

USER GUIDE

HYCU Data Protection for Enterprise Clouds

Version: 4.2.0

Product release date: December 2020

Document release date: December 2020



Legal notices

Copyright notice

© 2020 HYCU. All rights reserved.

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

Trademarks

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

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

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

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.

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 vSphere®, VMware vSphere® Data Protection™, 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	12
HYCU key features and benefits	13
Data protection environment overview	14
HYCU data protection	15
2 Deploying the HYCU virtual appliance	16
Sizing resources for your HYCU backup infrastructure	17
Adjusting firewall configuration	18
Deploying HYCU to a Nutanix AHV cluster	22
Deployment tasks	22
Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment	26
Logging on to HYCU	28
3 Establishing a data protection environment	30
Adding sources	31
Adding a Nutanix cluster	31
Adding a vCenter Server	33
Adding a Nutanix Files server	34
Adding a physical machine	37
Setting up targets	38
Setting up an NFS target	38
Setting up an SMB target	40
Setting up a Nutanix target	42
Setting up an iSCSI target	44
Setting up an AWS S3/Compatible target	46
Setting up an Azure target	48
Setting up a Google Cloud target	50
Setting up a tape target	52
Defining your backup strategy	54
Taking advantage of predefined policies	55

Creating a custom policy	55
Setting a default policy	65
4 Protecting virtual machines	66
Planning virtual machine protection	67
Preparing your data protection environment	67
Preparing for disaster recovery	69
Physical machine specifics	71
HYCU Protégé specifics	72
Enabling access to data	74
Setting up virtual machine backup configuration options	77
Backing up virtual machines	79
Restoring virtual machines	80
Restore options	81
Restoring a virtual machine	82
Cloning a virtual machine	85
Restoring virtual disks	90
Cloning virtual disks	91
Exporting virtual disks	93
Restoring individual files	97
5 Protecting applications	102
Enabling access to application data	102
Planning application protection	105
Backing up applications	110
Restoring whole applications	111
Restore options	112
Restoring a virtual machine	112
Cloning a virtual machine	116
Restoring SQL Server databases	119
Restoring Exchange Server databases, mailboxes, and public folders	122
Restoring Oracle database instances and tablespaces	125

6 Protecting file shares	128
Backing up file shares	128
Restoring file share data	129
7 Recovering your data protection environment	133
Preparing for disaster recovery	133
Deploying a temporary HYCU backup controller	134
Importing targets	135
Performing disaster recovery	136
Restoring the HYCU backup controller to the original source	137
Restoring the HYCU backup controller to a different source	138
Recreating HYCU instances	140
8 Performing daily tasks	142
Using the HYCU dashboard	143
Checking the status of jobs	144
Viewing events	146
Sending email notifications	146
Using HYCU reports	147
Getting started with reporting	148
Viewing reports	149
Generating reports	150
Scheduling reports	151
Exporting and importing reports	152
Viewing entity details	152
Viewing the backup status of entities	154
Filtering data	155
Applying the main filter	155
Applying the detail filter	156
Filtering options in the Applications panel	156
Filtering options in the Virtual Machines panel	157
Filtering options in the Volume Groups panel	158

Filtering options in the Shares panel	159
Filtering options in the Policies panel	159
Filtering options in the Targets panel	159
Filtering options in the Jobs panel	160
Filtering options in the Events panel	160
Filtering options in the Self-Service panel	161
Exporting the contents of the panel	161
Managing targets	161
Viewing target information	162
Editing a target	163
Activating or deactivating a target	164
Increasing the size of an iSCSI target	164
Deleting a target	165
Managing policies	165
Viewing policy information	165
Editing a policy	166
Deleting a policy	167
Performing a manual backup	167
Expiring backups	168
Expiring backups automatically	168
Expiring backups manually	168
Archiving data manually	170
Recreating snapshots	171
Adjusting the HYCU virtual machine resources	171
9 Managing users	173
HYCU groups	173
User roles	174
Setting up a user environment	176
Creating a user	176
Adding a user to a group	178
Creating a self-service group	178

Setting ownership of virtual machines	179
Setting ownership of file shares	180
Activating or deactivating users or self-service groups	180
Switching to another group	181
Updating your user profile	182
10 Administering	183
Configuring Active Directory authentication	184
Enabling certificate authentication	185
Adding a cloud account	185
Adding a Google Cloud Platform service account	186
Adding an Azure service principal	188
Configuring target encryption	189
Exporting an encryption key	189
Importing an encryption key	189
Managing HYCU instances	189
Creating a HYCU instance by using the HYCU web user interface	190
Viewing HYCU instance information	191
Deleting a HYCU instance	191
Setting the iSCSI Initiator secret	191
Licensing	192
Creating a license request	193
Requesting and retrieving licenses	194
Activating licenses	194
Setting up logging	195
Configuring your network	197
Changing network settings	197
Limiting network bandwidth	198
Setting power options	200
Configuring an SMTP server	200
Configuring SSL certificates	201
Creating a self-signed certificate	202

Importing a custom certificate	202
Sharing telemetry data with HYCU	203
Upgrading HYCU	204
Upgrading HYCU on a Nutanix AHV cluster	206
Upgrading HYCU on a Nutanix ESXi cluster	208
Upgrading HYCU in a vSphere environment	212
Applying HYCU hotfixes	216
Applying a hotfix by using the HYCU web user interface	217
Applying a hotfix by using the shell script	219
Removing HYCU	220
11 Tuning your data protection environment	223
Accessing the HYCU backup controller virtual machine by using SSH	224
Enabling HTTPS for WinRM connections	225
Configuring FIPS-compliant mode for HYCU	226
Enabling FIPS-compliant mode for HYCU	227
Disabling FIPS-compliant mode for HYCU	227
Setting up LDAPS authentication	228
Securing SMTP connections	228
Importing an SSL certificate for the STARTTLS security mode	228
Importing an SSL certificate for the SSL/TLS security mode	229
Setting up HYCU to use multiple networks	229
Setting up HYCU to use multiple networks on a Nutanix AHV or ESXi cluster	230
Setting up HYCU to use multiple networks in a vSphere environment	231
Increasing the size of the HYCU virtual disks	232
Increasing the size of the HYCU disks in a Nutanix AHV cluster	232
Increasing the size of the HYCU disks in a Nutanix ESXi cluster or vSphere environment	232
Assigning privileges to a vSphere user	233
Using the HYCU REST API Explorer	235
Using the command-line interface	235
Using the pre and post scripts	236

12 Monitoring data protection environments	237
Using the HYCU Manager console	237
Monitoring your HYCU controllers	238
Adding a HYCU controller	239
Viewing information about HYCU controllers	240
Viewing events	241
Performing administration tasks	241
Managing users	242
13 Employing Nutanix Mine with HYCU	245
Registering HYCU with Nutanix Prism	245
Accessing HYCU from the Nutanix Prism web console	246
Viewing the Nutanix Mine with HYCU dashboard	247
14 HYCU Protégé	249
Protecting data across on-premises and Google Cloud Platform environments	249
Migrating virtual machines across different environments	250
Performing disaster recovery of data to Google Cloud Platform	254
Protecting data across on-premises and Azure environments	256
Migrating virtual machines across different environments	256
Performing disaster recovery of data to Azure	262
A Customizing HYCU configuration settings	264
Snapshot settings	265
Utilization threshold settings	266
Display settings	266
SQL Server application settings	266
Settings for aborting jobs	267
HTTPS for WinRM configuration settings	267
Nutanix Files settings	267
Data rehydration settings	268
HYCU backup controller restore settings	269
User management settings	269

B Restoring to an environment with a different hypervisor	270
Restoring a virtual machine from a Nutanix ESXi cluster or a vSphere environment to a Nutanix AHV cluster	271
Restoring a virtual machine from a vSphere environment to a Nutanix ESXi cluster	274
Restoring a virtual machine from a Nutanix AHV cluster or a Nutanix ESXi cluster to a vSphere environment	275

Chapter 1

About HYCU

HYCU Data Protection for Enterprise Clouds (HYCU) is a high performing backup and recovery solution for Nutanix, VMware, and physical machine environments. It is the first data protection solution that is fully integrated with Nutanix, making data protection easy to deploy and simple to use.




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 the Nutanix Prism web console (for Nutanix AHV clusters) or the vSphere (Web) Client (for Nutanix ESXi clusters and vSphere environments).

- **Provides new-found visibility**

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

- **Protects data in a few minutes**

Data protection of virtual machines, physical machines, applications, file shares, 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.

- **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.

- **Offers a scalable backup for Nutanix Files**

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.

- **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.

- **Provides business continuity of your data protection environment across different infrastructures**

HYCU Protégé ensures data resilience by using the SpinUp functionality to migrate protected data between the on-premises and cloud infrastructures (Google Cloud Platform or Azure). In the event of a disaster, HYCU Protégé provides disaster recovery of mission-critical data to cloud.

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.
Sources	Environments for which HYCU provides data protection—Nutanix clusters, vSphere environments, Nutanix Files servers, and physical machines.
Entities	Objects to which you can assign a policy and for which you therefore provide data protection—virtual and physical machines, applications, and file shares. 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:

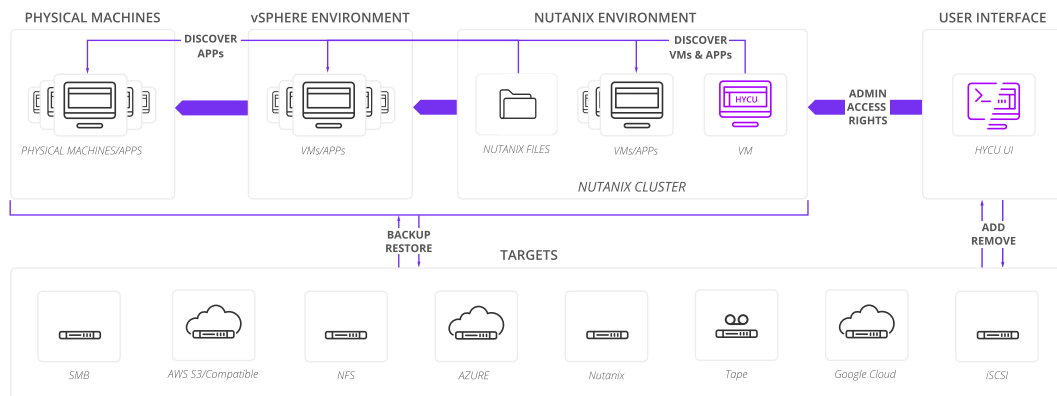


Figure 1-2: HYCU architecture

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 and physical machines, applications running on them, file shares on Nutanix Files servers, volume groups, and 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 and physical machines, 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 66](#)
- [“Protecting applications” on page 102](#)
- [“Protecting file shares” on page 128](#)

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, or a vSphere environment for which you want to provide data protection.

Deployment modes

Mode	Select this mode if you want to...
HYCU Backup Controller	Protect virtual machines (including virtual machine templates), physical machines, applications, file shares, and volume groups. 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	Protect file shares. A HYCU instance is a virtual machine that HYCU uses to perform data protection operations for Nutanix Files, taking the load off the HYCU backup controller.
HYCU Manager	Manage HYCU controllers. 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</i>	"Adjusting firewall configuration" on page 18

Task	Instructions
<i>on your network</i> . Open relevant ports in each involved firewall.	
3. Deploy the HYCU virtual appliance to a source.	“Deploying HYCU to a Nutanix AHV cluster” on page 22 or “Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment” on page 26

After you successfully deploy the HYCU virtual appliance, you can access HYCU by using a supported web browser. For details on how to log on to HYCU, see [“Logging on to HYCU” on page 28](#).

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:

Size of environment	System requirements		
	Storage	CPU cores	Memory
Small (100–200 VMs)	20–40 GiB	6	8 GiB
Medium (200–500 VMs)	100–200 GiB	8	32 GiB

Size of environment	System requirements		
	Storage	CPU cores	Memory
Large (more than 500 VMs)	The figures vary. Contact HYCU Customer 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.

- *For deploying in the HYCU Backup Controller mode:* Targets:

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 Nutanix, vSphere, or HYCU communication endpoints. For HYCU to operate properly, you must adjust the firewall rules and open the ports listed in the table that follows.

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.

Purpose	Communication endpoints		Ports at destination	Protocols
	Source	Destination		
Use of the HYCU interface	System where HYCU interface is accessed	HYCU backup controller	8443	TCP
Access to the HYCU backup controller by using SSH	System where HYCU interface is accessed	HYCU backup controller	22	TCP
Use of a DNS server	HYCU backup controller, HYCU instance	DNS server	53	TCP UDP
Use of an NTP server	HYCU backup controller, HYCU instance	NTP server	123	UDP

Purpose	Communication endpoints		Ports at destination	Protocols
	Source	Destination		
Discovery of VMs running Linux and applications on them	HYCU backup controller	VMs	22 ^a	TCP
	VMs	HYCU backup controller	8443	
Discovery of VMs running Windows and applications on them	HYCU backup controller	VMs	5985 5986	TCP
	VMs	HYCU backup controller	8443	
Backup	HYCU backup controller	Nutanix Controller VMs	3205 3260	TCP
Backup and restore of file shares	HYCU instance	Nutanix Files server	445 ^b 2049 ^c 9440	TCP
	HYCU backup controller	HYCU instance	8443	
	HYCU instance	HYCU backup controller		
Backup of data to an NFS v4 target	HYCU backup controller, HYCU instance	NFS v4 server	2049	TCP UDP
Backup of data to an NFS v3 target	HYCU backup controller, HYCU instance	NFS v3 server	111 2049 mountd port ^d	TCP UDP
Backup of data to an SMB target	HYCU backup controller, HYCU instance	SMB server	445	TCP
Backup of data to an iSCSI target	HYCU backup controller	iSCSI server	3260	TCP
Backup of data to a cloud target	HYCU backup controller, HYCU instance	Cloud server	443	TCP
Archive of data to a	HYCU backup	QStar server	111	TCP

Purpose	Communication endpoints		Ports at destination	Protocols
	Source	Destination		
QStar NFS target	controller, HYCU instance		2049 mountd port ^d 18082 ^e	
Archive of data to a QStar SMB target	HYCU backup controller, HYCU instance	QStar server	445 18082 ^e	TCP
Restore from backups created with the Fast Restore policy option enabled	HYCU backup controller	Nutanix Controller VMs	3205	TCP
Restore of applications or individual files	System where HYCU interface is accessed	HYCU backup controller	445	TCP
Restore of files from a snapshot to a Windows virtual machine	VMs	Nutanix Controller VMs	860 3260	TCP
Restore of files from a target to a Windows virtual machine	VMs	HYCU backup controller	139 445	TCP
Restore of files to a Linux virtual machine	HYCU backup controller	VMs	22	TCP
Restore of files to an SMB share	HYCU backup controller	System with an SMB share	445	TCP
Restore of files to an NFS share	HYCU backup controller	System with an NFS share	NFS4: 2049 NFS3: 111, mountd port ^d	TCP
Restore of files to the local machine	System where the HYCU interface is accessed	HYCU backup controller	8443	TCP
Restore of	HYCU backup controller	VMs	860 3205	TCP

Purpose	Communication endpoints		Ports at destination	Protocols
	Source	Destination		
applications or individual files from backups on an iSCSI target that were created with the Fast Restore policy option enabled	VMs	HYCU backup controller	3260	
		Cluster virtual server (cluster virtual IP address) ^f		
		iSCSI target discovery portal (iSCSI Data Services IP address) ^g		
		Nutanix Controller VMs ^g		
Data protection of entities in a Nutanix AHV or Nutanix ESXi cluster or on a Nutanix Files server ^h	HYCU backup controller	Cluster virtual server (cluster virtual IP address)	9440	TCP
		Nutanix Controller VMs		
Data protection of virtual machines in a Nutanix cluster or volume groups ⁱ	HYCU backup controller	Cluster virtual server (cluster virtual IP address) ^f	3205	TCP
		iSCSI target discovery portal (iSCSI Data Services IP address) ^g	3260	
Backup of entities in a vSphere environment	HYCU backup controller	ESXi hosts	902	TCP
		vCenter Server	443	
Sharing telemetry data with HYCU	HYCU backup controller	Telemetry endpoint: callhome.hycu.com	443	TCP
		Amazon S3 AWS endpoint: s3.eu-central-1.amazonaws.com		
Use of an LDAP server	HYCU backup controller	LDAP server	LDAP: 389 LDAPS: 636	TCP
Use of an SMTP server for sending	HYCU backup controller	SMTP server	25 ^j	TCP

Purpose	Communication endpoints		Ports at destination	Protocols
	Source	Destination		
email notifications				

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

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

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

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

^e 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.

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

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

^h HYCU uses the Nutanix REST API v3.

ⁱ HYCU accesses Nutanix Volumes.

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

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 17](#).

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” on the next page
2. Create a virtual machine for HYCU deployment.	“Creating a virtual machine for HYCU deployment on a Nutanix AHV cluster” on page 24
3. Configure HYCU on the created virtual	“Configuring HYCU on the virtual machine”

Task	Instructions
machine.	on page 25

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. Log on 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-4.2.0-3634`


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. Set the time zone as required.
 - 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 disks in a Nutanix AHV cluster” on page 232](#).
 - 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 disks in a Nutanix AHV cluster” on page 232](#).
 - iv. Click **Add**.
4. In the Network Adapters (NIC) section, click **Add New NIC**, and then, in the Create NIC

dialog box, do the following:

- a. From the VLAN Name drop-down menu, select a VLAN.
 - b. 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 16](#).

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 189](#).
 - 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 prebuilt Instant-on license. This license expires automatically after 45 days and cannot be reused. Therefore, make sure to obtain a permanent license within this 45-day period. For instructions, see [“Licensing” on page 192](#).

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.


 **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 17](#).
- *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.

Procedure

1. Log on 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	<p>Browse your file system for the HYCU OVF package.</p> <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**.

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 16](#).
 - 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 189](#).
 - 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 prebuilt Instant-on license. This license expires automatically after 45 days and cannot be reused. Therefore, make sure to obtain a permanent license within this 45-day period. For instructions, see [“Licensing” on page 192](#).

Logging on 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 logon page, enter your logon name and password. You can use the default user name and password for initial access to HYCU:

User name: **admin**

Password: **admin**

For security purposes, it is highly recommended that you change the default password.

 **Note** Keep in mind that the level of access depends on your user permissions. For details, see [“Managing users” on page 173](#).

After you log on 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 235](#).

Chapter 3

Establishing a data protection environment

After you deploy the HYCU virtual appliance and log on 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 38](#)

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 56](#).


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 69](#).

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 or in vSphere environments, file shares on Nutanix Files servers, or physical machines and applications running on physical machines. For instructions on how to add a specific source, see one of the following sections:

- [“Adding a Nutanix cluster” below](#)
- [“Adding a vCenter Server” on page 33](#)
- [“Adding a Nutanix Files server” on page 34](#)
- [“Adding a physical machine” on page 37](#)

 **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, each of which hosts a series of virtual machines running applications. Adding one or more Nutanix clusters to HYCU is the first step to protecting your virtual machine data.

Prerequisites

- *For Nutanix ESXi clusters:* Your cluster is registered to the vCenter Server through the Prism web console. For details on how to do this, see Nutanix documentation.

- *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 is 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 64](#).

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 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.

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	<p>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 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 64.</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 it 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 VC 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.
- You can delete any Nutanix cluster, provided it does not have any dependencies. Therefore, deleting a Nutanix cluster that is specified as the central site cluster in your policy is not possible 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, 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.

Prerequisite

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

[page 233](#).

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 include in the backup are running.

3. Enter the user name and password of a user with specific privileges for vCenter Servers.
4. Click **Save**.

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

Adding a Nutanix Files server

HYCU enables you to protect SMB and NFS file shares on Nutanix Files servers. You can add one or more Nutanix Files servers that host file shares you want to include in the backup.

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 Nutanix Files data protection operations, taking the load off the HYCU backup controller. You can have one or more HYCU instances in your data protection environment, depending on your business needs. For details on HYCU instances, see [“HYCU instances” on page 36](#).

Prerequisite

HYCU can access a Nutanix Files server. For details, 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 **Nutanix Files** tab, and then click **+ New**.
2. Enter the name of the Nutanix Files server in the following format:

`https://<NutanixFilesServerHostname>:<Port>`

Entering the port is optional if the default value is used, 9440.

 **Important** Make sure the name of the Nutanix Files server is unique.

3. In the Nutanix Files Server Credentials section, enter the user name and password of a user with Nutanix Files server administrative rights.


The default user name and password are the following:

User name: **admin**

Password: **Nutanix/4u**



If you created a new user for accessing the REST API, you can use this user's credentials instead.

4. Use the **Use SMB protocol for accessing shares** switch if you plan to protect SMB file shares, and then, in the SMB Backup Credentials section, enter the user name and password of a server or backup administrator with access to all SMB file shares within the Nutanix Files server.

 **Note** You cannot assign credentials to each share individually.

5. Use the **Use NFSv4 protocol for accessing shares** switch if you plan to protect NFS file shares.
6. Click **Save**.


You can later do the following:

- Edit any of the existing Nutanix Files servers. To do so, select the Nutanix Files server, click  **Edit**, make the required modifications, and then click **Save**.
- Delete the Nutanix Files servers that you do not need anymore as follows:
 - a. Select the Nutanix Files server, and then click  **Delete**.
 - b. In the Remove Nutanix Files dialog box that opens, do the following:
 - If you want to delete also the corresponding HYCU instances, use the **Remove unused HYCU instances** switch.
 - If you want to delete snapshots created by HYCU, use the **Delete snapshots** switch.
 - c. Click **Yes**.

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. You can create a new user

(recommended) or use the default administrator to access the REST API.

 **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, do the following:

- To create a new user to access the REST API, follow these steps:
 1. Establish a connection to the Nutanix cluster:

```
ssh @<NutanixClusterHostname>
```

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

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

You can later use the newly created user and password as Nutanix Files server credentials when adding a Nutanix Files server to HYCU.

- If you are using the default administrator to access the REST API, follow these steps:
 1. Establish a connection to the Nutanix Files server:

```
ssh <NutanixFilesServerHostname>
```

When requested, enter the following credentials:

User name: **nutanix**

Password: **Nutanix/4u**

2. Run the following command to reset the password for all Nutanix Files server nodes to the default one (that is, Nutanix/4u):

```
allssh reset_admin_password.py
```

You can later set this password to a different one by running the following commands:

```
allssh "sudo truncate -s 0 /etc/security/opasswd"
```

```
allssh "sudo faillock --user admin --reset"
```

```
allssh "echo <NewPassword> | sudo passwd --stdin admin"
```

HYCU instances

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.

You can have one or more HYCU instances on your Nutanix cluster. 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. When distributing the load among multiple HYCU instances, HYCU automatically prioritizes the HYCU instances that are running on the same Nutanix cluster as the Nutanix Files server and the HYCU backup controller. 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 264](#).

You can create a HYCU instance in one of the following ways:

- By deploying the HYCU virtual appliance and selecting the HYCU Instance mode. For details, see [“Deploying the HYCU virtual appliance” on page 16](#).
- By using the HYCU web user interface. For details, see [“Creating a HYCU instance by using the HYCU web user interface” on page 190](#).


Considerations

- You can create a HYCU instance before or after adding a Nutanix Files 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 64 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 desired values. For details on how to customize the HYCU configuration settings, see [“Customizing HYCU configuration settings” on page 264](#).
- 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.

If you later decide to remove any HYCU instance from your data protection environment, you can do it as described in [“Deleting a HYCU instance” on page 191](#).

Adding a physical machine

Adding one or more physical machines to HYCU is as the first step to protecting your physical machine data.



 **Important** Protection of Linux physical machines is a preview feature and is not intended for use in a production environment.


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 **Physical Machines** tab, and then click **+ New**.
2. Enter the name of the physical machine.
3. Enter the host name or IP address of the physical machine.
4. Click **Save**.

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

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

Setting up targets

Targets are locations where the protected data is stored. HYCU allows you to store your data to an NFS share, an SMB share, an iSCSI storage device, a Nutanix cluster, AWS S3 or S3-compatible storage, Azure storage, Google Cloud Platform storage, and tape.

 **Note** A Nutanix Files share can be used as an NFS or SMB target. If you plan to use the Nutanix Files share only as a target and not as a source, there is no need to add the Nutanix Files server to HYCU.

