

USER GUIDE

HYCU Data Protection for Nutanix

Version: 3.0.0

Product release date: April 2018

Document release date: April 2018



Legal notices

Copyright notice

© 2017–2018 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.

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.

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	8
HYCU key features and benefits	9
HYCU backup environment overview	9
2 Deploying the HYCU virtual appliance	11
Sizing your backup infrastructure for HYCU	11
Firewall settings	12
Deploying HYCU on a Nutanix AHV cluster	12
Uploading the HYCU virtual appliance image to a Nutanix AHV cluster	14
Creating a virtual machine for HYCU deployment on a Nutanix AHV cluster ..	14
Configuring HYCU on the virtual machine	15
Deploying HYCU to a Nutanix ESXi cluster	16
Logging on to HYCU	17
3 Establishing a backup environment	19
Adding Nutanix clusters	20
Setting up backup targets	21
How to set up an AWS S3/Compatible target	21
How to set up an Azure target	22
How to set up an NFS target	23
How to set up an SMB target	24
How to set up an iSCSI target	26
Defining your backup policy strategy	27
Applying a predefined backup policy	28
Creating a custom backup policy	29
Setting a default backup policy	36
4 Protecting data	37
Enabling access to data	39
Assigning credentials to virtual machines	40
Assigning credentials to applications	41

Backing up virtual machines	42
Backing up applications	43
Performing a manual backup	45
5 Restoring data	46
Restoring virtual machine data	46
Restoring an entire virtual machine	46
Restoring individual files	50
Restoring application data	52
Restoring a whole application	52
Restoring application items	56
6 Protecting the HYCU backup controller	62
Backing up the HYCU backup controller	62
Recovering the HYCU backup controller	63
7 Performing daily tasks	66
Using the HYCU dashboard	66
Checking the status of jobs	68
Viewing events	69
Viewing backup source details	69
Viewing the backup status of backup sources	70
Filtering data in panels	71
Filtering options in the Applications panel	72
Filtering options in the Virtual Machines panel	72
Filtering options in the Policies panel	73
Filtering options in the Targets panel	74
Filtering options in the Jobs panel	74
Filtering options in the Events panel	75
Filtering options in the Self-Service panel	75
Managing backup targets	75
Viewing backup target information	76
Editing a backup target	77
Activating or deactivating a backup target	77

Increasing the size of an iSCSI backup target	78
Deleting a backup target	78
Managing backup policies	78
Viewing backup policy information	79
Editing a backup policy	79
Deleting a backup policy	80
Expiring backups manually	80
Adjusting the HYCU virtual machine resources	81
8 Managing HYCU users	82
Setting up user groups and users	83
Creating a new user group	83
Adding a new user	83
Activating or deactivating a user group or a user	84
Setting ownership of virtual machines	84
Assigning owners to virtual machines	84
Removing owners from virtual machines	85
9 Administering	86
Licensing	87
Creating a license request	88
Requesting and retrieving licenses	88
Activating licenses	89
Upgrading HYCU	90
Upgrading HYCU on a Nutanix AHV cluster	90
Upgrading HYCU on a Nutanix ESXi cluster	92
Changing network settings	94
Changing the HYCU listening port number	95
Configuring the SSL certificate	95
Creating a new self-signed certificate	96
Importing a CA certificate	96
Configuring FIPS-compliant mode for HYCU	96
Enabling FIPS-compliant mode for HYCU	97

Disabling FIPS-compliant mode for HYCU	97
Setting the iSCSI Initiator secret	98
Configuring backup target encryption	98
Setting power options	98
Configuring Active Directory authentication	99
Setting up logging	99
Accessing the HYCU backup controller virtual machine by using SSH	100
Using the command-line interface	102
Using the HYCU REST API Explorer	102
Enabling HTTPS for WinRM connections	102
Increasing the size of the HYCU virtual disk	103
Increasing the HYCU disk size in a Nutanix AHV cluster	104
Increasing the HYCU disk size in a Nutanix ESXi cluster	104
Removing HYCU	105
A Customizing HYCU configuration settings	106
How to customize HYCU configuration settings	106
Email notification settings	107
Snapshot settings	109
Utilization threshold settings	109
Display settings	109
SQL Server application settings	110
Settings for aborting jobs	110
Azure account settings	110
B Restoring to a different hypervisor	111
Restoring a Nutanix ESXi virtual machine to a Nutanix AHV cluster	111

