the file server and the HYCU instances that are running on the same Nutanix cluster as the HYCU backup controller. It also takes into account the number of jobs that are already running on each HYCU instance.

Raising the value of this setting gives higher priority to the HYCU instances that are running on the same Nutanix cluster as the file server.

`afs.instance.bc.cluster.priority`

HYCU uses an internal algorithm to distribute the load among multiple HYCU instances. It prioritizes the HYCU instances that are running on the same Nutanix cluster as the file server and the HYCU instances that are running on the same Nutanix cluster as the HYCU backup controller. It also takes into account the number of jobs that are already running on each HYCU instance.

Settings/descriptions

Raising the value of this setting gives higher priority to the HYCU instances that are running on the same Nutanix cluster as the HYCU backup controller.

Data rehydration settings

You can use the following settings to configure HYCU to perform data rehydration:

Settings/descriptions

`target.azure.blob.rehydration.enable`

HYCU is preconfigured to perform data rehydration before performing the restore if backup data or a copy of backup data is stored in the Azure archive storage tier. During a rehydration task, the data is moved from the archive storage tier to the hot storage tier from which HYCU can restore data. HYCU does not move data back to the archive storage tier afterward. The default value is `true`.

`target.azure.blob.rehydration.threads`

Number of blobs that can be rehydrated in parallel. The default value is 20.

Disaster recovery settings

You can use the following settings to enable additional scenarios for disaster recovery or adjust automatic target synchronization:

Settings/descriptions

`clone.enabled.for.hycu.dr`

HYCU is preconfigured to prevent creating clones of the HYCU backup controller (the virtual machine itself or its virtual disks).

⚠ Caution Do not activate a clone of the HYCU backup controller while the original HYCU backup controller is still active. If such activation happens, data loss may occur. All currently running backups fail and their status is set to Error. The corresponding restore points are then automatically removed by the HYCU cleaning process.

If set to `true`, cloning of the HYCU backup controller is enabled and the

Settings/descriptions

respective restore options become available in the HYCU web user interface.

`synchronize.target.catalog.interval.minutes`

When the recovery HYCU backup controller is in recovery mode, automatic target synchronization is by default performed every 60 minutes. Setting the value to 0 disables automatic target synchronization.

User management settings

You can use the following setting to completely prevent deleting protected data when changing ownership of virtual machines and file shares:

Setting/description

`force.keep.backups.on.owner.change`

If set to `true` (the default value is `false`), data protected by specific owners is never deleted—even if the option to delete such data is specified when changing ownership of virtual machines and file shares in any of the HYCU interfaces.

Appendix B

After restoring a virtual machine to a different source

A virtual machine can be restored to a different source by using the Clone VM restore option as described in [“Cloning a virtual machine” on page 171](#).

However, depending on your virtual machine original environment and target environment, you might have to perform some additional steps after the restore:

VM original environment	VM target environment	Additional steps
Nutanix ESXi, vSphere, AWS GovCloud (US), or Azure Government, or servers	Nutanix AHV	See “After restoring a virtual machine to a Nutanix AHV cluster” on the next page .
vSphere	Nutanix ESXi	See “After restoring a virtual machine to a Nutanix ESXi cluster” on page 532 .
Nutanix AHV, Nutanix ESXi, AWS GovCloud (US), Azure Government, or servers	vSphere	<ul style="list-style-type: none"><i>Only if restoring a virtual machine with more than one disk.</i> After the restore, additional disks will be offline. Make sure to bring them back online.<i>Only if restoring a Windows virtual machine from AWS GovCloud (US).</i> Make sure that the latest version of VMware Tools is installed on the virtual machine. For instructions, see VMware documentation.

VM original environment	VM target environment	Additional steps
		<ul style="list-style-type: none"> Only if the restored virtual machine has more than one disk. Check the hard drive boot order of the restored virtual machine. If it differs from the one on the original virtual machine, change the boot order in BIOS.