The approach to set 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:

- [“Setting up an NFS target” below](#)
- [“Setting up an SMB target” on page 40](#)
- [“Setting up a Nutanix target” on page 42](#)
- [“Setting up an iSCSI target” on page 44](#)
- [“Setting up an AWS S3/Compatible target” on page 46](#)
- [“Setting up an Azure target” on page 48](#)
- [“Setting up a Google Cloud target” on page 50](#)
- [“Setting up a tape target” on page 52](#)

Setting up an NFS target

Prerequisites


- The service is configured and accessible for the HYCU backup controller and the HYCU instances.
- There is enough free space on the target for storing the data.

- If deduplication is enabled on the target, the target is 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) are 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 physical machine data:* The target is accessible from the physical machine.

Limitations

- Target compression is not supported for file shares.
- *For protecting physical machine data:*
 - You can store only Linux physical machine backups to this type of target.
 - Target encryption and compression are not supported.

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. In the General section, do the following:
 - a. Enter the name of a target and, optionally, its description.
 - b. *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.
 - c. 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.
 - d. Use the **Use for archiving** switch if you want this target to be reserved for data archives.

 **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.

 **Important** The target that you use for archiving data cannot be used for


backing up data or storing copies of backup data.

- e. 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.

3. In the Target section, do the following:

- a. From the Type drop-down menu, select **NFS**.
- b. 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).
- c. 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, and file shares, export the encryption key to a file and keep this file safe. For instructions, see [“Exporting an encryption key” on page 189](#).

4. Click **Save**.

The target is added to the list of targets. For details on managing targets, see [“Managing targets” on page 161](#).

Setting up an SMB target


Prerequisites

- The service is configured and accessible for the HYCU backup controller and the HYCU instances.
- There is enough free space on the target for storing the data.
- If deduplication is enabled on the target, the target is 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 is used. For a list of supported SMB versions, see the *HYCU Compatibility Matrix*.
- *For protecting physical machine data:* The target is accessible from the physical machine.

Limitations

- Target compression is not supported for file shares.
- *For protecting physical machine data:*
 - You can store only Windows physical machine backups to this type of target.
 - Target encryption and compression are not supported.


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. In the General section, do the following:
 - a. Enter the name of a target and, optionally, its description.
 - b. *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.
 - c. 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.
 - d. Use 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.
 - e. 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.

3. In the Target section, do the following:

- a. From the Type drop-down menu, select **SMB**.
- b. *Optional.* Enter the domain and user credentials.
- c. 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).
- d. 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, and file shares, export the encryption key to a file and keep this file safe. For instructions, see [“Exporting an encryption key” on page 189](#).

4. Click **Save**.

The target is added to the list of targets. For details on managing targets, see [“Managing targets” on page 161](#).

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 physical machine backups to 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. In the General section, do the following:
 - a. Enter the name of a target and, optionally, its description.
 - b. *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.
 - c. 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.
 - d. Use 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.
 - e. 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.
3. In the Target section, do the following:

- a. From the Type drop-down menu, select **Nutanix**.
 - b. Enter the name of the Nutanix cluster in the following URL format:
`https://<ServerName>:<Port>`
 - c. 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.
- d. 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.

- e. 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 and applications, export the encryption key to a file and keep this file on safe. For instructions, see ["Exporting an encryption key" on page 189](#).

4. Click **Save**.

The target is added to the list of targets. For details on managing targets, see ["Managing targets" on page 161](#).

Setting up an iSCSI target

Prerequisites

- The service is configured and accessible.
- The target has not been initialized yet.
- The HYCU iSCSI Initiator secret is 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 physical machine backups to 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 Customer 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. In the General section, do the following:
 - a. Enter the name of a target and, optionally, its description.
 - b. *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.
 - c. 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.
 - d. Use 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.
 - e. 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.
3. In the Target section, do the following:
 - a. From the Type drop-down menu, select **iSCSI**.
 - b. 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.

- c. 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 and applications, export the encryption key to a file and keep this file safe. For instructions, see [“Exporting an encryption key” on page 189](#).

4. If the iSCSI server requires CHAP authentication, in the CHAP section, do the following:
 - a. Use the switch to turn the CHAP authentication option on, 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 191](#).


5. Click **Save**.

The target is added to the list of targets. For details on managing targets, see [“Managing targets” on page 161](#).

Setting up an AWS S3/Compatible target

Prerequisites

- The service is configured and accessible.
- The S3 bucket is 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*.

 **Important** Using Object Lock (WORM) is supported for AWS S3 and Nutanix Objects targets. In this case, make sure that versioning is set to Enabled when creating an S3 bucket for HYCU. For details, see AWS or Nutanix documentation.

- The following minimum required AWS S3 permissions are specified: `s3:GetObject`, `s3:DeleteObject`, `s3:PutObject`, `s3:ListBucket`, `s3:GetBucketAcl`, `s3:ListBucketMultipartUploads`, `s3:GetBucketLocation`, `s3:GetBucketObjectLockConfiguration`, `s3:DeleteObjectVersion`, `s3:ListBucketVersions`, and `s3:GetBucketVersioning`.
- *For S3-compatible targets:* If you want to provide secure HTTPS access, make sure the required CA-signed certificate is imported as follows:
 1. Open a remote session to the HYCU backup controller virtual machine:

```
ssh hycu@<HYCUBackupControllerIPAddress>
```

When requested, enter the default password.

2. Import the required CA-signed certificate:

```
keytool -importcert -keystore /etc/pki/ca-  
trust/extracted/java/cacerts  
-file <CertificatePathname>
```

```
keytool -importcert -keystore /etc/pki/cert-  
templates/cacerts.template  
-file <CertificatePathname>
```


- *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

- The size of files created on this type of target is limited to 5 TiB. Therefore, make sure the size of backup data, copies of backup data, and data archives does not exceed this limit.
- HYCU does not support the AWS S3 targets that use the Glacier storage class.
- HYCU currently supports only AWS S3 Signature Version 4.
- Target compression is not supported for file shares.
- Storing physical machine backups to this type of target is not supported.

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. In the General section, do the following:
 - a. Enter the name of a target and, optionally, its description.
 - b. In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
 - c. 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.
 - d. Use 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.

- e. 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.

3. In the Target section, do the following:

- a. From the Type drop-down menu, select **AWS S3/Compatible** or **AWS Government**.
- b. Enter the service endpoint URL, the bucket name, the access key ID, and the secret access key. The access key and the secret access key are used to authenticate Amazon API service calls.
- c. 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`).
- d. 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, and file shares, export the encryption key to a file and keep this file safe. For instructions, see [“Exporting an encryption key” on page 189](#).

4. Click **Save**.

The target is added to the list of targets. For details on managing targets, see [“Managing targets” on page 161](#).

Setting up an Azure target

Prerequisite

The service is configured and accessible.

Limitations

- The size of files created on this type of target is limited to 5 TiB. Therefore, make sure the size of backup data, copies of backup data, and data archives does not exceed this


limit.

- Target compression is not supported for file shares.
- Storing physical machine backups to this type of target is not supported.

Consideration

Your data on the Azure target can be stored in the hot, cool, and archive storage tiers. To ensure the data stored in the archive storage tier is also restored, HYCU is configured to perform data rehydration before performing a restore. For details, see [“Data rehydration settings” on page 268](#).


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. In the General section, do the following:
 - a. Enter the name of a target and, optionally, its description.
 - b. In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
 - c. 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.
 - d. Use 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.

For details on how HYCU manages archiving data to Azure, see [“Archiving data to the Azure archive storage tier” on page 64](#).


- e. 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.


3. In the Target section, do the following:

a. From the Type drop-down menu, select **AZURE**, **AZURE Government**, or **AZURE China**.

b. Enter the storage account name, the secret access key, and the container name.

 **Note** If the container does not exist, it is created automatically.

c. 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, and file shares, export the encryption key to a file and keep this file safe. For instructions, see [“Exporting an encryption key” on page 189](#).

4. Click **Save**.

The target is added to the list of targets. For details on managing targets, see [“Managing targets” on page 161](#).

Setting up a Google Cloud target

Prerequisites

- A Google Cloud Platform service account is created and then added to HYCU. For instructions on how to add a cloud account to HYCU, see [“Adding a Google Cloud Platform service account” on page 186](#).
- A Google Cloud Storage bucket is created in the project that is linked to the created Google Cloud Platform service account you added to HYCU.
- The service is configured and accessible.


Limitations

- The size of files created on this type of target is limited to 5 TiB. Therefore, make sure the size of backup data, copies of backup data, and data archives does not exceed this limit.
- Target compression is not supported for file shares.
- Storing physical machine backups to this type of target is not supported.

Consideration

To ensure your data is stored most cost-efficiently, HYCU stores backup data and copies of backup data in the Google Cloud Platform storage class that is optimal for the retention period set in your policy. Therefore, such data can be stored in a different storage class than the one set as the bucket's default storage class. However, this does not apply to the standard storage class. If a bucket's default storage class is set to standard, backup data and copies of backup data are always stored in the standard storage class.


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. In the General section, do the following:
 - a. Enter the name of a target and, optionally, its description.
 - b. In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
 - c. 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.
 - d. Use 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.

For details on how HYCU manages archiving data to Google Cloud Platform, see [“Archiving data to the Google Cloud Platform archive storage class” on page 64](#).
 - e. 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.
3. In the Target section, do the following:
 - a. From the Type drop-down menu, select **Google Cloud**.
 - b. 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 Platform service account you added to HYCU.
 - c. From the Cloud account drop-down menu, select the Google Cloud Platform service account you added to HYCU.
 - d. 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, and file shares, export the encryption key to a file and keep this file safe. For instructions, see [“Exporting an encryption key” on page 189](#).

4. Click **Save**.

The target is added to the list of targets. For details on managing targets, see [“Managing targets” on page 161](#).

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 is sufficient for storing archive data.
- The QStar cache is large enough.
- There is enough free space for storing archive data on QStar.

For details, see QStar documentation.

Limitation


Target compression is not supported—archive data cannot be compressed before it is stored on the target.

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. In the General section, do the following:
 - a. Enter the name of a target and, optionally, its description.
 - b. *Optional.* In the Size field, enter the maximum space that should be reserved for archive data (in MiB, GiB, or TiB).
 - c. 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.

- d. Make sure the **Use for archiving** option is enabled.
 - e. Make sure the **Enable Compression** option is disabled.
3. In the Target section, from the Type drop-down menu, select one of the following tape targets and follow the instructions:

Target type	Instructions
QStar NFS	<ol style="list-style-type: none"> a. Provide user credentials that HYCU will use to access the shared folder and make web service calls. b. Enter the name of the Integral Volume set where you want to archive data. c. 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: <code>https://<QStarServer>:<Port></code> d. <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. e. Use the Target encryption switch if you want the data stored on this target to be encrypted. <div style="border-left: 2px solid #007bff; padding-left: 10px; margin-top: 10px;"> <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, and file shares, export the encryption key to a file and keep this file safe. For instructions, see “Exporting an encryption key” on page 189. </div>
QStar SMB	<ol style="list-style-type: none"> a. <i>Optional.</i> Specify the domain in which the account that has access permissions on the shared folder is registered. b. Provide user credentials that HYCU will use to access the shared folder and make web service calls. c. Enter the name of the Integral Volume set where you want to archive data. 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

Target type	Instructions
	<p>will be used to access the QStar server in the following format: <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, and file shares, export the encryption key to a file and keep this file safe. For instructions, see “Exporting an encryption key” on page 189.

4. Click **Save**.

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

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” below](#).

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 65](#).
- You can set up the automatic assignment of policies to virtual machines. For details, see [“Setting up automatic policy assignment” on page 64](#).

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 desired 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
Backup window	Allows you to start all backup jobs within specific time frames to improve effectiveness and avoid overload of your environment.

Policy option	Description
Copy	Allows you to create a copy of backup data.
Archiving	Allows you to preserve your data for future reference.
Fast restore ^a	<p><i>Available only for Nutanix clusters.</i> Allows you to restore virtual machine or application data to the original storage container in a fast way by keeping local snapshots on the Nutanix cluster for the specified retention time.</p> <p>With this option enabled, HYCU will keep more than one snapshot on the Nutanix cluster, depending on your retention settings. This will allow you to restore virtual machine or application data in a fast way, reducing downtime.</p>
Backup from replica	<p><i>Available only for Nutanix clusters.</i> Allows you to back up your virtual machines from their replicas in remote office/branch office (ROBO) environments.</p> <p> Important Make sure that the schedule interval you set for the Nutanix protection domains that include the virtual machines you want to protect is less than or equal to the RPO set in the HYCU policy.</p> <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 through the Nutanix Prism web console, see Nutanix documentation.</p>
Auto-assignment	Allows you to set up the automatic assignment of policies to virtual machines. You do this by first applying categories or custom attributes to virtual machines in Nutanix Prism or VMware vSphere and then specifying the corresponding metadata in HYCU policies.

^a For Nutanix ESXi clusters with AOS version 5.10: If besides HYCU you use a Nutanix protection domain to protect virtual machines, a fast restore is performed only if the required snapshot is available in the protection domain.


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 enable the Backup window policy option, make sure you have created a backup window. For details on how to do this, see [“Creating a backup window” on page 60](#).
- 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 62](#).
- If you plan to back up virtual machines from their replicas in ROBO environments, make sure these prerequisites are met:
 - A protection domain that includes the virtual machines you want to protect is created and the specified schedule interval is less than or equal to the RPO set in the HYCU policy. For details on protecting virtual machines through the Nutanix Prism web console, see Nutanix documentation.
 - Both the central site Nutanix cluster and the branch office site cluster are added to HYCU. For details, see [“Adding a Nutanix cluster” on page 31](#).
- If you plan to enable the Auto-assignment policy option, make sure you are familiar with all the prerequisites and considerations described in [“Setting up automatic policy assignment” on page 64](#).


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 appears.
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*)
- **Backup window**
- **Copy**
- **Archiving**
- **Fast restore**
- **Backup from replica**
- **Auto-assignment**

 **Important** The Backup from replica and Fast restore options are not available for vSphere virtual machines and applications.

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. 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, or days) for the

data. The retention period defines when a restore point will be expired. For details on expiring backups, see [“Expiring backups” on page 168](#).

 **Note** Only if you use Object Lock on AWS 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 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.

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



- e. 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.

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

Enabled option	Procedure
Backup window	To specify a backup window, in the Backup section, 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 60 .

Enabled option	Procedure
	<p>If you do not select a backup window, the Always option is shown, which means that your backups are allowed to run at any time.</p>
Copy	<p>To create a copy of backup data, in the Copy section, do the following:</p> <ol style="list-style-type: none"> Set a retention period (in months, weeks, or days) for the copy of backup data. From the 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>
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 62.</p>
Fast restore ^a	<p><i>Available only for Nutanix clusters.</i> To keep more than one snapshot on the Nutanix cluster, 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 Nutanix cluster.</p> <p> Note The snapshot retention period cannot be shorter than the RPO or longer than the backup retention period.</p>
Backup from replica ^a	<p><i>Available only for Nutanix clusters.</i> To back up virtual machines from their replicas, in the Backup from replica section, from the Central site cluster drop-down menu, select the cluster on which the replicas of your virtual machines reside.</p>
Auto-assignment	<p>To set up automatic policy assignment, in the Auto-assignment section, enter a metadata key and value, and then click Add.</p> <p>During the next virtual machine synchronization, the policy is automatically assigned to all the virtual machines that have the</p>

Enabled option	Procedure
	<p>corresponding category or custom attribute values applied in Nutanix Prism or VMware vSphere.</p> <p> Note 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  Synchronize in the Virtual Machines panel.</p>

^a For Nutanix ESXi clusters with AOS version 5.10: If besides HYCU you use a Nutanix protection domain to protect virtual machines, a fast restore is performed only if the required snapshot is available in the protection domain.


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 165](#).


Creating a backup window

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


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

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


Accessing the Policies panel

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


Procedure


1. In the Policies panel, click  **Backup Window**.
2. In the Backup Window dialog box that appears, click **+ New**. The New dialog box appears.
3. Enter a name for your backup window and, optionally, a 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** 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.

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**.
8. In the Backup Window dialog box, click **Close**.

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

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

- Specify a backup window when creating a new policy. For details, see [“Creating a policy” on page 56](#).
- 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 selected the Bronze policy 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 week days.

The screenshot shows the 'Backup Window - New' configuration window. The 'NAME' field is 'window_bronze' and the 'DESCRIPTION (OPTIONAL)' is 'window_bronze_desc'. The 'TIME ZONE' is set to 'Europe/Ljubljana (GMT+02:00)'. Below this, it states 'YOUR LOCAL TIMEZONE IS Europe/Budapest (GMT+02:00)' and 'CONTROLLER TIMEZONE IS Europe/Ljubljana (GMT+02:00)'. The 'SCHEDULE' section shows a 24-hour timeline for each day of the week. For Monday through Friday, there are two teal bars representing 'Incremental Only' backups: one from 18:00 to 24:00 and another from 00:00 to 06:00. For Saturday and Sunday, there is a single blue bar representing 'Full/Incremental' backups from 00:00 to 24:00. The 'TIME FRAMES' section at the bottom lists the specific time ranges for each day: MON (00:00-06:00), MON (18:00-24:00), TUE (00:00-06:00), TUE (18:00-24:00), WED (00:00-06:00), WED (18:00-24:00), THU (00:00-06:00), THU (18:00-24:00), FRI (00:00-06:00), FRI (18:00-24:00), SAT (00:00-24:00), and SUN (00:00-24:00). Buttons for 'Close', 'Back', and 'Save' are at the bottom right.

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 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 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 Platform archive storage class:* Data archives created with any of the previous versions of HYCU are not moved to the archive storage tier.

Accessing the Policies panel


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

Procedure


1. In the Policies panel, click  **Archiving**.
2. In the Archiving dialog box that appears, click  **New**.
3. In the New dialog box that appears, enter a name for your data archive and, optionally, a description.
4. Add any of the desired archiving options to the list of the enabled options by clicking it. The following options are available:

Weekly	Allows you to create a weekly archive of data.
Monthly	Allows you to create a monthly archive of data.
Yearly	Allows you to create a yearly archive of data.

5. Specify the hour and the minute when the archive job should begin running.



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

6. Provide information about when to archive data, the retention period to be used, and one or more archive targets.


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

For details on how HYCU manages archiving data to Azure or Google Cloud Platform, see [“Archiving data to the Azure archive storage tier” on the next page](#) or [“Archiving data to the Google Cloud Platform archive storage class” on the next page](#).

7. 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 56](#).
- 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 170](#).

Archiving data to the Azure archive storage tier

HYCU automatically moves each 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.

 **Important** 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.

Archiving data to the Google Cloud Platform archive storage class

HYCU automatically moves each data archive that has a retention period set to at least 365 days to the Google Cloud Platform archive storage class during the next archive synchronization. By moving data archives to the archive storage class, HYCU ensures your data is stored most cost-efficiently because the archive storage class is optimized for storing data that is not accessed frequently and is stored for at least 365 days.

Setting up automatic policy assignment

You can set up the automatic assignment of policies to virtual machines by first applying categories or custom attributes to virtual machines in Nutanix Prism or VMware vSphere and then specifying the matching metadata in HYCU policies. If 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.

Considerations

- If you want a predefined policy to be automatically assigned to a virtual machine, when specifying the values for the category or the custom attribute and the metadata, you can use the name of the policy (Gold, Silver, or Bronze). You can also use the Exclude value if you want the virtual machine to be excluded from a backup.
- Assigning policies automatically does not affect virtual machines to which you assigned policies by clicking the **Assign** or **Set Default** button.
- If the comparison of the custom attributes and metadata 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 custom attribute value is kept on the restored virtual machine only if the virtual machine is backed up with HYCU version 4.2.0 and the original custom attribute still exists in VMware vSphere.


Procedure

VMs reside...	You apply the following to VMs...
On a Nutanix ESXi cluster	Categories in the Nutanix Prism web console For instructions on how to apply a category to a virtual machine, see Nutanix documentation.
On a Nutanix ESXi cluster or in a vSphere environment	Custom attributes in the vSphere (Web) Client For instructions on how to apply a custom attribute to a virtual machine, see VMware documentation.


Setting a default policy


You can select one of the predefined or custom policies to be the default policy for your data protection environment. After you set a default policy, it is assigned to all existing entities that do not have an assigned policy yet, and to all newly discovered ones.

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 default, and then click  **Set Default**.
2. In the Set Default Policy dialog box that appears, do one of the following:
 - Click **Yes** if you want to assign the default policy to all entities that do not have an assigned policy (that is, existing and newly discovered ones).
 - Click **No** if you want to assign the default policy only to newly discovered entities.



If you later decide not to use this policy as the default one, click  **Clear Default**.

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 group (logically related virtual disks that are attached to a virtual machine).</p> <p> Tip After you back up virtual machines, you can view all protected volume groups attached to virtual machines and their details in the Protected Volume Groups panel. The details that you can view for volume groups are the same as for any entity. Therefore, the same information applies to volume groups as described for entities in “Viewing entity details” on page 152.</p> <p>Accessing the Protected Volume Groups panel</p> <p>To access the Protected Volume Groups panel, in the navigation pane, click  Volume Groups.</p>
vSphere environment	Virtual machine template (virtual machine that is used as a template to create other virtual machines).

The preparation steps and instructions for protecting virtual machines (including the HYCU backup controller) and physical machines 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 79](#)
- [“Restoring virtual machines” on page 80](#)
- [“Restoring individual files” on page 97](#)

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 69](#)
- [“Physical machine specifics” on page 71](#)
- [“HYCU Protégé specifics” on page 72](#)
- [“Enabling access to data” on page 74](#)
- [“Setting up virtual machine backup configuration options” on page 77](#)

Preparing your data protection environment

Prerequisites

- *For vSphere environments:* VMware Tools of the latest version is installed on virtual machines.
- *For Nutanix clusters:* When backing up virtual machines that have volume groups attached and you want the volume groups to be backed up as well, your AOS version is 5.10 or later. Otherwise, only the virtual machines will be backed up.
- *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 as well, 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.

Limitations

- Only a backup of local fixed disks and volume groups that are attached to virtual machines on the same Nutanix cluster directly or by using iSCSI 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.
- Backing up vSphere virtual machines that have NVMe controllers added is not supported.
- *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 ESXi clusters:* If you enabled the Backing up from replica policy option, backing up virtual machines that have disks on different containers is not supported.

- *For Nutanix AHV clusters:* With a Nutanix protection domain configured with NearSync, HYCU can create snapshots only if the AOS version is 5.10 or later.

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 vSphere environments:* 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.
- *For Nutanix clusters:* Archiving is performed from a snapshot if the snapshot is available on the original location (the cluster on which the original virtual machine is running or the central site Nutanix cluster if you are using the Backup from replica option). Otherwise, archiving is performed from the target.
- *For Nutanix ESXi clusters:* If the snapshot that HYCU used to perform a full backup is missing on a Nutanix cluster (for example, because the HYCU protection domain was deleted from Prism), 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 264](#).
- *For Nutanix ESXi clusters with AOS version 5.10 and 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.

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. Therefore, having the same volume group attached to the virtual machines inside the same Nutanix protection domain and at the same time to other virtual machines is not recommended.
- *For virtual machines in a ROBO environment and virtual machines that are part of a protection domain on Nutanix ESXi cluster with AOS version 5.10:* To ensure that applications on such 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.
- *Only if you plan to restore a virtual machine to an environment with a different hypervisor.* Follow these recommendations:
 - *For restoring a virtual machine from a Nutanix ESXi cluster or from a vSphere environment 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 [“Restoring a virtual machine from a Nutanix ESXi cluster or a vSphere environment to a Nutanix AHV cluster” on page 271](#)):
 - *For Windows virtual machines:* The Nutanix VirtIO package is installed on the virtual machine. If you have NGT installed on your virtual machine on the Nutanix ESXi cluster, there is no need to install the Nutanix VirtIO package because it is already installed as part of NGT installation.
 - *For Linux virtual machines on Nutanix ESXi clusters:* NGT is installed on your virtual machine.
 - *For Linux virtual machines in vSphere environments:* The VirtIO drivers are added to the guest OS kernel.
 - *For restoring a Linux virtual machine from a Nutanix AHV cluster to a vSphere environment:* Install Nutanix Guest Tools (NGT) on the virtual machine to ensure that it starts properly after the restore.

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.

 **Important** To further increase safety, we recommend 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.

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, and file shares 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
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)
AWS 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 Platform service account
QStar NFS	<ul style="list-style-type: none"> • User name • Password (if used)

	<ul style="list-style-type: none"> • Integral volume set name • Web service endpoint • 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)

Consideration

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.

Physical machine specifics

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



Important Protection of Linux physical machines is a preview feature and is not intended for use in a production environment.