Chapter 1

About HYCU

HYCU Data Protection for Nutanix (HYCU) is a high performing backup and recovery solution for Nutanix. It is the first data protection solution that is fully integrated with Nutanix, and makes data protection easy to deploy and simple to use.

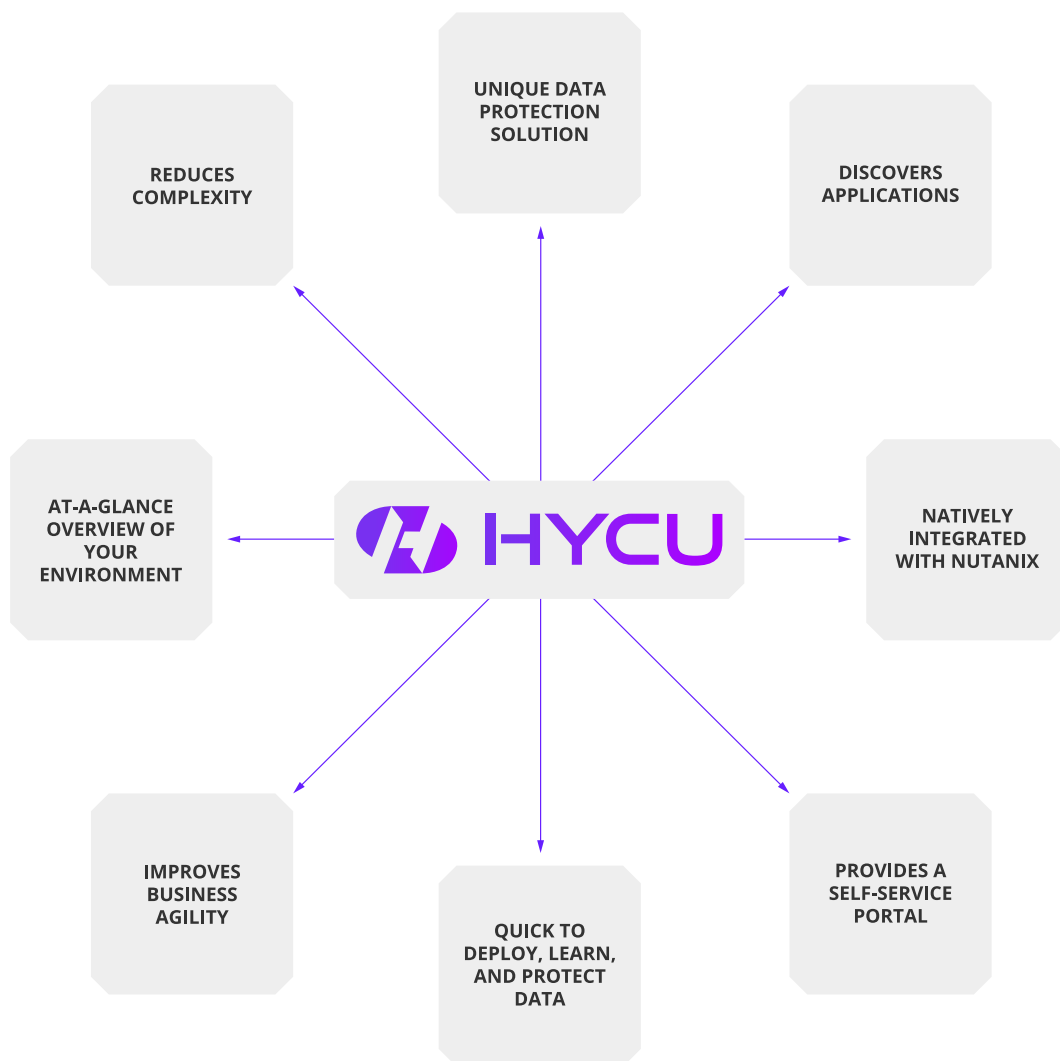


Figure 1-1: Introduction to HYCU

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 VMware vSphere Web Client (for Nutanix ESXi clusters).