After restoring a virtual machine to a Nutanix AHV cluster

Considerations

- Only if restoring a virtual machine with more than one disk from a vSphere environment to a Nutanix AHV cluster. After the restore, additional disks will be offline. Make sure to bring them back online.
- If you have not followed the recommendations described in “[Preparing for the restore to a different source](#)” on page 135, your virtual machine will not boot after the restore, and you must perform the following additional steps:
 1. Make sure that the restored virtual machine is turned off.
 2. As the administrator or the root user, sign in to the Nutanix AHV cluster by using SSH.
 3. List the virtual machine details:

```
acli vm.get <VMName>
```

4. Take a note of the current bus and index values in the `disk_list` section.
5. Clone the existing disk to a new disk on the compatible bus:

```
acli vm.disk_create <VMName> bus=<BusType>
clone_from_vmdisk=vm:<VMName>:<CurrentBus>.<CurrentIndex>
```

In this instance, `<VMName>` is the name of the restored virtual machine, `<BusType>` is `scsi`, `ide`, or `sata`, `<CurrentBus>` is the bus value from

the `disk_list` section, and `<CurrentIndex>` is the index value from the `disk_list` section.

If the original virtual machine has the SATA or SCSI disks, clone them to the SATA disks. For example:

```
acli vm.disk_create test-vm bus=sata
clone_from_vmdisk=vm:test-vm:scsi.0
```

If the original virtual machine has the IDE disks, clone them to the IDE disks. For example:

```
acli vm.disk_create test-vm bus=ide
clone_from_vmdisk=vm:test-vm:ide.0
```

After you perform the previous procedure for all the disks, follow these steps:

1. Sign in to the Nutanix Prism web console.
2. In the menu bar, click **Home**, and then select **VM**.
3. Click the **Table** tab to display the VM Table view.
4. From the list of virtual machines, select the restored virtual machine, and click **Update**.
5. Delete the source disks, and then select the boot disk and click **Save**.
6. Click **Power on** to turn on the restored virtual machine.
7. Install the Nutanix Guest Tools software bundle of the latest version on the virtual machine.
8. *Recommended for virtual machines that had the SCSI disks.* Clone the controller back to the SCSI controller.

For details on how to update a virtual machine on a Nutanix cluster, see Nutanix documentation.

After restoring a virtual machine to a Nutanix ESXi cluster

If after restoring a virtual machine from a vSphere environment to a Nutanix ESXi cluster the virtual machine does not start, you must perform additional steps.

 **Note** You can use either the vSphere Web Client or the vSphere Client as the interface for performing the steps. As an example, you are guided through the steps that you must perform if you are using the vSphere Web Client.

Steps

- If the type of controller on the restored virtual machine is not the same as it was on the original virtual machine, do the following:
 1. Sign in to the vSphere Web Client.
 2. Click the **VMs** tab, and then right-click the restored virtual machine and select **Edit Settings**.
 3. On the Virtual Hardware tab, modify the controller settings so that they match the ones on the original virtual machine.
- If the virtual machine uses UEFI firmware, you may need to select the boot file manually. In this case, do the following:
 1. Sign in to the vSphere Web Client.
 2. Access the EFI Boot Manager menu, and then do the following:
 - a. Select the **Enter setup** option.
 - b. Enter the boot maintenance manager by selecting **Boot option maintenance menu**.
 - c. Use the **Boot from a File** option to browse for a boot file.
 - d. Find a device whose name contains the GPT string that represents the boot partition, and then press **Enter** to open it.
 - e. Navigate to the EFI boot file that you can find at the following location:
 - Windows: \EFI\Microsoft\Boot\bootmgrfw.efi
 - Linux: /EFI/<OSName>/grubx64.efi
 - f. Press **Enter** to resume booting.

Provide feedback

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

info@hycu.com

We will be glad to hear from you!