Prerequisites

- Access to the file system data is enabled. For instructions, see [“Enabling access to data” on page 74](#).
- Sufficient disk space 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 physical machines:*
 - 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 physical machine to a Nutanix AHV cluster:* Make sure the Nutanix VirtIO package is installed on the physical machine before you back it up. For detailed information about installing Nutanix VirtIO, see Nutanix documentation.

- *For Linux physical machines:*
 - Access to the physical machine through SSH is enabled.
 - Sufficient space in the volume group is available for LVM snapshots. It is recommended that at least 20 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.
 - *For cloning a Linux physical machine:* 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"
```

Limitation

Protecting physical machines that use Virtual Data Optimizer (VDO) is not supported.

HYCU Protégé specifics

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

Prerequisites

- *For migration of virtual machines and applications to cloud:* Configure your environment to provide a successful cloud readiness check during the virtual and physical 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.
 - *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 following Linux utilities are available: sudo, iptables, ip, and lsinitrd.
- 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.
- The following drivers must be included in initramfs:
 - *Migration to Google Cloud Platform:* virtio drivers (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_net virtio_scsi"
```

- *Migration to Azure:* Hyper-V drivers (hv_vmbus, hv_storvsc, and hv_netvsc)
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 to Google Cloud Platform:* The Nutanix VirtIO package is installed on the virtual machines that you plan to migrate.

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

- *For migration of virtual machines from Azure to a Nutanix AHV cluster:* The Nutanix VirtIO package is installed on the virtual machines.
- *For disaster recovery to cloud:* Configure your environment to provide the Migration/DR-ready status for the virtual machines. A virtual machine has the Migration/DR-ready status if the following is true:
 - All backups in the current backup chain are stored on one of the cloud targets (Google Cloud or Azure).
 - A successful cloud readiness check is performed during the virtual machine backup.

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

Limitations

- Migrating virtual machines with different types of storage controllers is not supported.
- *For migration of virtual machines from Google Cloud Platform:* 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.

- Migrating physical machines that use UEFI firmware to cloud is not supported.

Recommendation

For migration of virtual machines and applications to cloud:

- *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 cloud readiness check status even if the network is not working.

Consideration

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
```

Enabling access to data

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

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

- You plan to protect physical machines.
- 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.
- You plan to use pre- and post-scripts.
- You plan to use HYCU Protégé 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 225](#).

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.

Consideration



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).

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>


Protocol options	Instructions
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"> In the Port field, enter the SSH server port number. From the Authentication type drop-down menu, select the type of authentication you want to be used and provide the required information: <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 <ol 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. <i>Only if the private key is encrypted.</i> 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"> From the Transport drop-down menu, select the type of transport you want to be used. In the Port field, enter the WinRM server port number. 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 or physical 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” below
Exclude selected disks or volume groups from data protection.	“Excluding disks from the backup” on the next page

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 236](#).

Prerequisites

- Access to the virtual machine file system is enabled. For instructions, see [“Enabling access to data” on page 74](#).
- 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 switches of your choice to specify the pre/post-snapshot and pre/post-backup scripts, and then enter the script path names. Enable one or more switches:
 - **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 disks from the backup

Prerequisite

You are an owner of the virtual machine whose disks you want to exclude from the backup. For instructions on how to set ownership of a virtual machine, see [“Setting ownership of virtual machines” on page 179](#).

Limitation


Only if you plan to restore individual files. If you exclude all virtual machine disks from the backup and leave only the attached volume groups, you will not be able to restore individual files.

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 virtual 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 152](#).
- *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.

Procedure


1. In the Virtual Machines panel, select the virtual machine whose disks and volume groups you want to exclude from data protection, and then select **Configuration**. The Configuration dialog box opens.
2. In the Exclude from backup tab, select the disks and volume groups that you want to exclude from the backup, and then click **Save**.

 **Important** For vSphere environments: If you plan to restore individual files, make sure not to exclude the operating system disk from the backup.

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

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.


Prerequisite

Only if you plan to protect physical machines or volume groups that are attached to a virtual machine by using iSCSI. Credentials are assigned to physical machines 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 74](#).

Nutanix cluster considerations



- 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.
- If during virtual machine synchronization, a virtual machine cannot be found on a Nutanix cluster, the status of this virtual machine or any discovered applications running on it is set to PENDING_REMOVAL. Such a virtual machine and its applications are grayed out in HYCU and you cannot perform any data protection actions for them. If during the time interval of two automatic virtual machine synchronization processes, the virtual machine is found on the Nutanix cluster, its status is changed to PROTECTED_DELETED. Otherwise, the virtual machine is removed from HYCU.


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  **Synchronize**. To narrow down the list of displayed virtual machines, you can use the filtering options described in [“Filtering data” on page 155](#).


2. Click  **Policies**. The Policies dialog box opens.
3. From the list of available policies, select the desired policy.
4. Click **Assign** to assign the policy to the selected 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 167](#).

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.

 **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

- *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 161](#).
- *For vSphere environments:* You have the required restore privileges assigned. For details, see [“Assigning privileges to a vSphere user” on page 233](#).
- 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 physical machines:* At least one Nutanix cluster or vCenter Server is added to HYCU to provide a storage container for storing the restore data. For details on how to add a Nutanix cluster to HYCU, see [“Adding a Nutanix cluster” on page 31](#). For details on how to add a vCenter Server to HYCU, see [“Adding a vCenter Server” on page 33](#).

Limitation




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.

Consideration


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 264](#).

Restore options

You can select among the following restore options:

Restore option	Description
Restore VM	<p>Enables you to restore a virtual machine. Select this option if you want to 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 a physical machine by using this option.</p>
Clone VM	<p>Enables you to restore a virtual machine by creating its clone. Select this option if you want to keep the original virtual machine. For instructions, see “Cloning a virtual machine” on page 85.</p>
Restore vDisks	<p>Enables you to restore virtual disks. Select this option if you want to replace the original virtual disks with the restored ones. For instructions, see “Restoring virtual disks” on page 90.</p> <p> Important You cannot restore physical machine 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 91.</p> <p> Important You cannot restore vSphere virtual machine disks by using this option.</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</p>

Restore option	Description
	with specific access permissions, or if you want to use the virtual disks later to restore data to a physical machine or to an environment with a hypervisor not supported by HYCU or not added to HYCU as a source. For instructions, see “Exporting virtual disks” on page 93 .

 **Note** By using the Clone VM option, you can also restore a virtual machine to an environment with a different hypervisor. For prerequisites, limitations, considerations, and/or additional steps that you should perform to successfully restore a virtual machine to an environment with a different hypervisor, see [“Restoring to an environment with a different hypervisor” on page 270](#).

Restoring a virtual machine

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


Limitations

- Restoring a virtual machine running on a Nutanix ESXi cluster by using the Restore VM option is supported only if your AOS version is 5.10 or later.
- Restoring physical machines by using the Restore VM option is not supported.

Considerations


- A restore is performed from the snapshot only if the snapshot is available on the original location (the source where the original virtual machine was running). Otherwise, a restore is performed from the target.
- *For volume groups:* If there are one or more volume groups 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.



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

1.  check box before the name of the virtual machine will not open the Detail view.
2. In the Detail view, select the desired restore point.
3. Click  **Restore VM**.
4. Select **Restore VM**, and then click **Next**.
5. 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, fast restore cannot be performed, because the restore will be performed from the target and not from the snapshot.

- b. 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. The maximum number of virtual CPUs is 1024.
 - In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine. The maximum number of cores per virtual CPU is 64.

 **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. The value that you specify must be a whole number and cannot be greater than 4096 GiB.
 - c. 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.
 - d. From the Restore instance drop-down menu, select which restore instance you want to use for a restore. Your restore point can contain one or more restore instances among which you can select:
 - **Automatic**: This type of restore ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
 - 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.


6. In the Network section, do the following:

Review the list of network adapters for the selected restore point. Under VM Network, you can view to which networks the virtual machine was connected at backup time. If any of the original networks is no longer available, N/A is shown. Depending on whether the original networks are available, do the following:

The original VM networks are...	Procedures
Available	<ul style="list-style-type: none"> To restore the virtual machine with the original network settings, leave the default values. To restore the virtual machine with different network settings, you can do the following: <ul style="list-style-type: none"> Add a new network adapter by clicking + New and selecting the desired 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 desired network. Delete the network adapter you do not need anymore by selecting it, and then clicking Delete.
Unavailable	<ul style="list-style-type: none"> To proceed with the restore, do one of the following: <ul style="list-style-type: none"> Edit the affected network adapter to connect the virtual machine to a new network by selecting the network adapter, and then clicking Edit and selecting the desired network. Delete the affected network adapter by selecting it, and then clicking Delete. Optionally, add a new network adapter by clicking + New and selecting the desired network.

 **Note** You can restore the virtual machine without a network adapter. Make sure to configure the network settings on the virtual machine afterward.

7. 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.

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. In this case, the original virtual machine will not be overwritten.

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 is added to HYCU. For details on how to do this, see [“Adding a Nutanix cluster” on page 31](#) or [“Adding a vCenter Server” on page 33](#).
- *For Linux physical machines:* References to the devices that are included in the backup (LVM volumes, swap partitions, and the boot device) in the `/etc/fstab` system configuration file entries must use universally unique identifiers (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`).

Limitations

- *For vSphere environments:* Attaching the ISO image to the restored virtual machine is not supported.
- Cloning physical machines that use UEFI firmware is not supported.

Considerations


- A restore is performed from the snapshot only if the snapshot is available on the original location (the source where the original virtual machine was running). Otherwise, a restore is performed from the target.
- *For volume groups:* If there are volume groups attached to the virtual machine that you are restoring by creating its clone, 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 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 [“Restoring a virtual machine from a vSphere environment to a Nutanix ESXi cluster” on page 274.](#)


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.
2. In the Detail view, select the desired restore point.
3. Click  **Restore VM**.
4. Select **Clone VM**, and then click **Next**.
5. In the General section, do the following:
 - a. From the Storage container drop-down menu, select where you want to restore the virtual machine.


 **Note** *For virtual machines:* 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:

- 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 hypervisor, additional prerequisites apply. For details, see [“Restoring to an environment with a different hypervisor” on page 270.](#)

- b. In the New VM name field, specify a new name for the virtual machine.
- 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. The maximum number of virtual CPUs is 1024.
 - In the Cores per vCPU field, enter the number of cores per virtual CPU for the


restored virtual machine. The maximum number of cores per virtual CPU is 64.

 **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. The value that you specify must be a whole number and cannot be greater than 4096 GiB.
- d. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore. If you turn the restored virtual machine on, the original virtual machine will be turned off automatically.
-  **Important** 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 the next page](#).
- e. From the Restore instance drop-down menu, select which restore instance you want to use for a restore. Your restore point can contain one or more restore instances among which you can select:
- **Automatic:** This type of restore ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
- 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.
6. In the Network section, do the following:
- a. Review the list of network adapters for the selected restore point. Under VM Network, you can view to which networks the virtual machine was connected at backup time. If any of the original networks is no longer available, N/A is shown. Depending on whether the original networks are available, do the following:

The original VM networks are...	Procedures
Available	<ul style="list-style-type: none"> • To restore the virtual machine with the original network

The original VM networks are...	Procedures
	<p>settings, leave the default values.</p> <ul style="list-style-type: none"> To restore the virtual machine with different network settings, you can do the following: <ul style="list-style-type: none"> Add a new network adapter by clicking + New and selecting the desired 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 desired network. Delete the network adapter you do not need anymore by selecting it, and then clicking Delete.
Unavailable	<ul style="list-style-type: none"> To proceed with the restore, do one of the following: <ul style="list-style-type: none"> Edit the affected network adapter to connect the virtual machine to a new network by selecting the network adapter, and then clicking Edit and selecting the desired network. Delete the affected network adapter by selecting it, and then clicking Delete. Optionally, add a new network adapter by clicking + New and selecting the desired network.

 **Note** You can restore 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.

7. Click **Restore**.

After cloning a virtual machine

After cloning a virtual machine, consider the following:

- 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, RHEL 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 physical machines:*
 - *Only if you cloned a Windows physical machine to a virtual machine running on 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.
 - *For Linux physical machines:*
 - Because the original boot loader of the physical machine is replaced with a temporary one during the backup, it is recommended to update the boot configuration after the restore. For details on how to do this, see [“Updating the boot configuration of Linux physical machines”](#) below.
 - *Only if you cloned the physical machine to a Nutanix ESXi cluster or a vSphere environment.* Make sure to change the storage controller on the cloned virtual machine to VMware Paravirtual SCSI. For details, see VMware documentation.

Updating the boot configuration of Linux physical machines

Procedure

1. 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 cloned physical machine to match the resume device UUID on the original physical machine. For example, if the resume device on the original physical machine is `resume=/dev/mapper/cl-swap`, the resume device on the cloned physical machine should be `resume=UUID=4044243b-612b-42bc-ba22-4736c4eadd6`.
2. *Optional.* If you want to speed up the boot process by skipping mounting non-existent volumes, do the following:

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 is an example of the warning message:

```
Non LVM volumes detected: Following volumes are not backupable:
/dev/sdf3:/test_mount.
```


In this example, comment the line that contains the `/test_mount` mountpoint in the `/etc/fstab` system configuration file.

3. Update the GRUB configuration by running the following command:

```
grub2-mkconfig -o /boot/grub2/grub.cfg
```

4. Install the GRUB 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 system.

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.


Limitation

Restoring physical machine disks by using the Restore vDisks option is not supported.

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 152](#).
- 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.


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 desired 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 VM**.
4. Select **Restore vDisks**, and then click **Next**.
5. From the list of virtual disks that are available for a restore, select the ones that you want to restore, and then click **Next**.

 **Important** *For volume groups:* You cannot select individual disks, but only the whole volume group.

6. 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.

7. From the Restore instance drop-down menu, select which restore instance you want to use for a restore. Your restore point can contain one or more restore instances 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. 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.

Limitation

Restoring vSphere virtual machine disks by using the Clone vDisks option 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 152](#).
- *For volume groups:* The original volume groups are kept alongside 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 desired 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 VM**.
4. Select **Clone vDisks**, and then click **Next**.
5. From the list of virtual disks that are available for a restore, select the ones that you want to restore, and then click **Next**.

 **Important** *For volume groups:* You cannot select individual disks, but only the whole volume group.

6. From the 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 physical machine.
7. 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.
 8. From the Restore instance drop-down menu, select which restore instance you want to use for a restore. Your restore point can contain one or more restore instances 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*)
 9. 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 physical machine. For details, see [“Restoring data to a physical machine” on page 95](#).


Prerequisite

For restoring virtual disks to an SMB share: The SMB server is configured to stop creating sparse files (the `strict allocate` parameter is set to `yes` in the `smb.conf` file).

Considerations

- A restore is performed from the snapshot only if the snapshot is available on the original location (the source where the original virtual machine was running). Otherwise, a restore is performed from the target.
- 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 152](#).


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 desired 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 VM**.

4. 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.

5. From the list of virtual disks that are available for a 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:

- **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).

7. From the Restore instance drop-down menu, select which restore instance you want to use for a restore. Your restore point can contain one or more restore instances 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. Click **Restore**.

After exporting virtual disks

After the restore of the virtual disks is complete, you can use them to restore data to a physical machine or to an environment with a hypervisor not supported by HYCU or not added to HYCU as a source.

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 hypervisor in your environment, the following files are created for each selected disk:

Hypervisor	Files
Nutanix AHV	<DiskName> (without extensions)
Nutanix ESXi	A raw image of the disk, including unallocated space as zeroes
vSphere	<ul style="list-style-type: none"> • <DiskName>-flat.vmdk A raw image of the disk • <DiskName>.vmdk A VMDK descriptor file, referencing <DiskName>-flat.vmdk

Restoring data to a physical machine

The procedure described in this section is an example of how to restore data to a Windows physical machine.

Prerequisites

- The physical machine 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 physical machine 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:

```
PhysicalDisk0 PhysicalDisk1
```

4. Verify each exported disk by running the following command:

```
fdisk -l <ExportedDiskName>
```

For example:

```
fdisk -l PhysicalDisk0
```

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 `PhysicalDisk0` matches the size of disk `/dev/sda`. Therefore, disk `PhysicalDisk0` can be restored to disk `/dev/sda`.

The following is an example of the output:

```
Disk PhysicalDisk0: 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
```