- **Provides new-found visibility**

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

- **Protects data in a few minutes**

Data protection of virtual machines and applications can be enabled in a few minutes after the deployment.

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

Predefined backup policies (Gold, Silver, and Bronze) that come with HYCU simplify the data protection implementation. However, if the needs of the backup environment require, a wide range of opportunities to customize backup 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, and ensures that the entire application data is protected and recovered to a consistent state.

- **Lets you choose targets and hypervisors**

Using data storage targets and hypervisors 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.

HYCU backup environment overview

The HYCU backup environment is a set of components that discover, analyze, and protect the specified data on Nutanix clusters, and present it in the web-based console. The HYCU

environment consists of the following parts:

HYCU backup controller	A virtual machine hosted by a Nutanix cluster where HYCU resides. The HYCU backup controller processes the data collected in the Nutanix environment and presents it in the web user interface.
HYCU interface	A system from which you access HYCU. It also provides administration tools to adjust the environment to your needs. The HYCU interface has an established connection to the HYCU backup controller. You can use the HYCU web user interface or command-line interface (hyCLI).
HYCU targets	Storage locations that HYCU uses for storing the protected data.
Nutanix clusters	A backup environment for which HYCU provides data protection.

The following diagram shows the HYCU environment and its components:

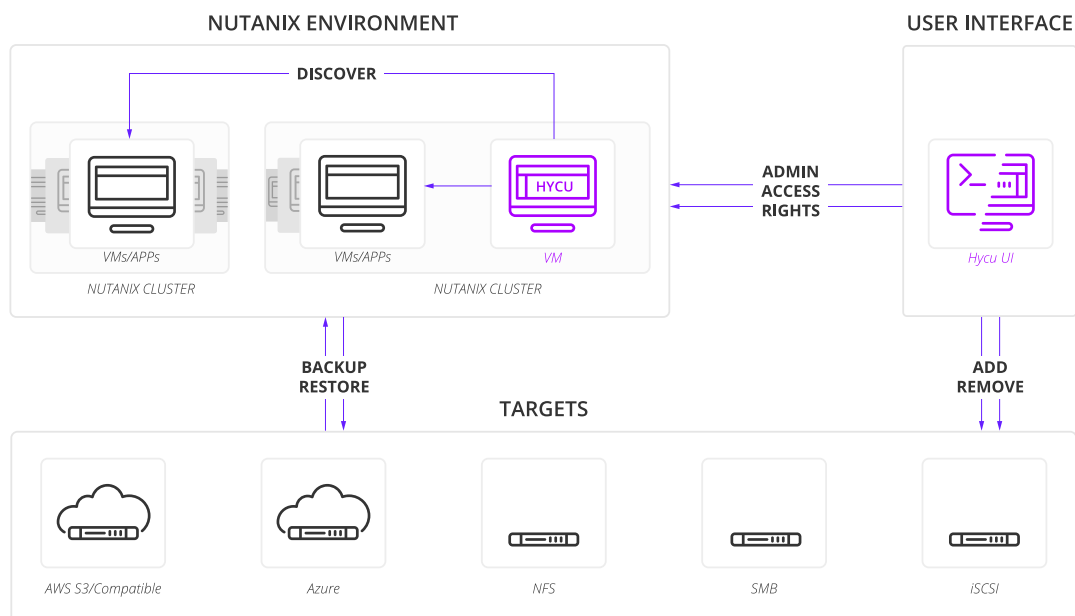


Figure 1-2: HYCU architecture

Chapter 2

Deploying the HYCU virtual appliance

The HYCU virtual appliance is a preconfigured HYCU software solution that you can easily deploy by using the Nutanix Prism web console (for Nutanix AHV clusters) or the VMware vSphere Client (for Nutanix ESXi clusters). The virtual appliance is distributed as a virtual disk image (for Nutanix AHV clusters) or an OVF package (for Nutanix ESXi clusters).