```
PhysicalDisk0p1 * 2048 718847 716800 350M 7
```

```
HPFS/NTFS/exFAT
```

```
PhysicalDisk0p2 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=PhysicalDisk0 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

```

6. Eject the Linux live CD and reboot the physical machine.

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 (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 171](#).

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 236](#).

Prerequisites

Windows virtual machines	<ul style="list-style-type: none"> • The NTFS, FAT, or FAT32 file system is used. • For improved restore performance, the startup type of the Microsoft iSCSI Initiator Service is not 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 is enabled and configured by using the <code>winrm quickconfig</code> command. ◦ A Windows operating system user account exists that has WinRM permissions granted and is 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 74. ◦ <i>For pre/post-restore scripts:</i> A script is available in the accessible folder and has one of the following extensions: bat, ps1, cmd.
Linux virtual machines	<ul style="list-style-type: none"> • The FAT32, xfs, ext4/ext3/ext2, reiserfs, or btrfs file system is used. • References in the <code>/etc/fstab</code> system configuration file entries use

	<p>universally unique identifiers (for example, UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5) rather than device names (for example, /dev/sda1) unless the entries refer to logical volumes (for example, /dev/mapper/ol-root).</p> <ul style="list-style-type: none"> • <i>For restoring files to a virtual machine:</i> <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 74. ◦ <i>For pre/post-restore scripts:</i> A script is available in the accessible folder and has the sh extension. You 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 are installed on the client virtual machine. For detailed information about installing VMware Tools, see VMware documentation. For detailed information about installing NGT, see Nutanix documentation.
vSphere environments	<ul style="list-style-type: none"> • At least one NFS, SMB, iSCSI, or Nutanix target is available on the HYCU backup controller. • You have the required restore privileges assigned. For details, see “Assigning privileges to a vSphere user” on page 233.
Only if restoring data from tape	<ul style="list-style-type: none"> • 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 161.

Limitations

- Restoring individual files on dual-boot systems 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.
- *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 4 GiB.
- You cannot restore individual files if you excluded all virtual machine disks from the backup and left only the attached volume groups.


Considerations

- 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 264](#).
- *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 152](#).
- *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.

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 91](#).


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 desired restore point.
3. Click  **Restore Files**. The Restore Files dialog box appears.
4. In the Restore Files dialog box, 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.

5. Select the restore option that defines the location for the file restore, click **Next**, and follow the instructions for the selected option:


Restore option	Instructions
Restore to virtual machine	<ol style="list-style-type: none"> a. In the General tab, select the location for restoring the files on the same or a different virtual machine, and provide the required information: <ul style="list-style-type: none"> • Original location Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file). • Alternate location Specify the path to an alternate location on the virtual machine in the following format: <div style="border: 1px solid #ccc; padding: 2px; margin-top: 5px;">C:\<Path></div> b. Use the Restore ACL switch if you want to restore the original access control list. <div style="border-left: 2px solid #007bff; padding-left: 10px; margin-top: 10px;">  Important If the virtual machine is not accessible due to various reasons (for 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. </div> c. <i>Optional.</i> Click the Pre/post scripts tab. Use the switches of your choice to specify the pre/post-restore scripts, and then enter the script path names. Enable one or more switches: <ul style="list-style-type: none"> • Run pre-restore script • Run post-restore script <div style="border-left: 2px solid #007bff; padding-left: 10px; margin-top: 10px;">  Note In the script path name field, a sample path name is displayed. Make sure to enter the valid script path name. </div> d. Click Save. e. Click Restore.
Restore to	<ol style="list-style-type: none"> a. Select NFS or SMB for the share type, and then specify the

Restore option	Instructions
fileshare	<p>path to a shared folder in the following format:</p> <pre>\\server\<Path></pre> <ol style="list-style-type: none"> <i>For SMB:</i> Optionally, provide user credentials to access the SMB share. Click Restore.
Download	<p>Click Download to save a data archive of the selected files locally.</p> <p> Important Make sure not to reload the current page or log out from the domain before the download process is finished.</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 physical machines 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 105](#)
- [“Backing up applications” on page 110](#)
- [“Restoring whole applications” on page 111](#)
- [“Restoring SQL Server databases” on page 119](#)
- [“Restoring Exchange Server databases, mailboxes, and public folders” on page 122](#)
- [“Restoring Oracle database instances and tablespaces” on page 125](#)

Enabling access to application data


After you assign credentials to virtual machines as described in [“Enabling access to data” on page 74](#), 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 and physical machines. 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:



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 110](#).

 **Note** Access to Active Directory and SAP HANA is always granted with

	the virtual machine credentials.
✖	<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 74. • 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"> • <i>For Windows 8 and 10:</i> 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 74
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 74
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.

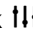
Consideration


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. 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 Organization Management role group and have the default permissions enabled.
 - 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 110](#).

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 mark restore points as expired.

Planning application protection

Before performing an application backup, get familiar with prerequisites, limitations, considerations, recommendations, and application-specific options to determine if your environment is ready for application data protection.

Prerequisites

- *For vSphere environments:* VMware Tools of the latest version is 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 are available on the system, are owned by root, and have permissions set to 0700: `/usr/local/sbin/pre_freeze` and `/usr/local/sbin/post_thaw`.
- *For HYCU Protégé:* Make sure all the prerequisites for migrating virtual machines and applications to cloud listed in [“HYCU Protégé specifics” on page 72](#) are fulfilled.
- *For physical machines:*
 - *For Windows:*
 - 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 Linux:* Access to the physical machine through SSH is enabled.
- *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.

Application-specific prerequisites

Application type	Prerequisites
SQL Server	<ul style="list-style-type: none"> • Databases 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 is disabled. For details, see Nutanix documentation. • <i>For restoring an SQL Server database to a point in time:</i> The database is online and is 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 are protected by HYCU or only the node with the synchronized databases of the Always On Availability Group (must be online when being protected). In the latter case, the risk of data loss is increased if the node goes offline or the databases get out of sync. • <i>For using a separate disk volume as backup storage for SQL Server</i>

Application type	Prerequisites
	<p><i>temporary files:</i> Make sure that a dedicated disk of a sufficient size is allocated. The volume should be able to store temporary files that are generated between two backups of your SQL Server database.</p> <ul style="list-style-type: none"> • <i>For SQL Server failover cluster:</i> <ul style="list-style-type: none"> ◦ All virtual machines where an SQL Server failover cluster resides are discovered by HYCU. ◦ Policies are assigned to all virtual machines on which the application instance is running.
Active Directory	<ul style="list-style-type: none"> • NGT is installed and enabled on the client virtual machine. For details on how to do this, see Nutanix documentation. • No volume groups are attached to the client virtual machine.
Exchange Server	<ul style="list-style-type: none"> • NGT is installed and enabled on the client virtual machine. For details on how to do this, see Nutanix documentation. • No volume groups are attached to the client virtual machine. • All databases are mounted. • The Active Directory application is 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 first need to 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 need to restore only the Exchange Server application.</p>
Oracle	<ul style="list-style-type: none"> • The SSH service is enabled on the Oracle server and is listening on port 22 for incoming connections. • The Oracle database user has the SYSDBA privilege. • The database is running in ARCHIVELOG mode. • Tablespaces are online. • Additional disk space must be provided for temporary files created between two database backups. For optimal restore performance, separate disks should be specified for the temporary and database files.
SAP HANA	<ul style="list-style-type: none"> • SAP HANA savepoints are enabled.

Application type	Prerequisites
	<ul style="list-style-type: none"> • <i>For multiple volume groups:</i> All data volumes and log volumes belong to the same volume group. <p><i>For distributed (multi-host) environments:</i></p> <ul style="list-style-type: none"> • All virtual machines where SAP HANA resides are discovered by HYCU. • Policies are 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 with AOS version 5.10:* Backing up applications running on virtual machines that are part of a Nutanix protection domain 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 Nutanix AHV clusters:* With a Nutanix protection domain configured with NearSync, HYCU can create snapshots only if the AOS version is 5.10 or later.

Application-specific limitations

Application type	Limitations
SQL Server	<ul style="list-style-type: none"> • The tempdb SQL Server system database is excluded from all backups. • Only a full backup of the master, model, and msdb 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 master, model, msdb, or tempdb 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"> • <i>For Nutanix clusters:</i> Protecting applications that are running on

Application type	Limitations
	<p>virtual machines with IDE disks is not possible.</p> <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.
Exchange Server	<ul style="list-style-type: none"> <i>For Nutanix clusters:</i> Protecting applications that are running on virtual machines with IDE disks is not possible. <i>For vSphere environments:</i> Log files are not truncated after a backup. Backing up the applications running on the volume groups or on the virtual machines with the attached volume groups 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:
 - Assign credentials to the upgraded application to enable HYCU to access it. For details, see [“Enabling access to application data” on page 102](#).
 - Assign a policy to the upgraded application to protect it. For details, see [“Backing up applications” on page 110](#).
 - 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.

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.

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 Configuration dialog box

To access the Configuration dialog box, 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**.

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> 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. 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. • Enter path for temporary translog and metadata files <i>(optional)</i> If specified, the backup copies of the SQL Server temporary files (transaction logs and metadata files) will be stored to this location. Else, the backup copies are stored to the /tmp/hycu folder. <p> Note For better restore performance, it is recommended to use a dedicated disk for storing backup copies of temporary files.</p>
Exchange Server	<ul style="list-style-type: none"> • Priority for Exchange Server restore requests Specifies the priority in which the restore requests for a mailbox restore are processed on the Exchange Server: Lowest, Lower, Low, Normal (the default value), High, Higher, Highest, Emergency. • Optimized Exchange Server DAG protection

	<p><i>Available only for Windows physical machines that are part of a database availability group (DAG).</i> Enable this option if you want to back up only the disks hosting the passive database copies with the highest activation preference number (including the system disk). If no passive database copies are available, active database copies will be backed up.</p> <p> Important Optimized Exchange Server DAG protection is effective only if separate databases are stored on separate disks.</p>
Oracle	<ul style="list-style-type: none"> • Back up and truncate Oracle archive logs (<i>enabled by default</i>) 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. 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. • Enter path for temporary Oracle files (<i>optional</i>) If specified, the backup copies the temporary Oracle files will be stored to this location. <p> Note For better restore performance, it is recommended to use a dedicated disk for storing backup copies of temporary files.</p>

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**.


Consideration


For Nutanix clusters: If during virtual machine synchronization, a virtual machine cannot be found on a Nutanix cluster, the status of this virtual machine or any discovered applications running on it is set to PENDING_REMOVAL. Such a virtual machine and its applications are grayed out in HYCU and you cannot perform any data protection actions for them. If during the time interval of two automatic virtual machine synchronization processes, the virtual machine is found on the Nutanix cluster, its status is changed to PROTECTED_DELETED. Otherwise, the virtual machine is 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 data” on page 155](#).


2. Click  **Policies**. The Policies dialog box appears.
3. From the list of available policies, select the desired 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 167](#).

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 vSphere environments:* You have the required restore privileges assigned. For details, see [“Assigning privileges to a vSphere user” on page 233](#).
- *For applications with status PROTECTED_DELETED whose backups are stored on the imported targets:* Discover these applications, for details, see [“Enabling access to application data” on page 102](#).
- *For physical machines:* At least one Nutanix cluster or vCenter Server is added to HYCU to provide a storage container for storing the restore data. For details on how to add a Nutanix cluster to HYCU, see [“Adding a Nutanix cluster” on page 31](#). For details on how to add a vCenter Server to HYCU, see [“Adding a vCenter Server” on page 33](#).
- *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 161](#).

Considerations


- 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 264](#).

- *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 [“Checking the status of jobs” on page 144](#).


Restore options

You can select between the following restore options:

Restore option	Description
Restore VM	<p>Enables you to restore an application by restoring the virtual machine on which it is running. Select this option if you want to replace the original virtual machine on which your application is running with the restored one. For instructions, see “Restoring a virtual machine” below.</p> <p> Important You cannot restore an SQL Server or Exchange Server application running on a physical machine by using this option.</p>
Clone VM	<p>Enables you to restore a virtual machine by creating its clone. Select this option if you want to keep the original virtual machine on which your application is running. For instructions, see “Cloning a virtual machine” on page 116.</p>

Restoring a virtual machine

HYCU enables you to restore an application by restoring the virtual machine on which it is running to its original or a new location. 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 167](#).


Limitations

- Restoring a virtual machine running on a Nutanix ESXi cluster by using the Restore VM option is supported only if your AOS version is 5.10 or later.
- Restoring SQL Server, Exchange Server and Oracle applications running on physical machines by using the Restore VM option is not supported.

Considerations



- A restore is performed from the snapshot only if the snapshot is available on the original location (the source where the original virtual machine was running). Otherwise, a restore is performed from the target.
- *For volume groups:* If there are one or more volume groups 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.


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 desired restore point, and then click  **Restore**.

 **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. Select **Restore whole server**, and then click **Next**.
4. Select **Restore VM**, and then click **Next**.
5. 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, fast restore cannot be performed, because the restore will be performed from the target and not from the snapshot.


- b. 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. The maximum number of virtual CPUs is 1024.
 - In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine. The maximum number of cores per virtual CPU is 64.

 **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. The value that you specify must be a whole number and cannot be greater than 4096 GiB.
 - c. 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.
 - d. From the Restore instance drop-down menu, select which restore instance you want to use for a restore. Your restore point can contain one or more restore instances among which you can select:
 - **Automatic:** This type of restore ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
 - 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.
6. In the Network section, do the following:

Review the list of network adapters for the selected restore point. Under VM Network, you can view to which networks the virtual machine was connected at backup time. If any of the original networks is no longer available, N/A is shown. Depending on whether the original networks are available, do the following:

The original VM networks are...	Procedures
Available	<ul style="list-style-type: none"> To restore the virtual machine with the original network settings, leave the default values. To restore the virtual machine with different network settings, you can do the following: <ul style="list-style-type: none"> Add a new network adapter by clicking + New and selecting the desired 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 desired network. Delete the network adapter you do not need anymore by selecting it, and then clicking Delete.
Unavailable	<ul style="list-style-type: none"> To proceed with the restore, do one of the following: <ul style="list-style-type: none"> Edit the affected network adapter to connect the virtual machine to a new network by selecting the network adapter, and then clicking Edit and selecting the desired network. Delete the affected network adapter by selecting it, and then clicking Delete. Optionally, add a new network adapter by clicking + New and selecting the desired network.

 **Note** You can restore the virtual machine without a network adapter. Make sure to configure the network settings on the virtual machine afterward.

7. 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.

After restoring a virtual machine

After restoring an Exchange Server or Active Directory application, reinstall NGT to ensure future successful backups of application data.

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. In this case, the original virtual machine will not be overwritten.

Prerequisite

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 is added to HYCU. For details on how to do this, see [“Adding a Nutanix cluster” on page 31](#) or [“Adding a vCenter Server” on page 33](#).


Considerations

- A restore is performed from the snapshot only if the snapshot is available on the original location (the source where the original virtual machine was running). Otherwise, a restore is performed from the target.
- *For volume groups:* If there are volume groups attached to the virtual machine that you are restoring by creating its clone, 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 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 [“Restoring a virtual machine from a vSphere environment to a Nutanix ESXi cluster” on page 274](#).

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 desired restore point, and then click  **Restore**.


 **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. Select **Restore whole server**, and then click **Next**.
4. Select **Clone VM**, and then click **Next**.
5. In the General section, do the following:
 - a. From the Storage container drop-down menu, select where you want to restore the virtual machine.

 **Note** *For virtual machines:* 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:






- 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 hypervisor, additional prerequisites apply. For details, see [“Restoring to an environment with a different hypervisor” on page 270](#).


- b. In the New VM name field, specify a new name for the virtual machine.
- 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. The maximum number of virtual CPUs is 1024.
 - In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine. The maximum number of cores per virtual CPU is 64.

 **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. The value that you specify must be a whole number and cannot be greater than 4096 GiB.
- d. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore. If you turn the restored virtual machine on, the original virtual machine will be turned off automatically.
-  **Important** 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 88](#).
- e. From the Restore instance drop-down menu, select which restore instance you want to use for a restore. Your restore point can contain one or more restore instances among which you can select:
- **Automatic:** This type of restore ensures the fastest restore to the latest state.
 - **Backup**
 - **Copy**
 - **Archive**
 - **Snapshot**
- 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.
6. In the Network section, do the following:
- a. Review the list of network adapters for the selected restore point. Under VM Network, you can view to which networks the virtual machine was connected at backup time. If any of the original networks is no longer available, N/A is shown. Depending on whether the original networks are available, do the following:

The original VM networks are...	Procedures
Available	<ul style="list-style-type: none"> • To restore the virtual machine with the original network settings, leave the default values. • To restore the virtual machine with different network settings, you can do the following: <ul style="list-style-type: none"> • Add a new network adapter by clicking + New and selecting the desired network.

The original VM networks are...	Procedures
	<ul style="list-style-type: none"> • 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 desired network. • Delete the network adapter you do not need anymore by selecting it, and then clicking  Delete.
Unavailable	<ul style="list-style-type: none"> • To proceed with the restore, do one of the following: <ul style="list-style-type: none"> • Edit the affected network adapter to connect the virtual machine to a new network by selecting the network adapter, and then clicking  Edit and selecting the desired network. • Delete the affected network adapter by selecting it, and then clicking  Delete. • Optionally, add a new network adapter by clicking  New and selecting the desired network.

 **Note** You can restore 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.

7. 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 88](#).

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 is set to full or bulk-logged.
- *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 161](#).

- *For restoring the whole SQL Server failover cluster instance:* The SQL Server service is 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 is not set to Disabled.
- *For physical machines:* At least one Nutanix cluster or vCenter Server is added to HYCU to provide a storage container for storing the restore data. For details on how to add a Nutanix cluster to HYCU, see [“Adding a Nutanix cluster” on page 31](#). For details on how to add a vCenter Server to HYCU, see [“Adding a vCenter Server” on page 33](#).


Limitations

- The restore of discovered applications is available for the NTFS, FAT, and FAT32 file systems.
- 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.

Considerations


- 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 264](#).
- *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.

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 desired 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 Destination storage container drop-down menu, select where you want to restore the databases.
8. Select the **Whole instance** check box if you want to restore the whole application instance or, from the list of databases that are available for a restore, select the ones that you want to restore.
9. *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.
10. Click **Next**.
11. 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.

12. Click **Restore**.
13. *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.
14. *Only if using SQL Server 2012 and 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.

15. *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 is restored to the original mailbox server. From there, the actual restore is performed to any mailbox or public folder within the organization.

Prerequisites

- *For restoring public folders:* The public folder exists in the public folder mailbox. If it does not exist, recreate it manually with the same name it had at backup time.
- *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 161](#).
- For improved restore performance, the startup type of the Microsoft iSCSI Initiator Service is not set to Disabled.
- *For physical machines:* At least one Nutanix cluster or vCenter Server is added to HYCU to provide a storage container for storing the restore data. For details on how to add a Nutanix cluster to HYCU, see [“Adding a Nutanix cluster” on page 31](#). For details on how to add a vCenter Server to HYCU, see [“Adding a vCenter Server” on page 33](#).


Limitations

- The restore of discovered applications is available for the NTFS, FAT, and FAT32 file systems.
- Restoring data to the hycu subfolder (the Restore to subfolder option) is currently not supported for public folders.

Consideration


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 264](#).

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 desired 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 MS Exchange Server dialog box appears.
4. Select which application items you want to restore:

- **Restore databases**

- a. From the Destination server drop-down menu, select the server for restoring the data. When specifying a destination 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 destination server on which the databases are currently active.

- b. From the Destination storage container drop-down menu, select the storage container for restoring the data.

- c. Select the **All databases** check box if you want to restore all databases or, from the list of databases that are available for a restore, select the ones that you want to restore.
- d. 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**

- a. From the Storage container drop-down menu, select where you want to restore the mailboxes and/or public folders.
- b. From the list of mailboxes and/or public folders that are available for a restore, select the ones that you want to restore, and then click **Next**.

 **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.

- c. Select where you want to restore data:
 - **Original mailbox**
 - **Alternate mailbox**, and then enter an alternate mailbox name.
- d. 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.
- e. *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.
- f. 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 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 do not write to standard output (STDOUT) or standard error (STDERR) for the user whose credentials are used for application discovery.
- *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 161](#).
- *For physical machines:* At least one Nutanix cluster or vCenter Server is added to HYCU to provide a storage container for storing the restore data. For details on how to add a Nutanix cluster to HYCU, see [“Adding a Nutanix cluster” on page 31](#). For details on how to add a vCenter Server to HYCU, see [“Adding a vCenter Server” on page 33](#).

Limitation

Tablespaces can be restored only from the latest restore point in the backup chain and cannot be restored to a point in time.

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.


- 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 264](#).

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 desired 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 Server 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 database**, and then click **Next**.
5. Select the **Whole instance** check box if you want to restore the whole database instance or, from the list of tablespaces that are available for a restore, select the ones that you want to restore.
6. *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.
7. Click **Restore**.

8. *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:

- [“Backing up file shares” below](#)
- [“Restoring file share data” on the next page](#)

Backing up file shares

A file-share backup allows a rapid backup of file shares by using parallel backup streams.

Prerequisite

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.

Limitations

- The iSCSI and Nutanix 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).
- *For encrypted file shares:* The backup of alternate data streams (ADS) is not supported.


Considerations

- You can change the number of incremental file share backups after which a full reindex is performed so that it suits the requirements of your environment. By doing so, you speed up the process of searching for the relevant files when you are restoring them. For details on how to customize HYCU configuration settings, see [“Customizing HYCU configuration settings” on page 264](#).
- If up to 100 file backups fail during the backup of a file share, the backup status of the file share is Completed with errors. You can customize this setting to suit your data

protection needs. For details on how to do this, see [“Customizing HYCU configuration settings” on page 264](#).



- *For Nutanix Files 3.6.0 or later:* When backing up a file share, HYCU also backs up any file shares that are inside the selected file share. Keep in mind that backing up nested file shares individually is not supported.



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  **Synchronize**. To narrow down the list of displayed file shares, you can use the filtering options described in [“Filtering data” on page 155](#).

2. *Only if you want to exclude particular file share folders from the backup.* Click  **Configuration**, and then, in the Configuration dialog box that opens, do the following:
 - a. In the Exclude folder path field, enter the full path (from the root of the file share) to the file share folder that you want to exclude from the backup (for example, /backup), and then click  **Add**. Repeat this step to add additional file share folders.

 **Note** The paths to all the file share folders 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  **Remove**.

 - b. Click **Save**.
3. Click  **Policies**. The Policies dialog box appears.
4. From the list of available policies, select the desired policy.
5. 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 167](#).

 **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 desired 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 Nutanix Files share, or to an external SMB or NFS share.

Prerequisites

- *For restoring data to a different Nutanix Files share:* The Nutanix Files server with the file share to which you want to restore data is added to HYCU. For details on how to do this, see [“Adding a Nutanix Files server” on page 34](#).
- *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 161](#).

Limitations

- The restore of alternate data streams (ADS) is supported only if you are restoring data from one Nutanix Files SMB share to another Nutanix Files SMB share.
- Symbolic links are restored only if you are restoring data from one NFS share to another NFS share.
- *Only if restoring files to an external share:*
 - Security information is not restored.
 - Restoring files or folders with newlines in their names is supported only for an NFS share set up on Unix.


Consideration

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 264](#).

Recommendation


For optimal restore performance, it is recommended that you restore data to a Nutanix Files 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 desired restore point.
3. Click  **Restore Files**. The Restore Files dialog box opens.
4. 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.

5. Depending on where you want to restore the selected files (to the original or a different Nutanix Files share, or to an external SMB or NFS share), select the desired restore option and follow the instructions:

Restore option	Instructions
Restore to Nutanix Files share	<ol style="list-style-type: none"> a. From the Share drop-down menu, select the Nutanix Files share to which you want to restore the files. b. Select the location on the Nutanix Files share where you want to restore the files, and then provide the required information: <ul style="list-style-type: none"> • Original location Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file).  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. • Alternate location (on the same share) Specify the path to an alternate location on the same share in the following format: <div style="background-color: #f0f0f0; padding: 2px; margin-top: 5px;"><code>/<Path></code></div> c. <i>Only if restoring files from one SMB share to another SMB share.</i> Use the Restore ACL switch if you want to restore the original access control list for the files.
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 Enter the path to the NFS shared folder in the following format:

Restore option	Instructions
	<div>\\server\<Path></div> <ul style="list-style-type: none">• SMB<ol style="list-style-type: none">Enter the path to the SMB shared folder in the following format:<div>\\server\<Path></div><i>Optional.</i> Provide user credentials to access the SMB share.

6. Click **Restore**.

Chapter 7

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, and file shares

Consideration

For Migration/DR-ready virtual machines and applications: You can recover your HYCU backup controller and protected data to cloud by using HYCU Protégé. For more information, see [“HYCU Protégé” on page 249](#).

Procedures

1. Prepare for disaster recovery. For instructions, see [“Preparing for disaster recovery” below](#).
2. Perform disaster recovery. For instructions, see [“Performing disaster recovery” on page 136](#).
3. *Only if HYCU is used for file share protection:* Reestablish connections of HYCU instances to the restored HYCU backup controller or recreate HYCU instances. For instructions, see [“Recreating HYCU instances” on page 140](#).

Preparing for disaster recovery

Prerequisites

- You know configuration parameters of the targets that store 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 69](#).
- The targets that store backup data of the entities you want to recover are accessible to the source where you plan to deploy a temporary HYCU backup controller.
- *Only if the backup of the original HYCU backup controller is stored on an iSCSI target.* The

iSCSI storage device is 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 and applications you want to recover is stored on a Google Cloud target.* A Google Cloud Platform service account is created and added to HYCU. For instructions on how to add a cloud account to HYCU, see [“Adding a Google Cloud Platform service account” on page 186](#).
- *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.* You exported the encryption target key from the original HYCU backup controller and the file containing the encryption key is available.

Procedure

Task	Instructions
1. Deploy a temporary HYCU backup controller.	“Deploying a temporary HYCU backup controller” below
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, and file shares.	“Importing targets” on the next page
3. Add a source to which you plan to restore your HYCU backup controller. If you plan to restore also virtual machines, applications, and file shares, add the sources to which you plan to restore them.	“Adding sources” on page 31

Deploying a temporary HYCU backup controller

Procedure

1. Log on to the Nutanix Prism web console (for Nutanix AHV clusters) or the vSphere (Web) Client (for Nutanix ESXi clusters and vSphere environments).
2. Deploy a temporary HYCU backup controller that you will use for restoring the original HYCU backup controller or other entities. Depending on the environment to which you want to deploy it, see one of the following:
 - *For a Nutanix AHV cluster:* [“Deploying HYCU to a Nutanix AHV cluster” on page 22](#).
 - *For a Nutanix ESXi cluster and a vSphere environment:* [“Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment” on page 26](#).
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 264](#).



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. Log on to the HYCU web user interface of the temporary 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 189](#).

Importing targets


Prerequisites

- The activities on the temporary HYCU backup controller and the original HYCU backup controller (if it still exists) are suspended and no jobs are running when you start importing the targets. For instructions, see [“Setting power options” on page 200](#).
- The temporary HYCU backup controller has either no targets or only imported targets.
- All targets you plan to import are deactivated and iSCSI and Nutanix targets are also unmounted on any other powered on HYCU backup controller until the import job is finished.


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.
- After you import the target, a volume group of the storage container that is associated with this target has access to the HYCU backup controller.


Accessing the Targets panel

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


Procedure

1. Log on to the HYCU web user interface of the temporary 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 Backup 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.

 **Important** Archive targets must be imported separately from other targets.

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 before all the required data is restored.
- The HYCU backup controller is listed in the Virtual Machines panel, and its status is PROTECTED_DELETED.
- For recovering virtual machines, applications, and file shares, consider the following:
 - The self-service groups existing in the original data protection environment are recreated on the temporary HYCU backup controller. The recreated self-service groups do not contain any users. To restore virtual machines, applications, and file shares, you need to create users and add them to the recreated user groups that have ownership over the virtual machines and file shares you want to restore. For instructions, see [“Setting up a user environment” on page 176](#).
 - 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 80](#).
 - 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 111](#).
 - 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 129](#).

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.x.	“Restoring the HYCU backup controller to the original source” below
The HYCU backup controller to a different source by using a restore point created with HYCU version 4.0.x.	<ul style="list-style-type: none"> For restoring the HYCU backup controller protected on a Nutanix cluster to a vSphere environment: “Restoring a virtual machine from a Nutanix AHV cluster or a Nutanix ESXi cluster to a vSphere environment” on page 275 For using all other combinations of source and destination environments when restoring the HYCU backup controller: “Restoring the HYCU backup controller to a different source” on the next page
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 93
Virtual machines	“Restoring virtual machines” on page 80
Applications	“Restoring whole applications” on page 111
File shares	“Restoring file share data” on page 129



Restoring the HYCU backup controller to the original source


Use this procedure when the original cluster of the HYCU backup controller is not damaged.

Prerequisites

- The temporary HYCU backup controller has network access to the cluster of the original HYCU backup controller.
- Depending on the cluster that you plan to restore the HYCU backup controller to, a corresponding source is added to HYCU.

Procedure

1. Log on to the HYCU web user interface of the temporary 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 VM**.
 5. Select **Restore VM**, and then click **Next**.
 6. From the Select a storage container drop-down menu, select where you want to restore the HYCU backup controller.
 7. Keep the **Power virtual machine on** switch on if you want the restored HYCU backup controller to be turned on automatically after the restore. The original HYCU backup controller is deleted automatically if it still exists.
 8. Click **Restore**. The activities of the restored HYCU backup controller are suspended automatically.
 9. Log out of the HYCU web user interface.
 10. Delete the temporary HYCU backup controller from its source. For instructions, see Nutanix or VMware documentation.
 11. Log on to the HYCU web user interface of the restored HYCU backup controller.
 12. Resume the activities of the HYCU backup controller. For instructions, see [“Setting power options” on page 200](#).
 13. *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 197](#).

 **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 a different source

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 temporary HYCU backup controller has network access to the cluster you plan to restore the original HYCU backup controller to.

- Depending on the cluster that you plan to restore the HYCU backup controller to, a corresponding source is added to HYCU.


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. Log on to the HYCU web user interface.
 - c. Suspend the activities of the HYCU backup controller. For instructions, see [“Setting power options” on page 200](#).
 - 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 data” on page 155](#).
2. *Only if the original HYCU backup controller still exists.* Do one of the following:
 - Remove the HYCU backup controller from the source.
From the Nutanix Prism web console or the vSphere (Web) Client, remove the HYCU backup controller from the source. For instructions, see Nutanix or VMware documentation.
 - Ensure that the activities of the HYCU backup controller are not resumed once its clone is deployed.
 3. Log on to the HYCU web user interface of the temporary HYCU backup controller.
 4. In the Virtual Machines panel, select the original HYCU backup controller.
 5. 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.

6. Click  **Restore VM**.
7. Select **Clone VM**, and then click **Next**.
8. From the Select a storage container drop-down menu, select where you want to restore the HYCU backup controller.
9. Keep the **Power virtual machine on** switch on if you want the restored HYCU backup controller to be turned on automatically after the restore.
10. Click **Restore**. The activities of the restored HYCU backup controller are suspended automatically.

11. Log out of the HYCU web user interface.
12. Delete the temporary HYCU backup controller from its source. For instructions, see Nutanix or VMware documentation.
13. Log on to the HYCU web user interface of the restored HYCU backup controller.
14. Resume the activities of the HYCU backup controller. For instructions, see [“Setting power options” on page 200](#).
15. *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 197](#).

 **Note** Make sure to enter the original IP address of the HYCU backup controller.

Recreating HYCU instances


If you are using HYCU for protection of Nutanix Files shares, your data protection environment includes at least one HYCU instance that is connected to the HYCU backup controller. Therefore, depending on the severity of the disaster, you may need to reestablish connections of HYCU instances to the restored HYCU backup controller or recreate HYCU instances.

Affected component	Required action on HYCU instances
HYCU backup controller	Reconnect all HYCU instances that were connected to the HYCU backup controller
HYCU instance	Restore the HYCU instance


Prerequisite

A restore of the HYCU backup controller is complete and the HYCU backup controller is turned on.

Procedure

1. Log on to the HYCU web user interface.
2. Click  **Administration**, and then select **Instances**.
3. For each HYCU instance, depending on its state, do one of the following:
 - HYCU instance still exists on the source:
 - a. *Only if the HYCU instance is turned on.* From the Nutanix Prism web console, turn the HYCU instance off.
 - b. From the Nutanix Prism web console, turn the HYCU instance on. It will establish a connection to the HYCU backup controller and will be reconfigured automatically.
 - HYCU instance is corrupted or no longer exists:

- a. *Only if you want to keep the name of the HYCU instance.* In the Instances dialog box, take a note of the VM name, Hostname, Source, and IP address option values for the HYCU instance.
- b. *Only if the original HYCU instance still exists and is corrupted.* From the Nutanix Prism web console, remove the corresponding virtual machine from the source.
- c. *Only if you want to use a new name for the HYCU instance.* Delete the HYCU instance through the HYCU web user interface. For instructions, see [“Managing HYCU instances” on page 189](#).
- d. Create a new HYCU instance. It is not required that you create it on the same source as the original HYCU instance. For instructions, see [“Adding a Nutanix Files server” on page 34](#).

 **Important** The HYCU instance must be created from the same HYCU virtual appliance image (OVF package) as your HYCU backup controller.

Only if you want to keep the name of the HYCU instance. Make sure that your new HYCU instance is configured with the same name, host name, and network settings as the original HYCU instance.

If—due to changes in your data protection environment—you realize that you do not need any of the HYCU instances anymore, you can remove them. For instructions, see [“Deleting a HYCU instance” on page 191](#).

Chapter 8

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 tasks that are running in my environment and get an insight into a specific task status.	"Checking the status of jobs" on page 144
View all events that occurred in my environment.	"Viewing events" on page 146
Configure HYCU to send email notifications when events occur.	"Sending email notifications" on page 146
Obtain reports on different aspects of the data protection environment.	"Using HYCU reports" on page 147
View entity details.	"Viewing entity details" on page 152
Narrow down the list of displayed items by applying filters.	"Filtering data" on page 155
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 161
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 161
View policy information, or edit or delete a policy.	"Managing policies" on page 165
Back up data manually.	"Performing a manual backup" on page 167
Mark a restore point as expired.	"Expiring backups" on page 168


I want to...	Procedure
Archive data manually.	“Archiving data manually” on page 170
Recreate a snapshot.	“Recreating snapshots” on page 171


In case of recognized problems in the Nutanix 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 virtual machine resources” on page 171](#).

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**.

 **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 virtual and physical machines in your environment, and the exact number of protected, unprotected, and migration/DR-ready virtual and physical machines. A virtual or physical machine is considered:</p> <ul style="list-style-type: none"> Protected: If it has at least one valid backup available. Migration/DR-ready: If all backups in the current backup chain are stored on one of the cloud targets (Google Cloud or Azure) and a successful cloud readiness check was performed during its latest backup. <p>For detailed information about protecting virtual and physical machines, see “Backing up virtual machines” on page 79.</p>
Applications	<p>Shows the percentage of protected applications, and the exact number of protected and unprotected applications. An application is considered protected if it has at least one valid backup available.</p>


Dashboard widget	Description
	For detailed information about protecting applications, see “Backing up applications” on page 110 .
HYCU Controller*	Shows the resource information about the virtual machine where the HYCU backup controller resides (storage, vCPU, and memory). For details about what to do if any of these values reaches a critical value (that is, if any of the values that are indicated by circles becomes red), see “Adjusting the HYCU virtual machine resources” on page 171 .
Backups	Shows the backup job success rate for the last seven days.
Targets*	Shows the number of existing targets, overall capacity utilization, and the utilization per target type. For detailed information about setting up targets, see “Setting up targets” on page 38 .
Policies	Shows the percentage of policies that are compliant and the exact number of compliant and non-compliant policies. A policy is considered compliant if all entities to which this policy is assigned are compliant with the policy settings. For detailed information about policies, see “Defining your backup strategy” on page 54 .
Jobs	Shows the number of jobs in the data protection environment in the last 48 hours. It also shows how many jobs succeeded, failed, are in progress or in a queue. For details, see “Checking the status of jobs” below .
Events	Shows the number of events in the data protection environment in the last 48 hours. It also shows the number of the events according to their status. For details, see “Viewing events” on page 146 .

* An infrastructure group administrator only.

Checking the status of jobs

You can use the Jobs panel to check the overall status of jobs.




Accessing the Jobs panel



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

In the Jobs panel, you can do the following:

- Check processes that are currently running.
- Check 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** To minimize the Detail view, click  **Minimize** or press **Spacebar**. To return it to its original size, click  **Maximize** or press **Spacebar**.

- *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 job by selecting it, and then clicking  **Abort Job**.

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.

 **Tip** You can update the list of jobs by clicking  **Refresh**.


The following table shows the job information:

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).
Started	When a job was started.
Finished	When a job finished.

Viewing events

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

Accessing the Events panel

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

The following information is available for each event:

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 in case of 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 desired event.

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


Sending email notifications

You can configure HYCU to send email notifications when events occur.


Prerequisite


Because HYCU uses SMTP to send email notifications, an SMTP server must be configured. For details, see [“Configuring an SMTP server” on page 200](#).

Accessing the Events panel

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

Procedure

1. In the Events panel, click  **Email Notifications**.
2. In the Email Notifications dialog box, click **+ New**. The New dialog box appears.
3. In the Subject field, enter a subject of the email message.
4. 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).



 **Tip** You can select all the available categories by selecting the **Select All** check box. If you want to remove any of the selected categories, expand the list, and then

click the one that you want to remove. To remove all the selected categories, click **X Clear All**.

5. From the Status drop-down menu, select the status of the events (Success, Warning, Failed).

Tip You can select all the available statuses by selecting the **Select All** check box. If you want to remove any of the selected statuses, expand the list, and then click the one that you want to remove. To remove all the selected statuses, click **X Clear All**.

6. In the Email address field, enter one or more email recipients that will get the email notifications. If you are entering more than one email address, make sure to press the Spacebar after entering each one.

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**).

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

- View reports. For details, see [“Viewing reports” on page 149](#).
- Generate reports. For details, see [“Generating reports” on page 150](#).
- Schedule reports. For details, see [“Scheduling reports” on page 151](#).

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 152](#).

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 and physical machines, applications, and shares 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*	List of targets with the corresponding number of jobs that were running during each of the last 24 hours.
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*	Total amount of protected data per owner.
Protected data per target*	Amount of the data protected in the last 24 hours per target.
Protected data timeline per target*	Daily amount of protected data per target.
Protected VM size per target *	List of protected virtual and physical machines 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*	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.

* 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 from scratch.	<ol style="list-style-type: none"> 1. In the Reports panel, click + New. The Report Configuration dialog box opens. 2. Enter a report name and, optionally, its description. 3. Specify the time range for the report. 4. Select the type of report. 5. Select the aggregation value that you want to use to perform a calculation on a set of collected data. 6. Distribute the report tags for the collected data that you want to include in your report between x-axis and y-axis to determine how the collected data will be presented in the report. 7. Click Save.
Edit an existing report and save it as a new report.	<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 desired report, and then double-click it or click  Preview .
View a saved report.	<ol style="list-style-type: none"> 1. In the Reports panel, from the list of reports, select the desired report. 2. In the Detail view that appears at the bottom of the screen, select the desired report version, and then double-click it or 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 the previous page](#).

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.
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 149](#).
- Delete the saved reports that you do not need anymore. To do so, select the desired report version, and then click  **Delete**.


Scheduling reports



You can use scheduling to generate 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 200](#).

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


 **Tip** The reports that are generated automatically are marked by  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

1. In the Reports panel, from the list of all reports, select the one that you want to export, and then click  **Export**.
2. Click **OK**.

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, physical machine, discovered application, and file share in the Detail view of the Virtual Machines, Applications, or Shares panel. The following details are available:

Summary	Shows detailed information about the selected entity.
Restore point	You can view the following information about each restore point: <ul style="list-style-type: none"> • Date and time when the restore point was created.

	<ul style="list-style-type: none"> • Restore instances: <ul style="list-style-type: none"> ◦ BCKP Backup: Available by default unless a backup is expired. <ul style="list-style-type: none"> ■ FULL Full: Visible if a full backup was performed. ■ INCR Incremental: Visible if an incremental backup was performed. ◦ ARCH Archive: Available if a data archive was created. ◦ COPY Copy: Available if a copy of backup data was created. ◦ SNAP Snapshot: Available if the source contains a local snapshot that enables you to perform a fast restore. ◦ SNAP Partial snapshot: <i>Nutanix clusters only</i>. Available if the source contains a local snapshot that enables you to perform a fast restore. Such a snapshot contains only a partial set of disks and can be used only for restoring application items. <p>If any virtual disks were excluded from a backup, the corresponding restore instance label is marked with a red circle.</p> <p>For example, FULL.</p> <p> Important If any of the restore instances 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 a compliance status indicated by a respective 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 a restore, the estimated time required for a restore, the number of copies and failed copies (if the Copy policy option is enabled), and the backup, snapshot, archive, and/or copy expiration time. In addition, if the compliance status of your entity is Failure, this list will also include a reason why it is not compliant.</p>
Backup status	For details, see “Viewing the backup status of entities” on the next page .
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 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
 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 a 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.


If you double-click a backup status icon, you are directed to the Jobs panel where you can check details about the related jobs.

Filtering data

HYCU provides you with two types of filters that you can apply—the main filter and the detail filter. After you apply any of the filters, only data that matches the filter criteria is displayed and you can easily find what you need.

Applying the main filter

Apply the main 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** This type of filter is available in the Applications, Virtual Machines, Volume Groups, Shares, Policies, Targets, Jobs, Events, and Self-Service panels.

Procedure


1. In the selected panel, click  **Main Filter**. The Main view side panel opens.
2. Select your filter criteria.
3. Click **Apply Filters**.

See one of the following sections for the details about the available filtering options:


- [“Filtering options in the Applications panel” on the next page](#)
- [“Filtering options in the Virtual Machines panel” on page 157](#)
- [“Filtering options in the Volume Groups panel” on page 158](#)
- [“Filtering options in the Shares panel” on page 159](#)
- [“Filtering options in the Policies panel” on page 159](#)
- [“Filtering options in the Targets panel” on page 159](#)
- [“Filtering options in the Jobs panel” on page 160](#)
- [“Filtering options in the Events panel” on page 160](#)
- [“Filtering options in the Self-Service panel” on page 161](#)

Applying the detail filter

Apply the detail filter when you want to focus on the information about the restore and backup data of the selected item.



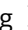

 **Note** This type of filter is available in the Applications, Virtual Machines, Volume Groups, and Shares panels.

Procedure

1. From the list of all items in the selected panel, select the item that you want to filter by restore and backup data.
2. In the Detail view that appears at the bottom of the screen, click  **Detail Filter**. The Detail view side panel opens.
3. Select your filter criteria.
4. Click **Apply Filters**.

See one of the following sections for the details about the available filtering options:

- [“Filtering options in the Applications panel” below](#)
- [“Filtering options in the Virtual Machines panel” on the next page](#)
- [“Filtering options in the Volume Groups panel” on page 158](#)
- [“Filtering options in the Shares panel” on page 159](#)

 **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.

Filtering options in the Applications panel

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

Filtering option	Action
Search	Enter a search term. You can filter by the name of the application.
Sources	From the drop-down menu, select the sources that host the virtual machines on which the applications are running or the physical machines on which the applications are running.
Policies	From the drop-down menu, select the policies that are assigned to the virtual or physical machines on which the applications are running.
Owners	From the drop-down menu, select the owners that are assigned to the virtual or physical machines on which the applications are running.

Filtering option	Action
Application types	From the drop-down menu, select the application types.
Compliance	Select one or more check boxes to filter by the compliance status.
Protection	Select one or more check boxes to filter by the protection status.
Discovery	Select one or more check boxes to filter by the application discovery status: <ul style="list-style-type: none"> • Success: One or more applications are discovered. • Failure: No applications were discovered. • Warning: Application discovery failed because the virtual or physical machine is offline or not reachable.

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

Filtering option	Action
Restore instances	From the drop-down menu, select one or more restore instances.
Restore point date	Select the time to filter by when the restore points were created.
Backup status	Select one or more check boxes to filter by the backup status.
Compliance	Select one or more check boxes to filter by the compliance status.

Filtering options in the Virtual Machines panel

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

Filtering option	Action
Search	Enter a search term. You can filter by the virtual or physical machine name, the HYCU UUID, or the source UUID.
Sources	From the drop-down menu, select the sources that host the virtual machines or the physical machines.
Credential groups	From the drop-down menu, select the credentials for the virtual or physical machines.
Policies	From the drop-down menu, select the policies that are assigned to the virtual or physical machines.
Owners	From the drop-down menu, select the owners that are assigned to the virtual or physical machines.
Compliance	Select one or more check boxes to filter by the compliance status.

Filtering option	Action
Discovery	Select one or more check boxes to filter by the application discovery status: <ul style="list-style-type: none"> • Success: One or more applications are discovered. • Failure: No applications were discovered. • Warning: Application discovery failed because the virtual or physical machine is offline or not reachable. • Undefined: Information about the application discovery status is not available.
Protection	Select one or more check boxes to filter by the protection status.
Disaster recovery readiness	Select the check box to filter by the migration/DR readiness status.

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

Filtering option	Action
Restore instances	From the drop-down menu, select one or more restore instances.
Restore point date	Select the time to filter by when the restore points were created.
Backup status	Select one or more check boxes to filter by the backup status.
Compliance	Select one or more check boxes to filter by the compliance status.

Filtering options in the Volume Groups panel

In the Main view side panel, select the following filtering option:

Filtering option	Action
Sources	From the drop-down menu, select the sources that host the virtual machines to which the volume groups are attached.

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

Filtering option	Action
Restore instances	From the drop-down menu, select one or more restore instances.
Restore point date	Select the time to filter by when the restore points were created.
Backup status	Select one or more check boxes to filter by the backup status.
Compliance	Select one or more check boxes to filter by the compliance status.

Filtering options in the Shares panel

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

Filtering option	Action
Search	Enter a search term. You can filter by the file share name.
File servers	From the drop-down menu, select the file servers that host the file shares.
Policies	From the drop-down menu, select the policies that are assigned to the file shares.
Owners	From the drop-down menu, select the owners that are assigned to the file shares.
Compliance	Select one or more check boxes to filter by the compliance status.
Protection	Select one or more check boxes to filter by the protection status of file shares.

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

Filtering option	Action
Restore instances	From the drop-down menu, select one or more restore instances.
Restore point date	Select the time to filter by when the restore points were created.
Backup status	Select one or more check boxes to filter by the backup status.
Compliance	Select one or more check boxes to filter by the compliance status.

Filtering options in the Policies panel

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

Filtering option	Action
Search	Enter a search term. You can filter by the name of the policy.
Compliance	Select one or more check boxes to filter by the compliance status.

Filtering options in the Targets panel

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

Filtering option	Action
Search	Enter a search term. You can filter by the name of the target.

Filtering option	Action
Target type	Select one or more check boxes to filter by the target type.
Health	Select one or more check boxes to filter by the health of the target.

Filtering options in the Jobs panel

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

Filtering option	Action
Search	Enter a search term. You can filter by the job name or the job UUID.
Status	Select one or more check boxes to filter by the status of the job.
Time range	Specify a time range to limit your search for jobs. You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, or Last week), or use the calendar to select a start date and hour and an end date and hour of the time range for jobs to be displayed.

Filtering options in the Events panel

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

Filtering option	Action
Message	Enter a text string to filter the list to include only the messages with the specified string.
Category	Enter a text string to filter the list to include only the categories with the specified string.
Username	From the drop-down menu, select the user name.
Status	Select one or more check boxes to filter by the status of the event.
Time range	Specify a time range to limit your search for events. You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, or Last week), or use the calendar to select a start date and hour and an end date and hour of the time range for events to be displayed.

Filtering options in the Self-Service panel


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

Filtering option	Action
Group name	Enter the group name.
Status	Select one of the following to filter by the status of the group or user (that is, which groups or users are allowed to log on to HYCU and which are not).


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  **Main Filter**, 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.
Type	Type of target (NFS, SMB, Nutanix, iSCSI, AWS S3/Compatible, AZURE, Google Cloud, QStar NFS, or QStar SMB). <div>  Note A tape target is represented by the  icon. </div>
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 virtual machine 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 or physical 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	Estimation of the amount of storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
Utilization	Percentage of the specified target size that is already used for storing protected data.
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

Target information	Description
	<p>data.</p> <ul style="list-style-type: none"> 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 before you restore all the required data. 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 the next page.</p>

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


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

- In the Targets panel, select the target that you want to edit, and then click  **Edit**. The Edit Target dialog box appears.
- Edit the selected target as required. For detailed information about target properties,

see [“Setting up targets” on page 38](#).

 **Important** If you want to change the NFS or SMB server name, IP address, or path to the shared folder, or the portal IP address of an iSCSI target, see [“Detaching storage and modifying target data” below](#).

3. Click **Save**.

Detaching storage and modifying target data



If you want to change the NFS or SMB server name, IP address, or path to the shared folder, or the portal IP address of an iSCSI target, you must make sure that the storage is detached from the HYCU backup controller to be able to perform the required modifications.

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 modifications 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 the previous page](#).
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. *For NFS, SMB, and iSCSI targets:* If you are deactivating the target to change the NFS or SMB server name, IP address, or path to the shared folder, or the portal IP address of an iSCSI target, enable the **Detach storage** switch. For details on detaching storage from the HYCU backup controller, see [“Detaching storage and modifying target data” above](#).
4. *For target deactivation:* Click **Yes** 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.


Increasing the size of an iSCSI target

HYCU enables you to increase the size of your iSCSI target by extending the HYCU logical volume.

Prerequisites