Make sure you first size the backup infrastructure for HYCU as described in [“Sizing your backup infrastructure for HYCU” below](#), and then follow the instructions in one of the following sections:

- [“Deploying HYCU on a Nutanix AHV cluster” on the next page](#)
- [“Deploying HYCU to a Nutanix ESXi cluster” on page 16](#)

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 17](#).

Sizing your backup infrastructure for HYCU

Before you start deploying the HYCU virtual appliance, size the backup infrastructure according to the following requirements:

- HYCU backup controller:

- Network connection

Make sure that you reserve an IP address for your virtual machine.

- System requirements

For HYCU backup controller deployment and configuration, at least 4 GB of RAM is required.

Ensure that your environment meets the following sizing requirements:

Size	Number of VMs	Storage	Cores	Memory
Small	100-200	20 GB-40 GB	4-6	4 GB

Size	Number of VMs	Storage	Cores	Memory
Medium	200-500	100 GB-200 GB	8-12	6-8 GB
Large	Contact HYCU Support.			

- HYCU web user interface:

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

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

- HYCU backup targets:

Make sure that destinations you want to use for storing your protected data are available and accessible.

Firewall settings

If a firewall is configured in your network infrastructure, make sure that the required ports are open on the following systems:

System	Port number	Description
HYCU backup controller	8443 (TCP)	For accessing the HYCU web user interface.
	22 (TCP)	For accessing the HYCU VM by using SSH.
	445 (TCP)	For a file or application restore.
HYCU backup target	2049 (TCP/UDP)	For accessing an NFS backup target.
	445 (TCP)	For accessing an SMB backup target.
	3260 (TCP)	For accessing an iSCSI backup target.
Nutanix clusters	9440 (TCP)	For accessing Nutanix REST API v3.
	3260 (TCP)	For accessing Acropolis Block Services (ABS).
Virtual machines	5985-5986 (TCP)	For accessing Windows virtual machines and applications (by using WinRM) for discovery.

Deploying HYCU on a Nutanix AHV cluster

HYCU deployment consists of several tasks that you must complete before you can start using HYCU for data protection. You must size the backup infrastructure for HYCU, upload the HYCU virtual machine image to a Nutanix cluster, create a virtual machine for your

HYCU deployment, and configure HYCU on the virtual machine, as shown in the following flowchart:






Figure 2-1: Overview of HYCU deployment tasks

To deploy HYCU, I need to...	Where can I find 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 the next page
3. Configure HYCU on the created virtual machine.	"Configuring HYCU on the virtual machine" on page 15

Uploading the HYCU virtual appliance image to a Nutanix AHV cluster

To upload the HYCU virtual appliance image to a Nutanix AHV cluster, follow these steps:

1. Log on to the Nutanix Prism web console by using your Nutanix logon credentials.
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. Optionally, enter an annotation.

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

`hycu-<version>-<revision>`

For example: `hycu-3.0.0-3634`

Make sure to leave out the `.qcow2` extension when entering the HYCU image name. If you enter the HYCU image name in a different format, you will not be able to use this image for an upgrade.


- b. From the Image Type drop-down menu, select **DISK**.
 - c. From the Storage Container drop-down menu, select a container for the image to be uploaded.
 - d. In the Image Source section, select one of the following:
 - **From URL**
Specify the location of the image file by using a URL.
 - **Upload a file**
Specify the location of the image file saved on your file system.
5. Click **Save**.
6. Click **Close** after the image is successfully uploaded.

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

To create a virtual machine for HYCU deployment on a Nutanix AHV cluster, in the Nutanix Prism web console, do the following:


1. In the menu bar, 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. Enter a virtual machine name and, optionally, its description.
- b. In the Compute Details section, enter the number of CPUs and cores per CPU, and a desired amount of RAM.
- c. In the Disks section, click **+Add New Disk**, and then, in the Add Disk dialog box, complete the following information:
 - 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 name of the image you uploaded.
 - v. In the Size (GiB) field, leave the default size of the virtual disk (10 GB).

 **Note** You can later increase the size of the HYCU virtual disk if needed. For details, see [“Increasing the HYCU disk size in a Nutanix AHV cluster” on page 104](#).

Click **Add**.

- d. Specify an additional data disk. Click **+Add New Disk**, and then, in the Add Disk dialog box, leave the default settings. In the Size (GiB) field, enter 32.

 **Note** You can later increase the size of the HYCU virtual disk if needed. For details, see [“Increasing the HYCU disk size in a Nutanix AHV cluster” on page 104](#).

Click **Add**.

4. In the Network Adapters (NIC) section, click **Add New NIC**, and then, in the Create NIC dialog box, select a VLAN from the VLAN Name drop-down menu. Click **Add**.
5. Click **Save**.


Configuring HYCU on the virtual machine

To configure HYCU on the created virtual machine, in the Nutanix Prism web console, do the following:

1. From the list of virtual machines, select the one you created, and then click **Power on**.
2. When the virtual machine is turned on, click **Launch Console**.


In the HYCU Network Configuration dialog box, do the following:

- a. Enter the values for the following:
 - *Optional*. Host name for the HYCU 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 (-).

- b. Click **OK**.
3. The progress of the HYCU backup controller configuration displays. After the HYCU backup controller is configured, confirm the summary message by clicking **OK**.

You can start using HYCU immediately with a prebuilt Instant-on license. This license expires automatically after 60 days and cannot be reused. Therefore, make sure to obtain a permanent license within this 60-day period. For instructions, see [“Licensing” on page 87](#).

Deploying HYCU to a Nutanix ESXi cluster

HYCU deployment includes creating the HYCU backup controller virtual machine in the VMware vSphere environment by deploying the HYCU virtual appliance package. Make sure you first size the backup infrastructure for HYCU.


Prerequisite

VMware vSphere 6.x only. Client Integration Plug-In is installed. For instructions, see the VMware documentation at: <https://docs.vmware.com/en/VMware-vSphere/6.0/com.vmware.vsphere.install.doc/GUID-CA16F78B-7890-4357-9760-AF8648806FE7.html>

To deploy HYCU to a Nutanix ESXi cluster, follow these steps:

1. Log on to the VMware vSphere Web Client.
2. Click the **VMs** tab, and then click **Deploy OVF Template....** The Deploy OVF Template dialog box opens.
3. In the Select template context, select **Local files**, and then click **Browse**.
4. Navigate to the HYCU virtual appliance package, select the HYCU disk image file (.vmdk) and the OVF template (.ovf), and then click **Open**.
5. In the Select template context, click **Next**.
6. In the Select name and location context, select the location where you want to deploy HYCU backup controller.

VMware vSphere 6.x only. Enter the HYCU image name in the format that should correspond to that of the HYCU image file you are uploading.

 **Important** The following format is required:
hycu-<version>-<revision>


For example: hycu-3.0.0-3634

Make sure to leave out the .qcow2 extension when entering the HYCU image name. If you enter the HYCU image name in a different format, you will not be able to use this image for an upgrade.


Click **Next**.

7. In the Select a resource context, select the Nutanix cluster where the HYCU backup controller will reside. Click **Next**.
8. In the Review details context, verify the template details. Click **Next**.
9. In the Select storage context and Select Network context, leave the default values, and then click **Next**.
10. In the Customize template context, enter the values for the following:

- *Optional*. Host name for the HYCU 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**.