- The size of the target has been increased on the iSCSI server.
- No backup or restore job is in progress on the selected target.
- No other maintenance task is 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 has already been started.

Procedure


1. In the Targets panel, select the target whose size you want to increase, and then click  **Extend**.
2. Click **Yes** to confirm that you want to increase the size of the selected target.


You will receive a message that indicates whether increasing the size of the iSCSI target completed successfully.

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 **Yes** 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.

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 152.</p>
VM Count	Total number of virtual and physical machines 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 desired 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


Consideration

When editing a policy that is assigned to several virtual machines, one of which is the HYCU backup controller, make sure that the policy remains applicable also for protecting the HYCU backup controller. For more information, see [“Preparing for disaster recovery” on page 69](#).

Procedure

1. In the Policies panel, select the policy that you want to edit, and then click  **Edit**. The Edit Policy dialog box appears.
2. Edit the selected policy as required. For detailed information about policy properties,

see [“Creating a policy” on page 56](#).

 **Important** *For vSphere environments:* You cannot edit the policy that is assigned to the vSphere virtual machines or applications in such a way that you enable the Backup from replica or Fast restore option. These options are not available for vSphere virtual machines or applications.


3. Click **Save**.

Deleting a policy

Consideration

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 **Yes** 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).

Consideration

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 264](#).

Procedure

1. In the Virtual Machines, Applications, or Shares panel, select which entities you want to back up.
2. Click  **Backup** to perform the backup of the selected entities.
3. Use 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 amount of changed data.
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.

Expiring backups

HYCU expires backups automatically according to the retention period that is set for the backup data in the policy. However, if there is a restore point (a backup) that you do not want to use for restoring data anymore, you can at any time expire it manually.

A restore point represents data that was backed up at a specified point in time. Your restore point can contain one or more restore instances—Backup, Copy, Snapshot, and Archive—that can be marked as expired also individually.

 **Tip** You can check the backup, copy, snapshot, and/or archive expiration time in the Detail view of the Virtual Machines, Applications, or Shares panel. For details, see [“Viewing entity details” on page 152](#).

Considerations

- If the most recent restore point is marked as expired, the next backup will be a full backup.
- When a restore point is marked as expired, any subsequent incremental backups within the same backup chain will also be marked as 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.
- The Backup and Copy restore instances are always expired together.

Expiring backups automatically

When any of the restore instances reaches its retention period, it is grayed out in the HYCU web user interface. Such restore instances are expired when the last restore instance in the backup chain reaches its retention period. This means that this data is not removed from HYCU or the target until all the restore instances in the backup chain are expired. However, if there is a restore point that contains the Archive restore instance, 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.

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 backups manually




You can mark as expired one of the following:

- Whole restore point:
Make sure that all restore instances are marked for expiration.
- One or more restore instances:
Make sure that only restore instances that you want to expire are marked for expiration.


Considerations

- An expire action cannot be undone.
- If you mark Backup and Copy for expiration, the associated snapshot is also expired, if there is one.

Depending on the entity for which you want to expire old backups, 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**.

Procedure

1. In the Virtual Machines, Applications, or Shares panel, select the entity for which you want to expire old backups.
2. In the Detail view that appears at the bottom of the screen, select the restore point that you want to mark as expired.
3. Click  **Expire**. The Expire dialog box appears.
4. Select the restore instances that you want to mark as expired:
 - **Backup and Copy**
 - **Snapshot**
 - **Archive**

The restore instances that are available for expiration are based on the options that you set in your policy. By selecting all the restore instances, you mark the whole restore point as expired.

5. Click **Yes** to confirm that you want the selected restore instances to be marked as expired.

 **Note** If you mark the whole restore point as expired, the backup status is

shown as Expired (🕒). This indicates that the restore point cannot be used for restoring data anymore.

The HYCU cleaning process removes the expired backups from the target.

Archiving data manually

HYCU archives your data automatically once 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 have the Administrator or Backup Operator user role assigned.
- The Archiving option is specified in the assigned policy and a data archive is created.

Consideration


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.


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**.

Procedure

1. In the Applications, Virtual Machines or Shares 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 desired 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 appears.
4. Select the desired archiving option.
5. Click **Run**.

Recreating snapshots


If you plan to restore individual files from a snapshot (and not directly from a target) and no snapshot is available for the selected virtual machine restore point, you can recreate it manually.


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 desired 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 appears.
4. From the Storage container drop-down menu, select where you want to recreate the snapshot.

 **Important** *For vSphere environments:* Because restoring individual files from the disks that reside on the vVol datastores is not supported, only the available VMFS or NFS datastores are shown.

5. From the Restore instance drop-down menu, select which restore instance you want to use for recreating the snapshot. Your restore point can contain one or more restore instances among which you can select:
 - **Automatic:** This option ensures the fastest snapshot creation.
 - **Backup**
 - **Copy**
 - **Archive**
6. Click **Recreate**.

Adjusting the HYCU virtual machine 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. To adjust the HYCU virtual machine resources, follow these steps:

1. Log on to Nutanix Prism. For details about the Prism web console, see Nutanix documentation.
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 your HYCU virtual machine, and then click **Power Off Actions** to shut down the virtual machine.
 **Important** Wait a moment for the virtual machine to shut down completely.
5. Click **Update**, and then, in the Update VM dialog box, modify the configuration as required, and click **Save**.
6. Click **Power on** to turn on the virtual machine.


Chapter 9

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 (HYCU users) or through an Active Directory server (AD users and AD groups).



Each user that logs on 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 the next page](#).

 **Note** User management concepts and procedures apply to both virtual and physical machines.

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 group	<p>Must be created by an infrastructure group administrator and represents a customer or 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 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 administrator, you can manage only the group you belong to. The following diagram shows which user-related actions you can perform:

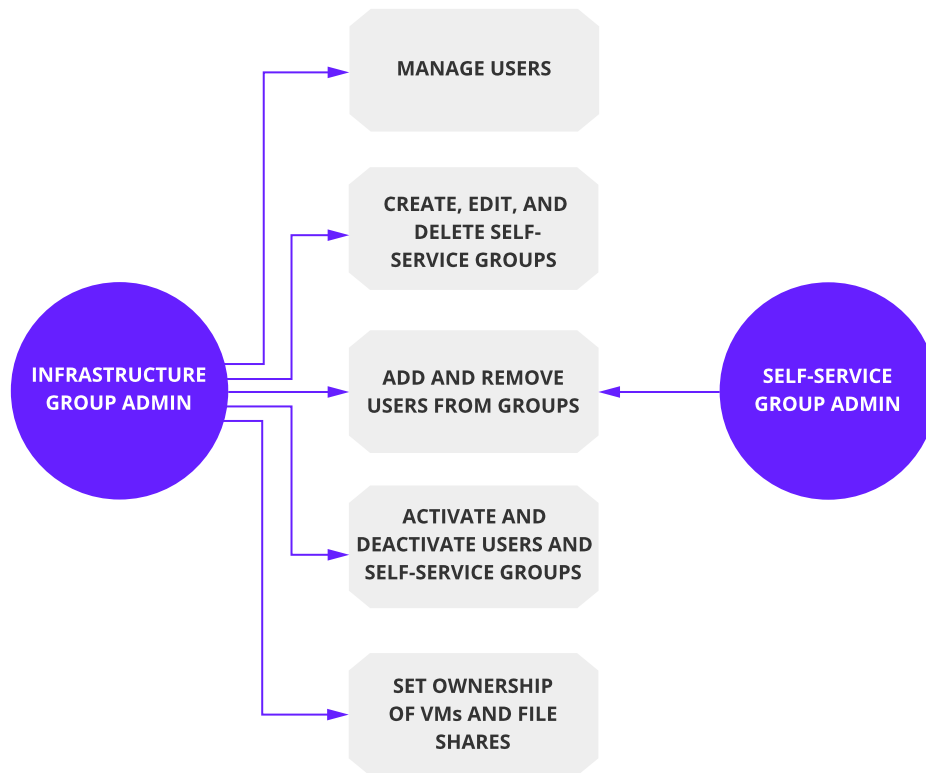



Figure 9-1: User management actions performed by the infrastructure and self-service group administrators

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 logged on 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 user management actions (creating, editing, and deleting users and self-service groups, setting ownership of virtual machines and file shares, enabling and disabling users and self-service groups to access HYCU, and adding and removing users from groups). • Perform all administrative actions in HYCU. • Perform all report management actions. 	<ul style="list-style-type: none"> • Assign policies. • Back up and restore virtual machines, applications, and file shares. • Add and remove users from groups. • Perform all report management actions. • Add, edit, and remove cloud accounts.
Viewer	<ul style="list-style-type: none"> • View information about applications, virtual machines, shares, 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, shares, 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 and file shares 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, and file shares.
Restore Operator	<ul style="list-style-type: none"> • View the same information as Viewer. • Restore virtual machines and file shares that are not owned by any self-service group, and restore applications. 	<ul style="list-style-type: none"> • View the same information as Viewer. • Restore virtual machines, applications, and file shares.

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" below
2. Add a user to a user group.	An infrastructure or a self-service group administrator	"Adding a user to a group" on page 178

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 178
Set ownership of virtual machines and file shares.	An infrastructure group administrator	"Setting ownership of virtual machines" on page 179 "Setting ownership of file shares" on page 180
Enable or disable specific groups or users from logging on to HYCU.	An infrastructure group administrator	"Activating or deactivating users or self-service groups" on page 180

Accessing the Self-Service panel

To access the Self-Service panel, in the navigation pane, click  **Self-Service**.

Creating a user


Prerequisite


For using Active Directory for authentication: Active Directory authentication is configured. For details on how to do this, see ["Configuring Active Directory authentication" on page 184](#).

Limitation

You cannot add the Active Directory primary group (usually the Domain Users group) as an AD group.

Procedure

1. In the Self-Service panel, click  **Manage Users**, and then click **+ New**. The Manage Users dialog box opens.
2. Enter a user name if you are adding a HYCU user or an AD user, or a common name if you are adding an AD group.


 **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 264](#).

3. From the Authentication type drop-down menu, select one of the following authentication types:

- **HYCU**

Enter a display name, the user password and, optionally, email address.

 **Note** The minimum password length is six characters.

- **AD user**



From the Active Directory drop-down menu, select the Active Directory the AD user belongs to.

- **AD group**

From the Active Directory drop-down menu, select the Active Directory the AD group belongs to.

4. Click **Save** followed by **Close**. The user is added to the list of all users.

You can later do the following:

- Edit any of the existing HYCU users by clicking  **Edit User** 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 logging on to HYCU. For details, see [“Activating or deactivating a user” on page 180](#).
- Delete any of the existing users by clicking  **Remove User**. 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. A self-service group is created. For details on how to do this, see [“Creating a self-service group” below](#).

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 174](#).
- 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 > 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**. The Add User to Group dialog box opens.

 **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**, **Viewer**, **Backup Operator**, or **Restore Operator**).
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

Procedure

1. In the Self-Service panel, click **+ New Group**. The New Group dialog box opens.
2. Enter a self-service group name and, optionally, its description.
3. Click **Save**.

You can later do the following:

- Add users to groups. For details, see [“Adding a user to a group” on the previous page](#).
- 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 and the Exclude policy (alongside of other policies already assigned to the virtual machines 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 264](#).
- Enable or disable specific self-service groups from logging on to HYCU. For details, see [“Activating or deactivating a self-service group” on page 181](#).
- Delete any of the existing self-service groups by clicking  **Delete**.


Setting ownership of virtual machines

By setting ownership of virtual machines, you enable specific groups to protect only the assigned virtual machines.


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 such virtual machines, and then clicking  **Owner** followed by **Unassign**.


Setting ownership of file shares

By setting ownership of file shares, you enable specific groups to protect only the assigned 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**.


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 logging on 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 logging on 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 log on 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 logged on 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 logged on to HYCU at the upper right of the screen.



Figure 9–2: Example of a self-service group, HYCU_group, under which a user, HYCU_group_member, is logged on to HYCU

2. From the list of all groups to which you belong, select the one to which you want to switch.

 **Tip** The group under which you are currently logged on to HYCU has  next to it.

3. Click **Switch**.

You are automatically switched to the group you selected.


Updating your user profile

As the currently logged-on user, you can edit your name and email address 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 176](#).

Procedure

1. Click  at the upper right of the screen, and then select **Update Profile**. The Update profile dialog box opens.
2. In the Name field, specify a new name.
3. In the Email field, enter the email address that you want to be associated with your user profile.
4. Click **Save**.

Chapter 10

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
Configure Active Directory authentication.	"Configuring Active Directory authentication" on the next page
Add cloud accounts to HYCU.	"Adding a cloud account" on page 185
Configure encryption for targets.	"Configuring target encryption" on page 189
Manage HYCU instances.	"Managing HYCU instances" on page 189
Set the iSCSI Initiator secret.	"Setting the iSCSI Initiator secret" on page 191
Obtain a permanent HYCU license.	"Licensing" on page 192
Configure log file settings to troubleshoot problems if HYCU does not perform as expected.	"Setting up logging" on page 195
Change network settings or enable network bandwidth throttling.	"Configuring your network" on page 197
Set power options.	"Setting power options" on page 200
Configure an SMTP server.	"Configuring an SMTP server" on page 200
Upgrade HYCU to a new available version.	"Upgrading HYCU" on page 204
Apply a HYCU hotfix.	"Applying HYCU hotfixes" on page 216
Configure the SSL certificate.	"Configuring SSL certificates" on page 201
Share telemetry diagnostic data with HYCU.	"Sharing telemetry data with HYCU" on page 203

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 220](#).

Configuring Active Directory authentication

In addition to standard HYCU authentication, you can also configure Active Directory authentication by adding one or more Active Directories as your authentication sources in HYCU. This allows users to log on to the HYCU web user interface with their Active Directory domain accounts or, if certificate authentication is enabled, with a client certificate or a smart card.

For details on how to enable certificate authentication, see [“Enabling certificate authentication” on the next page](#).

Prerequisite

For using LDAPS for user authentication: LDAPS authentication is set up. For details, see [“Setting up LDAPS authentication” on page 228](#).

Accessing the Active Directory dialog box

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

Procedure

1. In the Active Directory dialog box, click **+ New**. The New dialog box appears.
2. In the Name field, specify a name for the Active Directory.
3. 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.

Example

If you enter `mycompany.com` as the FQDN and `mc` as the alias domain name, the user will be able to log on to HYCU with `<Username>@mycompany.com` or `mc\<Username>`.


You can enter more than one FQDN or domain alias name. In this case, press the Spacebar after entering each one.

4. In the Provider URL field, enter the URL of the corresponding LDAP server in one of the following formats:

- `ldap://<LDAPServerHostnameorIPAddress>:<Port>`

When using the LDAP protocol, the default port is 389. Entering the port is optional if the default value is used.



- *Only if LDAPS authentication is set up.* `ldaps://<LDAPServerHostname>:<Port>`

 **Important** Make sure that the LDAP server hostname 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.

When using the LDAPS protocol, the default port is 636. Entering the port is optional if the default value is used.

You can enter more than one URL. In this case, press the Spacebar after entering each one.

5. *Only if you plan to enable certificate authentication.* Enable the **Use service account** option, and then enter the user name and password of the service account that HYCU will use to log on to the Active Directory and authorize users.
6. Click **Save**.

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

After Active Directory authentication is configured, the AD user or AD group authentication type can be specified when creating a new user. For details, see [“Creating a user” on page 176](#).

Enabling certificate authentication


By enabling certificate authentication, you allow users to log on 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 is added to HYCU.
- A CA-signed certificate is imported to HYCU. For details on how to do this, see [“Importing a custom certificate” on page 202](#).

Procedure

1. In the Active Directory 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.

 **Important** When you enable or disable certificate authentication, all affected users that are logged on the HYCU web user interface will lose their connections and will be required to log on again.

Adding a cloud account

You must add one or more cloud accounts to HYCU before performing any of the following data protection tasks:

- Storing data to a Google Cloud target.
- Migrating data protected with HYCU from the on-premises environment to cloud.
- Migrating data protected with HYCU Data Protection as a Service for GCP (HYCU for GCP) or HYCU Data Protection as a Service for Azure (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 GCP and HYCU for Azure data protection environments in HYCU Manager.

Consideration

Migrating virtual machines across different infrastructures, performing disaster recovery of data to cloud, and monitoring cloud data protection environments are supported only if you own a HYCU Protégé license and have an active subscription for HYCU for GCP or HYCU for Azure.

Depending on whether your data is protected with HYCU for GCP or HYCU for Azure, add one or more cloud accounts to HYCU:

Data is protected with...	Cloud account	Instructions
HYCU for GCP	Service account	“Adding a Google Cloud Platform service account” below
HYCU for Azure	Service principal	“Adding an Azure service principal” on page 188

Adding a Google Cloud Platform service account

The type of Google Cloud Platform service account that you add to HYCU depends on what data protection tasks you want to perform:

I want to...	Service account to add
Store data to a Google Cloud target.	An account that has access to the buckets where you want to store your backup data.
Migrate data protected with HYCU for GCP from Google Cloud Platform to the on-premises environment.	An account that is imported to HYCU for GCP 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 Platform.	An account that is imported to HYCU for GCP and has the Storage Admin and Compute Admin roles assigned on the projects where you want to migrate your virtual machines.
Perform disaster recovery of data to Google Cloud Platform in the event of a disaster.	An account that is imported to HYCU for GCP and has the Storage Admin and Compute Admin roles assigned on the project where you want to perform a disaster recovery.
Monitor my HYCU for GCP 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 is configured in the Google Cloud Platform service suite.
- The following APIs are enabled on the Google Cloud Platform 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 them, see Google Cloud documentation.

- The service account is granted the following roles in the Google Cloud Platform service suite: Compute Admin (roles/compute.admin), Storage Admin (roles/storage.admin), and Service Account User (roles/iam.serviceAccountUser) on the project with your protected instances.
- You 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 Selection dialog box, click **Create a GCP cloud account**, and then click **Next**. The GCP Authentication dialog box appears.
2. Browse for the JSON file with the service account information. In the Service account authentication field, the file name is displayed.
3. In the Name field, you can change the account service name.
4. Click **Upload**.

After you are notified about a successful service account upload, its name appears in the Cloud Accounts dialog.

5. Click **Close**.

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 is created in Azure and added to HYCU for Azure. For details, see HYCU for Azure documentation.

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.	<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
Migrate data protected with HYCU from the on-premises environment to Azure.	
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 **+ New**. The Cloud Selection dialog box appears.
2. Select **Add Azure service principal**, and then click **Next**. The Azure Authentication dialog box appears.
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 ID of the application's (HYCU for Azure) registration in the Azure Active Directory.
6. In the Secret key field, enter the secret that is associated with the application ID.
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.

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**.

You are notified about a successful import of the encryption key.

Managing HYCU instances

All existing HYCU instances in your data protection environment are listed in the 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, and delete HYCU instances.

For details on HYCU instances, see [“HYCU instances” on page 36](#).

Accessing the Instances dialog box

To access the Instances dialog box, click  **Administration**, and then select **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 is present on the Nutanix cluster in the following format:

`hycu-<Version>-<Revision>`

For example, `hycu-4.2.0-3634`.


- *For creating a HYCU instance on a Nutanix ESXi cluster:*
 - A user with specific privileges for vCenter Servers is specified. For details on which privileges must be assigned to a vSphere user, see [“Assigning privileges to a vSphere user” on page 233](#).
 - The HYCU OVF package is imported to the vCenter Server content library and its format is as follows:

`hycu-<Version>-<Revision>`

For example, `hycu-4.2.0-3634`.


Procedure

1. In the Instances dialog box, click **+ New**. The New dialog box opens.
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.
4. In the Deployment section, do the following:
 - a. From the Destination drop-down menu, select a Nutanix cluster on which your HYCU instance will reside.
 - b. From the Network drop-down menu, select a VLAN.

- c. From the Datastore drop-down menu, select a datastore.

 **Tip** If you select **Select automatically**, HYCU will select the datastore with the most available space.

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.
Source	Nutanix cluster on which the HYCU instance resides (visible only if it is added to HYCU).
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, hycu-4.2.0-3634).
IP address	IP address currently assigned to the HYCU instance.

Deleting a HYCU instance

Procedure

1. In the Instances dialog box, from the list of HYCU instances, select the one that you want to delete, and then click  **Delete**.
2. In the Remove Instance dialog box, click **Yes** 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 38](#).

Accessing the iSCSI Initiator dialog box

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

To set the iSCSI Initiator secret, follow these steps:

1. In the iSCSI Initiator dialog box, enter the secret.
2. Click **Save**.

Licensing

After you deploy the HYCU virtual appliance, you can start using HYCU immediately with a prebuilt Instant-on license. This license expires automatically after 45 days and cannot be reused. Therefore, make sure to obtain a permanent license within this 45-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 all sources that you plan to protect by using HYCU.
 - VM-based licenses

Licenses are based on the number of protected virtual machines on all sources and physical machines that you plan to protect by using HYCU.
- Nutanix Files 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 server that you plan to protect by using HYCU resides.
 - Capacity-based licenses

Licenses are based on the capacity of Nutanix Files shares, which is calculated automatically as an overall size (in terabytes) of all protected Nutanix Files shares.
- HYCU Protégé 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 GCP 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.
- *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” below](#).
3. Request and obtain licenses from the web licensing portal. For details, see [“Requesting and retrieving licenses” on the next page](#).
4. Activate the licenses to start using HYCU. For details, see [“Activating licenses” on the next page](#).

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 31](#).

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
PID nutanixbackup
ND C0F90A56-3FCC-4437-A49C-EFBA9B
NRP 3
QTY 127
AFS 3
AFSCAP 4
VER V1N
HSUD FA8A5061C61F6BA5CE5A9B2C007EE
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 **Sign in**, 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 perpetual 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 type (Standard licenses, Nutanix Files licenses, or both) 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
PID nutanixbackup
ND C0F90A56-3FCC-4437-A49C-EFBA9BD8FC0F
NRP 3
AFSCAP 5
EXP 02.08.2020
VER V1N
LK D29CB215357FED55304012B02143CA9437ED5D8FC556
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 grown environment. Contact your HYCU Sales representative.

You can check the following information related to licensing:

- License type
- Backup controller ID
- Status
- Actual and licensed number of sockets
- Licensed number of sockets for Nutanix Files
- Actual and licensed Nutanix Files capacity
- Actual and licensed number of protected virtual and physical machines

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 Customer Support: Sharing telemetry data with HYCU is enabled.

For instructions, see [“Sharing telemetry data with HYCU” on page 203](#).

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 file 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 file to HYCU Customer 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 10 MiB, whereas the maximum log file size is 10 GiB.
Number of log files	The number of log files. The default number is 9.
Level	The following logging levels are available: <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>	The following levels are available: <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>	The following levels are available: <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 then 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 the next page](#)

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 logged out automatically and your session will restart.


Limitation

Multiple network adapters on the same network are 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 desired network, select this network, and then click  **Set Main**.


Procedure

1. In the Networks dialog box, 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. Change the IP address, the gateway, the domain name, the netmask, and the DNS server as required.


3. *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.

4. *Only if the Enable listening on this port switch is enabled.* Do the following:
 - a. 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.

- b. From the SSL certificates drop-down menu, select the SSL certificate that you want to use for this network. If the appropriate certificate is not on the list, you can import or generate a needed certificate by clicking **Manage**. For instructions on how to generate and import SSL certificates, see [“Configuring SSL certificates” on page 201](#).

 **Note** If the Enable listening on this port switch is disabled, you can also specify the SSL certificate that you want to use for this network.

5. 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**. The New dialog box appears.
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. *Optional.* From the Throttling window drop-down menu, select the throttling window that you want to be used for limiting bandwidth. You can also create a new throttling window or edit existing ones by clicking **Manage**. For details on how to create a throttling window, see [“Creating a throttling window” below](#).



Important

If you define multiple sites with the same IP addresses, make sure the throttling windows 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 **Windows**. The Throttling Window dialog box appears.
2. Click **+ New**. The New dialog box appears.
3. 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 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 week days 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**).

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	<p>Pauses all HYCU backup controller activities.</p> <p>If you want the HYCU backup controller activities to automatically resume after a specified amount of time, in the 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 jobs.</p>
Suspend Cleanup	Pauses the cleanup of targets. The snapshot cleanup is not affected.
Resume	Allows HYCU backup controller activities to continue.


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 is imported to HYCU. For details on how to do this, see [“Securing SMTP connections” on page 228](#).

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. Click **Save**.

You can now configure HYCU to send email notifications. For details on how to do this, see [“Sending email notifications” on page 146](#).

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, such as the certificate name, the certificate common name, the certificate expiry date, and the certificate key size.

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 197](#).

Recommendation

It is recommended to replace the self-signed certificate that is generated automatically during HYCU deployment with a CA-signed certificate.

Procedures


Depending on whether you want to create a self-signed certificate or import a custom certificate to HYCU, see one of the following sections:

- [“Creating a self-signed certificate” below](#)
- [“Importing a custom certificate” below](#)

Creating a self-signed certificate

Procedure

1. In the SSL Certificates dialog box, click **Generate**. The Generate dialog box appears.
2. Provide the following certificate-related information:
 - Name
 - Common name
 - Organization
 - Organization unit
 - Location
 - Country
 - Key size

 **Important** The maximum number of characters in each field is 64.

3. 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 three years and that you must maintain the validity of the certificate.

Importing a custom certificate

Prerequisites

- The certificate is compliant with the PKCS#7 standard and encoded in the PEM format.
- All certificate files are unencrypted.
- *For importing an SSL key pair:* The private key and the certificate are available.
- *For importing a CA-signed certificate:* The CA-signed certificate or trust chain certificates are available.

Consideration



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.

Procedure

1. In the SSL Certificates dialog box, click **Import**. The Import dialog box appears.
2. Depending on whether you want to import an SSL key pair or a CA-signed certificate, click one of the following tabs and follow the instructions:

Tab	Instructions
SSL keypair	<ol style="list-style-type: none"> a. Enter a name for your certificate. b. Browse for the following files: <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 RSA algorithm and as a PEM file in PKCS#1 or PKCS#8 format. The recommended private key sizes are 2048 and 4096 bits.</p>
CA certificate/chain	<ol style="list-style-type: none"> a. Enter a name for your certificate. b. Browse for the file with the CA-signed certificate or trust chain certificates.

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), and other detailed information on your specific infrastructure, and then sends this data to

HYCU Customer 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 have a valid HYCU Customer Support user account.

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 Customer Support. Later, the telemetry diagnostic data is sent to HYCU Customer 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 Customer Support. For more information, see [“Setting up logging” on page 195](#).

Upgrading HYCU

You can upgrade HYCU when a new software release version is available.

Prerequisites

- The source where the HYCU backup controller resides is added to HYCU.
- The HYCU backup controller activities are suspended. For instructions on how to achieve this, see [“Setting power options” on page 200](#).
- Jobs that you do not want to be aborted are finished (the upgrade process aborts all currently running jobs).

- The HYCU data disk is 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 232](#).

Considerations

- *For Nutanix clusters:* 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. The old HYCU backup controller (virtual machine) will remain on the Nutanix cluster and will be renamed to `<HYCUBackupControllerName>_version_<OldHYCUVersion>`. You can safely delete it and remove it from Nutanix protection domain after a successful upgrade.
- Any users that have been logged on to the HYCU web user interface of the HYCU virtual machine 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 hotfix packages from the hotfix directory on the HYCU virtual machine.
- *For Nutanix ESXi clusters with AOS version 5.15 or later:* After upgrading HYCU, the first backup of virtual machines and the applications running on them will be full.
- *For S3-compatible targets:* After upgrading HYCU, if you want to provide secure HTTPS access, make sure the required CA-signed certificate is imported as follows:
 1. Open a remote session to the HYCU backup controller virtual machine:

```
ssh hycu@<HYCUBackupControllerIPAddress>
```

When requested, enter the default password.

2. Import the required CA certificate:

```
keytool -importcert -keystore /etc/pki/ca-trust/extracted/java/cacerts
-file <CertificatePathname>
```

```
keytool -importcert -keystore /etc/pki/cert-templates/cacerts.template
-file <CertificatePathname>
```

Procedures

- [“Upgrading HYCU on a Nutanix AHV cluster” on the next page](#)
- [“Upgrading HYCU on a Nutanix ESXi cluster” on page 208](#)
- [“Upgrading HYCU in a vSphere environment” on page 212](#)

Upgrading HYCU on a Nutanix AHV cluster

Prerequisites

- A snapshot of the HYCU backup controller is created by using the Nutanix protection domain. For details, see Nutanix documentation.
- The HYCU system disk is 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 is ACTIVE in the Nutanix Prism image service.

For details, see Nutanix documentation.

Consideration

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-4.2.0-3634`.

Otherwise, follow the HYCU upgrade procedure to perform the HYCU instance upgrade.

Procedure

1. Log on 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-4.2.0-3634`

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. Log on 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.
 - c. From the list of the available versions, select the one to which you want to upgrade HYCU.
-  **Note** You can also check whether any newer version is available on the HYCU Customer Support portal by clicking the **Check for new version** link.
- 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 191](#).
 - 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 190](#).

You will be logged 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 log on to the HYCU web user interface.

 **Important** Before you log on to the HYCU web user interface again, make sure to perform a hard reload of its webpage in your web browser.

After you make sure HYCU was upgraded successfully, you can safely delete the old HYCU backup controller virtual machine from the Nutanix AHV cluster.

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	Instructions
By importing the HYCU OVF package to a content library.	“Upgrading HYCU by importing the HYCU OVF package to a content library” below
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 210

If HYCU is used for file share protection, the HYCU instances that are connected to your HYCU backup controller must also be upgraded. For details, see [“Upgrading HYCU instances” on page 212](#).


Prerequisites

- A snapshot of the HYCU backup controller is created by using the Nutanix protection domain. For details, see Nutanix documentation.
- Any HYCU snapshots created by using VMware vSphere are removed.

Consideration

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

 **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. Log on 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.
------------	--


Local file	<p>Browse your file system for the HYCU OVF package.</p> <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-4.2.0-3634.

2. Log on to the HYCU web user interface, and then do as follows:
 - a. Click  **Administration**, and then select **Software Upgrade**.
 - b. *Only if the credentials for the vCenter Server are not provided.* In the Software Upgrade dialog box, do the following:
 - i. Enter the name of the vCenter Server registered with the Nutanix ESXi cluster on which the HYCU backup controller resides.
 - ii. Enter the user name and password of a vSphere user with the required upgrade privileges. For details on upgrade privileges, see ["Assigning privileges to a vSphere user" on page 233](#).
 - iii. Click **Next**.
 - c. Check the current version of HYCU and all available versions, and then, from the list of the available versions, select the one 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. *Only if database optimization has not been run yet.* Click **Database Optimization**, and then click **Yes** to confirm that you want to run the database optimization job.

 **Note** The database optimization job may take several minutes.
 - e. Click **Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

You will be logged 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 log on to the HYCU web user interface.

 **Important** Before you log on to the HYCU web user interface again, make sure to

perform a hard reload of its webpage in your web browser.


After you make sure HYCU was upgraded successfully, you can safely delete the old HYCU backup controller virtual machine from the Nutanix ESXi cluster.

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. Log on 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.
Local file	Browse your file system for the HYCU OVF package.  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**.

- 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-4.2.0-3634.

- 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. Log on to the HYCU web user interface, and then do as follows:
 - a. Click  **Administration**, and then select **Software Upgrade**.
 - b. *Only if the credentials for the vCenter Server are not provided.* In the Software Upgrade dialog box, do the following:
 - i. Enter the name of the vCenter Server registered with the Nutanix ESXi cluster on which the HYCU backup controller resides.
 - ii. Enter the user name and password of a vSphere user with the required upgrade privileges. For details on upgrade privileges, see [“Assigning privileges to a vSphere user” on page 233](#).
 - iii. Click **Next**.
 - c. Check the current version of HYCU and all available versions, and then, from the list of the available versions, select the one 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. *Only if database optimization has not been run yet.* Click **Database Optimization**, and then click **Yes** to confirm that you want to run the database optimization job.

 **Note** The database optimization job may take several minutes.
 - e. Click **Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

You will be logged 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 log on to the HYCU web user interface.

 **Important** Before you log on to the HYCU web user interface again, make sure to perform a hard reload of its webpage in your web browser.

After you make sure HYCU was upgraded successfully, you can safely delete the old HYCU backup controller virtual machine from the Nutanix ESXi cluster.

Upgrading HYCU instances


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-4.2.0-3634`.

Otherwise, upgrade the HYCU instances manually as follows:

1. Remove the existing HYCU instances. For details on how to do this, see [“Deleting a HYCU instance” on page 191](#).
2. 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 190](#).

 **Note** If you made any changes to the default user credentials, after the HYCU instance upgrade, you can use only the default operating system user credentials:

User name: **hycu**

Password: **hycu/4u**

Later you can 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	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 214


Prerequisites

- As a vSphere user, you have the required upgrade privileges. For details on upgrade privileges, see [“Assigning privileges to a vSphere user” on page 233](#).
- *For importing the HYCU OVF package to a content library:* A content library is 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.
- After you upgrade HYCU, in some vSphere environments you might get an error message that there is a MAC address conflict. You can safely ignore this message.
- 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 Customer Support](#).

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. Log on 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.
Local file	Browse your file system for the HYCU OVF package.  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**.

- 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-4.2.0-3634.

2. Log on 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, check the current version of HYCU and all available versions.
 - c. From the list of the available versions, select the one 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. *Only if database optimization has not been run yet.* Click **Database Optimization**, and then click **Yes** to confirm that you want to run the database optimization job.

 **Note** The database optimization job may take several minutes.
 - e. Click **Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

You will be logged 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_<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 log on to the HYCU web user interface.

 **Important** Before you log on to the HYCU web user interface again, make sure to perform a hard reload of its webpage in your web browser.

After you make sure HYCU was upgraded successfully, you can safely delete the old HYCU backup controller virtual machine from the vSphere environment.


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. Log on 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.
Local file	Browse your file system for the HYCU OVF package.  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**.

- 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-4.2.0-3634.

- 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. Log on 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, check the current version of HYCU and all available versions.
- c. From the list of the available versions, select the one 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. *Only if database optimization has not been run yet.* Click **Database Optimization**, and then click **Yes** to confirm that you want to run the database optimization job.


 **Note** The database optimization job may take several minutes.

- e. Click **Upgrade**, and then click **Yes** to confirm that you want to upgrade HYCU.

You will be logged 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_<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 log on to the HYCU web user interface.

 **Important** Before you log on to the HYCU web user interface again, make sure to perform a hard reload of its webpage in your web browser.

After you make sure HYCU was upgraded successfully, you can safely delete the old HYCU backup controller virtual machine from the vSphere environment.

Applying HYCU hotfixes

After you receive a HYCU hotfix from HYCU Customer Support, you can apply it to your current product version. A hotfix can be applied only to an installed compatible product version. For example, a hotfix labeled 1.2.3-4567 can be applied to the product version 1.2.3 whereas a hotfix labeled 1.2.4-5678 cannot.


 **Note** Each HYCU hotfix addresses a cumulative set of issues.

Prerequisites

- *For applying a hotfix to a HYCU backup controller:* The HYCU backup controller activities are suspended. For instructions on how to do this, see [“Setting power options” on page 200](#).
- Jobs that you do not want to be aborted are finished (the hotfix 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 data” on page 155](#).

- *For applying a hotfix to a HYCU instance:* The same hotfix is applied to the corresponding HYCU backup controller.
- *For applying a hotfix by using the shell script:* You know credentials of an operating system user account that has administrative user rights on the HYCU virtual machine where you plan to apply the hotfix.

 **Important** Unless instructed otherwise by HYCU Customer Support, you must apply the same hotfix to all your HYCU virtual machines: HYCU backup controllers, HYCU instances, and HYCU Managers.

Considerations

- The hotfix 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 a hotfix to a HYCU backup controller or a HYCU Manager:* Any users that have been logged on to the HYCU web user interface of the HYCU virtual machine where the hotfix 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 a hotfix to a HYCU backup controller, back up the HYCU backup controller. For instructions, see [“Backing up virtual machines” on page 79](#).

You can apply HYCU hotfixes:

- From the HYCU web user interface
Use this method if you want to apply a hotfix to a HYCU backup controller, a HYCU instance, or a HYCU Manager. For instructions, see [“Applying a hotfix by using the HYCU web user interface” below](#).
- By using the shell script
Use this method if you are unable to log on to the HYCU web user interface. For instructions, see [“Applying a hotfix by using the shell script” on page 219](#).


Applying a hotfix by using the HYCU web user interface



From the HYCU web user interface, you can apply a hotfix to any kind of HYCU virtual machine by using the following procedures:


- [“Applying a hotfix to a HYCU backup controller or a HYCU Manager” on the next page](#)
- [“Applying a hotfix to a HYCU instance” on the next page](#)

Applying a hotfix to a HYCU backup controller or a HYCU Manager

Procedure


1. Log on to the HYCU web user interface.
2. Click  **Administration**, and then select **Software Upgrade**.
3. In the Software Upgrade dialog box, click the **Hotfixes** tab.
4. In the Hotfix Label column, check if the package of the desired hotfix is already added to the HYCU backup controller or the HYCU Manager, and then do one of the following:
 - If the hotfix label is not present, follow these steps:
 - a. Click **+ Add**.
 - b. In the Add Hotfix dialog box, click **Browse**. browse for the hotfix package (in the ZIP format), select it, and then click **Open**.
 - c. Click **Add Hotfix**.
 - If the hotfix label is present, select it.
5. Click **Apply Hotfix**.
6. Verify that the displayed digital fingerprint matches the one that you were given by HYCU Customer Support.
7. Click **Yes** to start the hotfix application process. You are automatically logged off from the HYCU web user interface, and can track the progress of applying the hotfix on the web user interface logon page.
8. When the process completes, perform a hard reload of the HYCU web user interface page in your web browser.
9. *Only if you applied a hotfix to a HYCU backup controller.* Do the following:
 - a. Log on 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 200](#).

 **Tip** Click  **Info** to review the list of issues that the hotfix resolves.


To delete an added hotfix package when the hotfix is not applied, in the Software Upgrade dialog box in the Hotfixes tab, select its entry from the list of added hotfix packages, and then click  **Delete**.


Applying a hotfix to a HYCU instance

Procedure

1. Log on to the HYCU web user interface.
2. Click  **Administration**, and then select **Instances**.

3. In the Instances dialog box, select the desired HYCU instance, and then click **X Hotfixes**.
4. In the Hotfix Label column, check if the package of the desired hotfix is already added to the HYCU instance, and then do one of the following:
 - If the hotfix label is not present, follow these steps:
 - a. Click **+ Add**.
 - b. In the Add Hotfix dialog box, click **Browse**, browse for the hotfix package (in the ZIP format), select it, and then click **Open**.
 - c. Click **Add Hotfix**.
 - If the hotfix label is present, select it.

 **Note** Each hotfix that is applied to a HYCU instance is first uploaded to the corresponding HYCU backup controller.

-  **Tip** Click **i Info** to review the list of issues that the hotfix resolves.
5. Click **Apply Hotfix**.
 6. Verify that the displayed digital fingerprint matches the one that you were given by HYCU Customer Support.
 7. Click **Yes** to start the hotfix application process. The HYCU instance status icon in the 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. Once the hotfix is applied, the HYCU instance status icon turns green.

To delete an added hotfix package when the hotfix is not applied, in the Hotfixes dialog box, select its entry from the list of added hotfix packages, and then click **🗑 Delete**.

Applying a hotfix by using the shell script

Procedure

1. Log on 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 hotfix.
2. Log on 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 hotfix packages that are already added to the HYCU virtual machine:

```
sudo ./HycuPatch.sh -list_patches
```

5. If the label of the desired hotfix is not present on the list, follow these steps:
 - a. Extract the contents of the hotfix package (in the ZIP format). The package contains the main hotfix 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 hotfix file matches the one that you were given by HYCU Customer Support.
 - c. Copy the main hotfix file in the archived TAR (`.tar.gz`) format to the following directory on the HYCU virtual machine:

```
/hycudata/opt/grizzly/hotfixes
```

 **Tip** Run the following command to review the list of issues that the hotfix resolves:

```
sudo ./HycuPatch.sh -patch_info <HotfixLabel>
```

6. Run the following command to apply the hotfix to the HYCU virtual machine:





```
sudo ./HycuPatch.sh -apply_patch <HotfixLabel>
```



7. *Only if you applied a hotfix to a HYCU backup controller.* Do the following:
 - a. Log on 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 200](#).

Removing HYCU

When you remove HYCU from your environment, you also need to perform additional cleanup tasks.


To remove HYCU, follow these steps:

1. Log on 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  **Policies**.
 - 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  **Policies**.
 - 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  **Policies**.
 - c. Click **Unassign**.
 - d. Click **Yes** to confirm that you want to unassign the policies from the selected file shares.
- 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 191](#).
 - b. Remove the Nutanix Files 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 <NutanixFilesServer> -u <Username> -p
<Password> -dnfs -all
```

In this instance, `<NutanixFilesServer>` is the name of the Nutanix Files server in the following format: `https://<ServerName>:<Port>`.


 **Important** By running this command, you will also remove all Nutanix Files 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 snapshots created by HYCU:

```
sudo perl HycuCleanup.pl -c <NutanixCluster> -u <Username> -p
<Password> -dvms -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 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 `bkpctr1` folder.
5. Log on to the Nutanix Prism web console or the vSphere (Web) Client, and then delete the HYCU backup controller virtual machine. For details on how to delete a virtual machine, see Nutanix or VMware documentation.

Chapter 11

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 225
Configure FIPS-compliant mode for HYCU.	“Configuring FIPS-compliant mode for HYCU” on page 226
Set up LDAPS authentication.	“Setting up LDAPS authentication” on page 228
Secure SMTP connections.	“Securing SMTP connections” on page 228
Set up HYCU to use multiple networks.	“Setting up HYCU to use multiple networks” on page 229
Increase the size of the HYCU virtual disks.	“Increasing the size of the HYCU virtual disks” on page 232
Assign required backup privileges to a vSphere user.	“Assigning privileges to a vSphere user” on page 233
Use the HYCU REST API to automate tasks.	“Using the HYCU REST API Explorer” on page 235
Use hyCLI.	“Using the command-line interface” on

I want to...	Procedure
	page 235
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 236

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.

 **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.

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 service sshd stop
```

When requested, enter the password for the hycu user.

3. Disable the SSH service:

```
sudo chkconfig sshd off
```

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 Nutanix Prism web console.

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 desired operation on the HYCU application server:

```
sudo service grizzly {start | stop | restart}
```

When requested, enter the password for the hycu user.

 **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.

Prerequisite

The `winrm.https.enabled` configuration setting is set to `true`. For details, see [“Customizing HYCU configuration settings” on page 264](#).

Procedure

For each virtual machine for which you want to enable HTTPS for WinRM connections, do the following:

1. Set up a virtual machine for WinRM over HTTPS by using PowerShell:

- a. Create a new self-signed certificate:

```
$cert = New-SelfSignedCertificate -DnsName "hostname"
-CertStoreLocation "Cert:\LocalMachine\My"
```

- b.
- Only if an HTTPS WinRM listener already exists.*
- Remove the existing HTTPS WinRM listener:

```
winrm delete winrm/config/Listener?Address=*&Transport=HTTPS
```

- c.
- Recommended.*
- Remove the HTTP WinRM listener if it exists:

```
winrm delete winrm/config/Listener?Address=*&Transport=HTTP
```

- d. Create an HTTPS WinRM listener that uses the self-signed certificate from step 1:

```
New-Item -Path WSMan:\LocalHost\Listener -Transport HTTPS -Address *
-CertificateThumbPrint $cert.Thumbprint -Force
```

- e. Add a new firewall rule to allow incoming connections on TCP port 5986, if it has not already been added:

```
New-NetFirewallRule -DisplayName 'Windows Remote Management
(HTTPS-In)' -Name 'Windows Remote Management (HTTPS-In)'
-Profile Any -LocalPort 5986 -Protocol TCP
```

2. Open a remote session to the HYCU backup controller, and then do the following:

- a. Run the
- `add_certificate.sh`
- script:

```
sudo /opt/grizzly/bin/add_certificate.sh <Hostname>
```

In this instance, `<Hostname>` is the host name of the virtual machine for which you want to establish an HTTPS connection.

- b. Enter the password to access the trust store. The default password is
- hycu/4u**
- .

After you run the `add_certificate.sh` script, it connects to the virtual machine, imports the self-signed certificate, and adds it to the trust store. You get the information about the certificate that you must check and confirm. If the certificate is valid and matches the information of the certificate on the virtual machine, type **y** followed by **Enter**. Otherwise, type **n** followed by **Enter** to reject the certificate.

Configuring FIPS-compliant mode for HYCU

HYCU can be configured to operate in a way that is compliant with the Federal Information Processing Standards (FIPS), which establish security requirements for cryptography modules (which encryption algorithms and methods for generating encryption keys can be used).

HYCU employs a FIPS-compliant security provider that uses a cryptographically strong random number generator (RNG). For the RNG to generate truly random numbers, an adequate source of entropy is required. Because HYCU is running on a virtual machine, the amount of entropy that is available to it is limited and therefore an additional hardware source of entropy is needed. This source is usually provided by the real CPU or chipset. To enable access to this hardware entropy source, an additional service (`rngd.service`) is enabled on the HYCU backup controller.

Depending on the nature of your business, you can either enable or disable FIPS-compliant mode for HYCU.

Limitation

When FIPS-compliant mode is enabled, you cannot assign credentials to Linux virtual machines, and consequently restore individual files.

Enabling FIPS-compliant mode for HYCU

To enable FIPS-compliant mode for HYCU, 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. Start the HYCU application server:

```
systemctl start grizzly.service
```

Disabling FIPS-compliant mode for HYCU

If the nature of your business changes, you can easily disable FIPS-compliant mode. To do so, 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. Start the HYCU application server:

```
systemctl start grizzly.service
```

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, the LDAPS server certificate must be imported to HYCU.

To import the LDAPS server certificate to HYCU, open a remote session to the HYCU backup controller, and then do the following:

1. Run the `add_certificate.sh` script:

```
sudo /opt/grizzly/bin/add_certificate.sh <Hostname> <Port>
```

In this instance, `<Hostname>` is the LDAPS server host name and `<Port>` is the LDAPS port (usually 636).

2. Enter the keystore password. The default password is **hycu/4u**.

After you run the `add_certificate.sh` script, it connects to the LDAPS server, imports the certificate, and adds it to the keystore.

You get the information about the certificate that you must check and confirm. If the certificate is valid and matches the information of the certificate on the LDAPS server, type **y** followed by **Enter**. Otherwise, type **n** followed by **Enter** to reject the certificate.

Securing SMTP connections

If you are using STARTTLS or SSL/TLS for SMTP connections, you must import an SSL certificate to HYCU. Depending on the protocol you use to secure email traffic, see one of the following sections:

- ["Importing an SSL certificate for the STARTTLS security mode" below](#)
- ["Importing an SSL certificate for the SSL/TLS security mode" on the next page](#)

Importing an SSL certificate for the STARTTLS security mode

Procedure

Open a remote session to the HYCU backup controller, and then do the following:

1. Run the `add_certificate_starttls.sh` script:

```
sudo /opt/grizzly/bin/add_certificate_starttls.sh <Hostname> <Port>
```

In this instance, `<Hostname>` is the SMTP server host name and `<Port>` is the port for authenticated SMTP connections (587 or 25).

2. Enter the keystore password. The default password is **hycu/4u**.

After you run the `add_certificate_starttls.sh` script, it connects to the SMTP server, imports the certificate, and adds it to the keystore.

You get the information about the certificate that you must check and confirm. If the certificate is valid and matches the information of the certificate on the SMTP server, type **y** followed by **Enter**. Otherwise, type **n** followed by **Enter** to reject the certificate.

Importing an SSL certificate for the SSL/TLS security mode

Procedure

Open a remote session to the HYCU backup controller, and then do the following:

1. Run the `add_certificate.sh` script:

```
sudo /opt/grizzly/bin/add_certificate.sh <Hostname> <Port>
```

In this instance, `<Hostname>` is the SMTP server host name and `<Port>` is the port for authenticated SMTP connections (465).

2. Enter the keystore password. The default password is **hycu/4u**.


After you run the `add_certificate.sh` script, it connects to the SMTP server, imports the certificate, and adds it to the keystore.

You get the information about the certificate that you must check and confirm. If the certificate is valid and matches the information of the certificate on the SMTP server, type **y** followed by **Enter**. Otherwise, type **n** followed by **Enter** to reject the certificate.

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.

Nutanix Files 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 Nutanix Files 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 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 ESXi cluster

Procedure

1. Log on 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, from the VLAN Name drop-down menu, select the required VLAN.
 - e. Click **Add**.
 - f. 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

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`.


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

 **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. Log on 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. 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`.

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. To do so, follow the instructions in one of the following sections:


- [“Increasing the size of the HYCU disks in a Nutanix AHV cluster” below](#)
- [“Increasing the size of the HYCU disks in a Nutanix ESXi cluster or vSphere environment” below](#)

Increasing the size of the HYCU disks in a Nutanix AHV cluster

To increase the size of the HYCU system disk and/or data disk in a Nutanix AHV cluster, follow these steps:


1. Log on 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 your HYCU backup controller, and then click **Power Off Actions** followed by **Power off** to shut it down.
-  **Important** Wait a moment for the virtual machine to shut down completely.
5. Click **Update**, and then do the following:
 - a. Navigate to the Disks section, and then click **Edit** next to the HYCU disk whose size you want to increase.
 - b. In the Size (GiB) field, increase the size of the disk as required.
 - c. *For increasing the size of both HYCU disks:* Repeat steps a and b for the other HYCU disk.
 - d. Click **Update**.
 6. Click **Power on** to turn on the HYCU backup controller.

Increasing the size of the HYCU disks in a Nutanix ESXi cluster or vSphere environment

 **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.

To increase the size of the HYCU system disk and/or data disk in a Nutanix ESXi cluster or vSphere environment, follow these steps:

1. Log on to the vSphere Web Client.
2. Click the **VMs** tab, and then navigate to your HYCU backup controller.
3. Right-click the HYCU backup controller, and then select **Power > Power Off** to shut it down.


 **Important** Wait a moment for the virtual machine to shut down completely.

4. Right-click the HYCU backup controller, and then select **Edit Settings**.
5. On the Virtual Hardware tab, increase the size of one or both HYCU disks by entering new values in the Hard disk 1 and/or Hard disk 2 fields, and then click **OK**.
6. Right-click the HYCU backup controller, and then select **Power > Power On** to turn it on.

For details on how to manage a virtual machine in a Nutanix AHV or ESXi cluster, see Nutanix documentation. For details on how to manage a virtual machine in a vSphere environment, see VMware documentation.

Assigning privileges to a vSphere user

You can assign required privileges to a user by using the vSphere (Web) Client.

 **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. Log on to the vSphere Web Client as an administrator.
2. On the Home page, click **Roles**.
3. Right-click the Roles tab information panel, and then click **Add**.
4. Type a name for the new role (for example, **HYCU**).
5. Select the required privileges for the role, and then click **OK**.

Privilege category	Backup privileges	Restore privileges	Upgrade and HYCU instance creation 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	Not applicable	<ul style="list-style-type: none"> • Create virtual 	Not applicable

Privilege category	Backup privileges	Restore privileges	Upgrade and HYCU instance creation privileges
operations		machine <ul style="list-style-type: none"> Delete virtual machine Reconfigure virtual machine 	
Network	Not applicable	<ul style="list-style-type: none"> Assign network Configure 	<ul style="list-style-type: none"> Assign network
vApp	Not applicable	<ul style="list-style-type: none"> Add virtual machine 	<ul style="list-style-type: none"> Import
Virtual Machine > Configuration	<ul style="list-style-type: none"> Disk change tracking Settings 	All privileges	<ul style="list-style-type: none"> Add existing disk Add new disk Add or remove device Settings Remove disk Rename
Virtual Machine > Interaction	<ul style="list-style-type: none"> Power On 	<ul style="list-style-type: none"> Answer question Power Off Power On 	<ul style="list-style-type: none"> Power On
Virtual Machine > Inventory	Not applicable	<ul style="list-style-type: none"> Create new Register Remove Unregister 	<ul style="list-style-type: none"> Create from existing Remove
Virtual Machine > Provisioning	<ul style="list-style-type: none"> 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 	<ul style="list-style-type: none"> Allow disk access 	<ul style="list-style-type: none"> Clone virtual machine
Virtual Machine > Snapshot	<ul style="list-style-type: none"> Create snapshot Remove snapshot 	Not applicable	Not applicable


Privilege category	Backup privileges	Restore privileges	Upgrade and HYCU instance creation privileges
management			
Resource	Not applicable	<ul style="list-style-type: none"> Assign virtual machine to resource pool 	Not applicable

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  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 desired 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  at the upper right of the screen, and then select **Download hyCLI**.
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.

 **Note** hyCLI log files are located at `.Hycu/log` in the user's home directory. You can change logging settings for hyCLI in the `logging.properties` files located in the directory containing the extracted files.

For detailed information about hyCLI, see the `README.txt` file that you can find in the directory 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
<code>HYCU_TARGET_NAME</code>	Target name
<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_PREEEXEC_RETURN_CODE</code>	<i>Available only for post scripts.</i> Exit code of the pre script.

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 77](#)
- [“Restoring individual files” on page 97](#)

Chapter 12

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 with easy access to the collected information. You can view this information for the on-premises (HYCU) and the following cloud data protection environments:

- HYCU Data Protection as a Service for GCP (HYCU for GCP) data protection environment
- HYCU Data Protection as a Service for Azure (HYCU for Azure) data protection environment


For details on how to protect data with HYCU for GCP or HYCU for Azure, see HYCU for GCP or HYCU for Azure documentation.

After you deploy the HYCU virtual appliance in the HYCU Manager mode, you can access HYCU Manager and take advantage of this intuitive visualization approach to quickly identify and address potential issues.

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**.

You can find the following information within each widget:

Console widget	Description
Virtual Machines	Number of all virtual machines and the number of protected and unprotected virtual machines in your data protection environments.
Applications	<i>For the HYCU data protection environments:</i> Number of all applications and the number of protected and unprotected applications in your

Console widget	Description
	data protection environments.
HYCU Controllers	Number of available and unavailable HYCU controllers in your data protection environments.
Backups	<p>Percentage of successful backups and the number of successful and migration/DR-ready backups in your data protection environments. You can safely ignore the Migration/DR-ready label if you do not plan to employ HYCU Protégé. A backup is migration/DR-ready if all backups in the current backup chain are stored on one of the cloud targets (Google Cloud or Azure) and a successful cloud readiness check was performed during the latest backup.</p> <p>For detailed information about backups, see “Backing up virtual machines” on page 79.</p>
Shares	<i>For the HYCU data protection environments:</i> Number of all file shares and the number of protected and unprotected file shares in your data protection environments.
Targets	Number of all targets and the number of free and used targets in your data protection environments.
Policies	Number of all policies and the number of compliant and non-compliant policies in your data protection environments. A policy is considered compliant if all entities to which this policy is assigned are compliant with the policy settings.

 **Important** By clicking a value in any of the widgets, you are directed to the HYCU Controllers panel where you can view a list of the HYCU controllers sorted by the value you clicked. For example, if you click the number of compliant policies, the HYCU controllers are sorted by the policy compliance percentage in descending order.

Monitoring your HYCU controllers


You can use the HYCU Controllers panel to add, edit, and remove the HYCU controllers, as well as to view the information about each of them. Depending on your data protection strategy, which can include the on-premises and cloud data protection environments, a list of your HYCU controllers can include the HYCU backup controllers (on-premises controllers) and the HYCU for GCP and/or HYCU for Azure protection sets (cloud controllers).

Consideration

Only if you are monitoring the HYCU for GCP data protection environments. The list of your HYCU controllers includes both the projects and the protection sets in which these projects are

included.

Accessing the HYCU Controllers panel

To access the HYCU Controllers panel, in the navigation pane, click  **HYCU Controllers**.

Adding a HYCU controller


Prerequisites


Only if you plan to monitor cloud data protection environments:

- A cloud account is added to HYCU. Depending on the cloud data protection environment that you want to monitor, see [“Adding a Google Cloud Platform service account” on page 186](#) or [“Adding an Azure service principal” on page 188](#).
- You own a HYCU Protégé license. For more information, see [“Licensing” on page 192](#).
- You have an active subscription for HYCU for GCP or HYCU for Azure. For details, see HYCU for GCP or HYCU for Azure documentation.
- *For HYCU for GCP data protection environments:* The projects included in the protection set that you plan to monitor are linked to the Google Cloud Platform billing account that was selected when subscribing to HYCU for GCP. For details, see HYCU for GCP documentation.



Procedure



1. In the HYCU Controllers panel, click **+ Add**. The New Controller dialog box opens.
2. Depending on which data protection environment you want to monitor, select one of the following options:

Option	Instructions
Add on-premises controller	<ol style="list-style-type: none"> a. Click Next. The Add On-Premises Controller dialog box opens. b. Enter the name of the HYCU backup controller. c. Enter the URL of the HYCU backup controller. d. Enter the user name and password of an infrastructure group administrator.
Add cloud controller	<ol style="list-style-type: none"> a. Click Next. The Add Cloud Controller dialog box opens. b. From the list of all available cloud controllers, select the HYCU for GCP and/or 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. <p> Tip You can see which GCP projects or Azure</p>

resource groups are included in each available protection set by clicking .



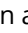

3. Click **Add**.

You can later edit any of the existing on-premises controllers (click  **Edit** and make the required modifications) or remove the HYCU controllers that you do not want to monitor anymore from HYCU Manager (click  **Delete**). If you use HYCU Manager to monitor also the cloud data protection environments, keep in mind that you cannot edit cloud controllers.

 **Tip** You can update data related to the data protection environments by clicking  **Synchronize**.

Viewing information about HYCU controllers

You can view the following information about each HYCU controller:

HYCU controller information	Description
Name	<p>Name of the HYCU controller.</p> <p>An on-premises 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 GCP</i>: The  icon and the name of the Google Cloud Platform service account and the HYCU for GCP protection set. • <i>HYCU for Azure</i>: The  icon and the name of the Azure service principal and the HYCU for Azure protection set. <p> Tip If you click the name of the HYCU controller, you are directed to the relevant web user interface.</p>
Version	HYCU software release version on the HYCU backup controller.
Status	Status of the HYCU controller (active or inactive).
Backups	Percentage of successful and failed backups.
Migration/DR-ready VMs	<i>Applicable only for the HYCU backup controllers.</i> Number of migration/DR-ready virtual and physical machines. A virtual or physical machine is migration/DR-ready if all backups in the current backup chain are stored on one of the cloud targets


HYCU controller information	Description
	(Google Cloud or Azure) and a successful cloud readiness check is performed during its latest backup.
VM protection	Percentage of protected and unprotected virtual machines.
App protection	<i>For HYCU data protection environments:</i> Percentage of protected and unprotected applications.
Share protection	<i>For HYCU data protection environments:</i> Percentage of protected and unprotected file shares.
Policy compliance	Percentage of compliant and non-compliant policies.
Target utilization	Percentage of used and free 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 161](#).

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 email notifications when events occur, and export the contents of the panel to a file in JSON or CSV format.


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.	“Viewing events” on page 146
Apply filters to events.	“Filtering data” on page 155
Configure HYCU to send email notifications when events occur.	“Sending email notifications” on page 146
Export event data.	“Exporting the contents of the panel” on page 161

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
Configure Active Directory authentication.	“Configuring Active Directory authentication” on page 184
Add cloud accounts to be able to monitor cloud data protection environments.	“Adding a cloud account” on page 185
Configure log file settings to troubleshoot problems if HYCU does not perform as expected.	“Setting up logging” on page 195
Change network settings.	“Changing network settings” on page 197
Configure an SMTP server.	“Configuring an SMTP server” on page 200
Upgrade HYCU to a new available version.	“Upgrading HYCU” on page 204
 Important Before upgrading, make sure you have added the source where your HYCU Manager virtual machine resides as described in “Adding sources” on page 31 .	
Configure the SSL certificate.	“Configuring SSL certificates” on page 201
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 235](#).
- Use the HYCU REST API Explorer. For details, see [“Using the HYCU REST API Explorer” on page 235](#).

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**. The New dialog box opens.
2. Enter a user name if you are adding a HYCU Manager user or an AD user, or a common name if you are adding an AD group.

 **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 264](#).

3. From the Authentication type drop-down menu, select one of the following authentication types:

- **HYCU**

Enter a display name, the user password and, optionally, email address.

 **Note** The minimum password length is six characters.


- **AD user**

From the Active Directory drop-down menu, select the Active Directory the AD user belongs to.



- **AD group**



From the Active Directory drop-down menu, select the Active Directory the AD group belongs to.

4. Click **Save**. The user is added to the list of users.

 **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.

You can later do the following:

- Edit any of the existing 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.
- Delete any of the existing users by clicking  **Delete**. Keep in mind that the built-in user cannot be deleted.
- Enable or disable specific users from logging on to the HYCU Manager console:

- 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**.

Chapter 13

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 42
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 have acquired a Nutanix Mine appliance.
- The HYCU backup controller resides on a Nutanix Mine cluster and this cluster is added to HYCU as a source. For details, see ["Deploying HYCU to a Nutanix AHV cluster" on page 22](#) and ["Adding a Nutanix cluster" on page 31](#).
- *For repeating the registration procedure:* Currently running jobs that you do not want to be aborted are finished.


Consideration


- All instructions that apply to the Nutanix AHV cluster, also apply 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.
 - AOS of the Nutanix Mine cluster was upgraded to a new version.
 - A new Controller VM 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. Log on 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. For details on protecting virtual and physical machines, see “Protecting virtual machines” on page 66 .
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. For details on protecting applications, see “Protecting applications” on page 102 .
Compliance	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 54 .
Backups	Backup success rates for the last seven days.
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 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 38 .
Target Summary	List of all targets other than Nutanix and S3-compatible targets in the data protection environment, and the information on how much space is used and available for storing data. For details on targets, see “Setting up targets” on page 38 .
HYCU Controller	Information on whether the HYCU backup controller is protected and

Dashboard widget	Description
	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 virtual machine resources” on page 171 .
Events	Number of events in the data protection environment in the last 48 hours according to their status (Success, Warning, and Failed). For details on events, see “Viewing events” on page 146 .
Jobs	Number of jobs in the data protection environment in the last 48 hours according to their status (Success, Warning, Failed, and In Progress). For details on jobs, see “Checking the status of jobs” on page 144 .

 **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 14

HYCU Protégé

The HYCU Protégé solution ensures business continuity of your data protection environment across different infrastructures. Besides storing backup data to Google Cloud or Azure targets, you can ensure data resilience by migrating virtual machines across the on-premises and cloud infrastructures (Google Cloud Platform or Azure). In the event of a disaster in your on-premises environment, HYCU Protégé provides disaster recovery of data to cloud.

Depending on your cloud environment, see one of the following sections:

- [“Protecting data across on-premises and Google Cloud Platform environments” below](#)
- [“Protecting data across on-premises and Azure environments” on page 256](#)

Protecting data across on-premises and Google Cloud Platform environments

HYCU Protégé ensures data resilience by using the SpinUp functionality to migrate protected data across the on-premises and Google Cloud Platform environments. In the event of a disaster, it provides disaster recovery of data to Google Cloud Platform.

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 Platform environments.	“Migrating virtual machines across different environments” on the next page
Perform disaster recovery of data to Google Cloud Platform.	“Performing disaster recovery of data to Google Cloud Platform” on page 254

Prerequisites

- You have an active subscription for HYCU for GCP. For instructions, see HYCU for GCP documentation.
- A Google Cloud Platform service account is added to HYCU. For instructions, see [“Adding a Google Cloud Platform service account” on page 186](#).
- You own a HYCU Protégé license. For instructions, see [“Licensing” on page 192](#).


Migrating virtual machines across different environments

You can migrate protected data across the on-premises and Google Cloud Platform environments:

- [“Migrating data to cloud” below](#)
- [“Migrating data from cloud” on page 252](#)

Migrating data to cloud

You can migrate virtual and physical machines as well as applications running on them to cloud by using the HYCU SpinUp functionality. Keep in mind that when you migrate an application, the whole virtual machine on which this application is running is migrated to cloud.

 **Note** The instructions for protecting virtual machine data apply also to physical machines 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 are protected and have a successful cloud readiness check during the backup. For more information, see [“HYCU Protégé specifics” on page 72](#).

Limitations


- *For Nutanix clusters:* You cannot migrate volume groups.
- *For vSphere environments:* You cannot migrate virtual machine templates.


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 VM to Cloud**. The SpinUp VM to Cloud dialog box appears.
4. Select **SpinUp VM to GCP**, and then click **Next**. The SpinUp VM to GCP dialog box appears.
5. From the Cloud account drop-down menu, select the Google Cloud Platform service account to which the project where you want to migrate the virtual machine is linked.
6. From the Project, Target region, and Target zone drop-down menus, select the required values, and then click **Next**. The VM Settings dialog box appears.
7. In the New VM name field, enter a name for the migrated virtual machine instance.
8. 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. The maximum number of vCPU cores that you can specify is 1024.
9. In the Memory field, set the amount of memory (in GiB) for the migrated virtual machine instance. The value that you specify must be a whole number and cannot be higher than 4096. The default value is the amount of memory in GiB of the original virtual machine.
10. From the Virtual machine type drop-down menu, select the machine type for the migrated virtual machine instance.

 **Note** The list contains machine types that match the specified number of virtual CPUs and amount of memory. If no such match exists, you can select the custom machine type. For more information about machine types, see Google Cloud Platform documentation.
11. 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). Depending on your data protection needs, you can leave the default network interface or do one of the following:
 - Add a new network interface:
 - a. Click **Add Network Interface**. The Network dialog box appears.
 - b. 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 desired external IP address resource. For details, see HYCU for GCP 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 desired IP address.
 - From the Internal address drop-down menu, select the name of the

desired internal IP address resource.

For details, see HYCU for GCP 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**.
12. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
 - **Linux**
 - **Windows**
 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. 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 for GCP. For details, see HYCU for GCP 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 for GCP. For details, see HYCU for GCP documentation.


Migrating data from cloud

You can migrate virtual machine instances from cloud by using the HYCU 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 VM from Cloud**. The SpinUp VM from Cloud dialog box appears.
2. Select **SpinUp VM from GCP**, and then click **Next**. The SpinUp VM from GCP dialog box opens.


3. From the Cloud account drop-down menu, select the Google Cloud Platform service account to which the project containing the virtual machine instance that you want to migrate is linked.
4. From the Project drop-down menu, select the Google Cloud Platform project to which the virtual machine instance that you want to migrate belongs.
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 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 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:
 - The number of virtual CPUs. The maximum number that you can specify is 1024.
 - The number of cores per virtual CPU. The maximum number that you can specify is 64.
 - The amount of memory (in GiB). The value that you specify must be a whole number and cannot be higher than 4096.

 **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 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.
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

- 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 66](#) and [“Protecting applications” on page 102](#).

Performing disaster recovery of data to Google Cloud Platform

You can perform disaster recovery of data from the on-premises environment to Google Cloud Platform 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 Platform project where you want to deploy your new HYCU backup controller.
 - To deploy Google Compute Engine VM instances to the Google Cloud Platform

project where you want to deploy your new HYCU backup controller.

- To set up a firewall rule in the Google Cloud Platform network where you plan to deploy your new HYCU backup controller.
- You have the HYCU virtual appliance image for the Google Cloud Platform service suite. To obtain the image and further instructions, contact HYCU Customer Support.
- The virtual machines that you want to migrate and the virtual machines on which applications that you want to migrate are running are protected and have the Migration/DR-ready status. For more information, see [“HYCU Protégé specifics” on page 72](#).

Consideration


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

Procedure

1. Open a web browser, go to the Google Cloud Platform Console webpage, and sign in to Google.
2. Select the Google Cloud Platform project where you want to deploy the HYCU backup controller.
3. In the Compute Engine browser, in the Images context, create a new Google Compute Engine image from the HYCU virtual appliance image. For instructions, see Google Cloud Platform documentation.
4. Based on this image, create a VM instance with an additional disk of 32 GB in size. A HYCU backup controller is deployed. For instructions, see Google Cloud Platform documentation.
5. 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 instructions, see Google Cloud Platform documentation.
6. Log on to HYCU by specifying the following URL:

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

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

7. Add a Google Cloud Platform service account with permissions to access the Google Cloud Storage buckets where backup data of the protected virtual machines is stored. For instructions, see [“Adding a Google Cloud Platform service account” on page 186](#).
8. Import the Google Cloud target with your backup data:
 - a. In the Targets panel, click  **Import**. The Import Target dialog box appears.
 - 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 Platform 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**.
9. Migrate your virtual machines or applications to cloud. For instructions, see [“Migrating data to cloud” on page 250](#).

Protecting data across on-premises and Azure environments

HYCU Protégé ensures data resilience by using 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.

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” below
Perform disaster recovery of data to Azure.	“Performing disaster recovery of data to Azure” on page 262

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 188](#).
- You own a HYCU Protégé license. For details, see [“Licensing” on page 192](#).
- A storage account is created in Azure in the same region and resource group as the virtual machine that you plan to migrate. This storage account must be dedicated exclusively to migration operations.


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 259](#)

Migrating data to cloud

You can migrate virtual and physical machines as well as applications running on them to cloud by using the HYCU SpinUp functionality. Keep in mind that when you migrate an application, the whole virtual machine on which this application is running is migrated to cloud.

 **Note** The instructions for protecting virtual machine data apply also to physical machines 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 are protected and have a successful cloud readiness check during the backup. For more information, see [“HYCU Protégé specifics” on page 72](#).



Limitations

- *For Nutanix clusters:* You cannot migrate volume groups.
- *For vSphere environments:* You cannot migrate virtual machine templates.

Considerations


- 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 Azure, 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 VM to Cloud**. The SpinUp VM to Cloud dialog box appears.
4. Select **SpinUp VM to Azure**, and then click **Next**. The SpinUp VM to Azure dialog box appears.
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 HYCU for Azure 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 Region drop-down menu, select the geographic region for the migrated virtual machine.
9. From the Storage account drop-down menu, select the storage account that is dedicated exclusively to migration operations.
10. Click **Next**. The VM Settings dialog box appears.
11. In the New VM name field, enter a name for the migrated virtual machine.
12. 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. The maximum number that you can specify is 1024.
13. In the Memory field, enter the amount of memory (in GiB) to be assigned to the migrated virtual machine. The value that you specify must be a whole number and cannot be higher than 4096.
14. From the Virtual machine type drop-down menu, select the virtual machine type.

 **Important** The list of available virtual machine types is based on the number of virtual CPU cores and the amount of memory that you specified. If no virtual machine type exactly corresponds to the specified values, the closest matches are shown.

15. Under Network interfaces, the default network interface is displayed and you can check to which network and subnet it is assigned (based on the selected resource group and region). Depending on your data protection needs, you can leave the default network interface or do one of the following:
 - Add a new network interface:
 - a. Click **Add Network Interface**. The Network dialog box appears.
 - b. From the Network drop-down menu, select the virtual network for the network interface.

 **Important** If a network interface already exists, you can select only

the virtual network to which the existing network interface is assigned. If you want to assign the new network interface to a different virtual network, you must first delete the existing network interface.

- c. From the Subnet drop-down menu, select a subnet within the selected virtual network to which the network interface will be assigned.
- d. Click **Add**.
 - Select another subnet for the existing network interface by selecting it, clicking  **Edit**, and then making the required modification.
 - Delete the existing network interface by selecting it, and then clicking  **Delete**.
16. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
 - **Linux**
 - **Windows**
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. 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 cloud by using the HYCU SpinUp functionality.


Limitation

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 VM from Cloud**. The SpinUp VM from Cloud dialog box appears.
2. Select **SpinUp VM from Azure**, and then click **Next**. The SpinUp VM from Azure dialog box appears.
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 appears.
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 maximum number that you can specify is 1024.
 - The number of cores to be assigned to each virtual CPU. The maximum number that you can specify is 64.
 - The amount of memory (in GiB). The value that you specify must be a whole number and cannot be higher than 4096.

 **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 appears.
 - 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 66](#) and [“Protecting applications” on page 102](#).

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.

Prerequisites

- You have the HYCU virtual appliance image for Azure. To obtain the image and further instructions, contact HYCU Customer Support.
- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate are protected and have the Migration/DR-ready status. For more information, see [“HYCU Protégé specifics” on page 72](#).

Consideration

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

Procedure


1. In Azure, do the following:
 - a. Deploy a HYCU backup controller:
 - i. Create a managed image from the HYCU virtual appliance image.
 - ii. 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.
 - b. Create a new firewall rule to allow ingress network traffic on TCP port 8443 from the entire subnet to which the HYCU backup controller belongs.


For instructions, see Azure documentation.

2. Log on 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 log on to HYCU and perform disaster recovery of data to Azure. For details on what credentials you can use to log on to HYCU or to access the HYCU backup controller by using SSH, see [“Logging on to HYCU” on page 28](#) or [“Accessing the HYCU backup controller virtual machine by using SSH” on page 224](#).


3. 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 appears.
 - b. From the Type drop-down menu, select **AZURE**, **AZURE Government**, or **AZURE China**.

- 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 appears.
 - 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:*
 - i. Select each target one by one and specify the values so that they match the original target configuration.
 - ii. For each target, click **Validate** to check the configuration.
 - iii. Click **Import**.
4. Migrate your virtual machines or applications to cloud. For instructions, see [“Migrating data to cloud” on page 257](#).

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 same folder, and then specify the desired 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 the next page](#)
- [“Utilization threshold settings” on page 266](#)
- [“Display settings” on page 266](#)
- [“SQL Server application settings” on page 266](#)
- [“Settings for aborting jobs” on page 267](#)
- [“HTTPS for WinRM configuration settings”](#)
- [“Nutanix Files settings” on page 267](#)
- [“Data rehydration settings” on page 268](#)
- [“HYCU backup controller restore settings” on page 269](#)
- [“User management settings” on page 269](#)

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 224](#).

2. Access and open the `config.properties` file by using one of the following text editors:

- Vim:

```
sudo vi /opt/grizzly/config.properties
```

- Nano:

```
sudo nano /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 when a new target is added to HYCU or an existing target is activated.
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:

Setting	Description
<code>max.snapshots.per.vm</code>	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.
<code>max.snapshots.per.cluster</code>	If the number of snapshots that are retained per Nutanix

Setting	Description
	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:

Setting	Description
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 162](#).

Display settings

You can use the following setting to customize the maximum number of displayed items:

Setting	Description
items.per.directory.in.flr	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
sql.translog.compress	During the backup of an SQL Server application, transaction

Setting	Description
	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:

Setting	Description
<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.

HTTPS for WinRM configuration settings

You can use the following settings to configure HTTPS for WinRM:

Setting	Description
<code>winrm.https.enabled</code>	HYCU is preconfigured to use HTTP for WinRM connections to virtual machines. If you want HYCU to use HTTPS instead, make sure to set the value for this setting to <code>true</code> , and then perform the procedure described in “Enabling HTTPS for WinRM connections” on page 225 .
<code>winrm.fallback.http</code>	<i>For configuring HTTPS if <code>winrm.https.enabled</code> is set to <code>true</code>:</i> If set to <code>true</code> , HYCU uses HTTP for WinRM connections to virtual machines if using HTTPS fails due to certificate issues.

Nutanix Files settings

You can use the following settings to configure file share backups:

Setting	Description
<code>afs.reindex.interval.count</code>	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.

Setting	Description
afs.partial.success.threshold.count	Number of failed file backups up to which the backup status of the corresponding file share is Completed with errors. The default value is 100. Value 0 disables the status.
afs.instance.afs.cluster.priority	<p>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 Nutanix Files 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.</p> <p>Raising the value of this setting gives higher priority to the HYCU instances that are running on the same Nutanix cluster as the Nutanix Files server.</p>
afs.instance.bc.cluster.priority	<p>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 Nutanix Files 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.</p> <p>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.</p>

Data rehydration settings


You can use the following settings to configure HYCU to perform data rehydration:

Setting	Description
target.azure.blob.rehydration.enable	HYCU is preconfigured to perform data rehydration before performing a restore if some or all of the restore point 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

Setting	Description
	the archive storage tier afterward. The default value is <code>true</code> .
<code>target.azure.blob.rehydration.threads</code>	Number of blobs that can be rehydrated in parallel. The default value is 20.

HYCU backup controller restore settings

You can use the following setting to enable additional scenarios for disaster recovery (restore of the HYCU backup controller):

Setting	Description
<code>clone.enabled.for.hycu.dr</code>	<p>HYCU is preconfigured to prevent creating clones of the HYCU backup controller (the virtual machine itself or its virtual disks).</p> <p> 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.</p> <p>If set to <code>true</code>, cloning of the HYCU backup controller is enabled and the respective restore options become available in the HYCU web user interface.</p>

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
<code>force.keep.backups.on.owner.change</code>	If set to <code>true</code> (the default value is <code>false</code>), 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

Restoring to an environment with a different hypervisor

This appendix describes prerequisites, limitations, considerations, and/or additional steps that you should perform to successfully restore a virtual machine to an environment that is based on a different hypervisor.

VM source environment	VM target environment	Restore option	Additional information
Nutanix ESXi or vSphere	Nutanix AHV	Clone VM	See “Restoring a virtual machine from a Nutanix ESXi cluster or a vSphere environment to a Nutanix AHV cluster” on the next page.
vSphere	Nutanix ESXi	Clone VM	See “Restoring a virtual machine from a vSphere environment to a Nutanix ESXi cluster” on page 274.
Nutanix AHV	Nutanix ESXi	Clone VM	A virtual machine on a Nutanix AHV cluster is restored to a Nutanix ESXi cluster as described in “Cloning a virtual machine” on page 85, and no additional actions are required.
Nutanix AHV or Nutanix ESXi	vSphere	Clone VM	See “Restoring a virtual machine from a Nutanix AHV cluster or a Nutanix ESXi cluster to a vSphere environment” on page 275.

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 details on how to restore a virtual machine, see [“Restoring virtual machines” on page 80](#).

- *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.

Restoring a virtual machine from a Nutanix ESXi cluster or a vSphere environment to a Nutanix AHV cluster

Prerequisite

A Nutanix AHV cluster is added to HYCU. For details on how to do this, see [“Adding a Nutanix cluster” on page 31](#).

Limitation

You can restore virtual machines that use UEFI firmware to a Nutanix AHV cluster only if your AOS version is 5.15 or later. For details, see Nutanix documentation.

Recommendations

To avoid having to perform manual steps after restoring a virtual machine on a Nutanix ESXi cluster or in a vSphere environment to a Nutanix AHV cluster, you should follow these recommendations before backing it up:

- *For Windows virtual machines:* The Nutanix VirtIO package is installed on the virtual machine. If you have NGT installed on your virtual machine on the Nutanix ESXi cluster, there is no need to install the Nutanix VirtIO package because it is already installed as part of NGT installation.
- *For Linux virtual machines on Nutanix ESXi clusters:* NGT is installed on your virtual machine.
- *For Linux virtual machines in vSphere environments:* The VirtIO drivers are added to the guest OS kernel.

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 the kernel, 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 the kernel:

```
97084 blocks
```

If the output is blank, the VirtIO drivers are not added to the kernel. To add the VirtIO drivers to the kernel, 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 the kernel, 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.

If you have not followed the above recommendations, 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, log on to the Nutanix AHV cluster by using SSH.
3. List the virtual machine details:

```
accli 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:

```
accli 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:

```
accli 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:

```
accli 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. Log on 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.

Restoring a virtual machine from a vSphere environment 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. Log on 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, edit the settings of the restored virtual machine as follows:
 1. Log on to the vSphere Web Client.
 2. Click the **VMs** tab, and then right-click the restored virtual machine and select **Edit Settings**.
 3. *For Linux virtual machines:* On the Virtual Hardware tab, change the controller to **VMware Paravirtual**.
 4. Click the **VM Options** tab, and then, under Boot Options, change the firmware to **EFI**.
 5. *Only if you need to select the boot file manually.* 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.

Restoring a virtual machine from a Nutanix AHV cluster or a Nutanix ESXi cluster to a vSphere environment

Procedure

1. Restore the virtual machine to a new location by creating its clone. For instructions, see [“Cloning a virtual machine” on page 85](#)
2. *Only if the original virtual machine resided on a Nutanix AHV cluster.* Modify the virtual machine configuration by specifying the appropriate guest operating system.
3. *Only if the restored virtual machine has more than one disk.* Check the order of virtual disks on the restored virtual machine. If the order differs from the one on the original virtual machine, make the necessary adjustments, including the selection of the correct boot disk.

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!