11. In the Ready to complete context, review the settings, and then click **Finish**.
12. In the list of virtual machines on the Nutanix ESXi cluster, select the newly created HYCU virtual machine, and then click **Power on**.

 **Note** Creating the HYCU virtual machine may take a few moments. The Power on icon is enabled only after a virtual machine is created.

You can start using HYCU immediately with a prebuilt Instant-on license. This license expires automatically after 60 days and cannot be reused. Therefore, make sure to obtain a permanent license within this 60-day period. For instructions, see [“Licensing” on page 87](#).

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*.

To log on to HYCU, follow these steps:

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

```
https://<server_name>:8443
```

In this instance, `<server_name>` 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** The level of access depends on your HYCU user permissions. For details, see [“Managing HYCU users” on page 82](#).

You can now start establishing your backup environment to enable data protection. For more information, see [“Establishing a backup environment” on page 19](#).

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 102](#).

Chapter 3

Establishing a backup environment

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

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

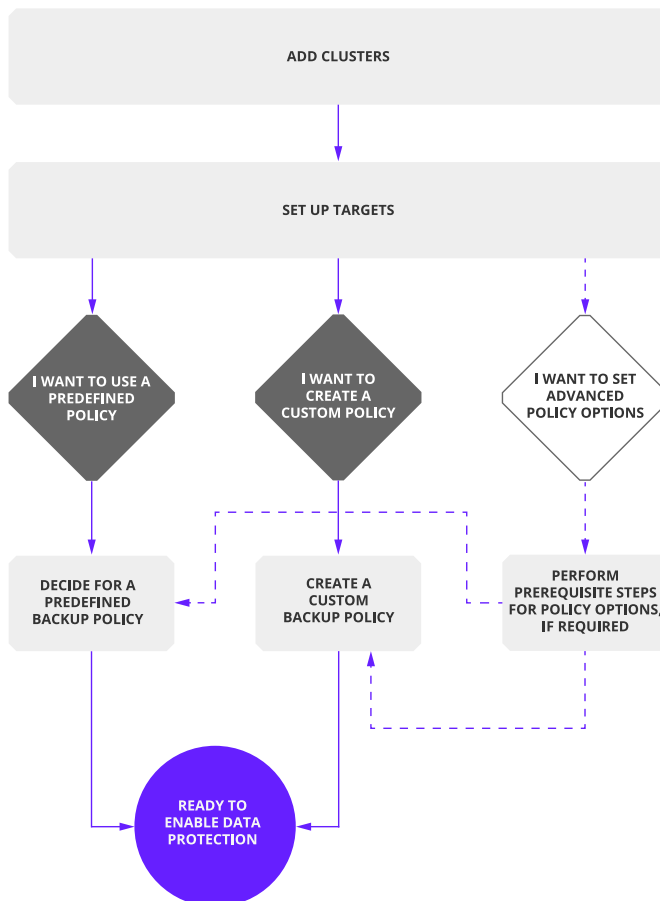


Figure 3–1: Establishing a backup environment

The tasks that are required to establish a backup environment can be performed only by an administrator, and are as follows:

- [“Adding Nutanix clusters” below](#)
- [“Setting up backup targets” on the next page](#)

Depending on the backup environment, the following task may also have to be performed:

- Creating your own backup policies
 - If you do not want to use any of the predefined backup policies that come with HYCU, you must create your own backup policies. For details, see [“How to create a custom backup policy” on page 30](#).

After the backup environment is established, data protection can be accomplished in several ways to fulfill the needs of particular business. For details, see [“Protecting data” on page 37](#).

Adding Nutanix clusters


A Nutanix environment consists of one or more Nutanix clusters, each of which hosts a series of virtual machines that run applications. You can add one or more Nutanix clusters that host virtual machines you want to include in the backup.

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.

Nutanix ESXi cluster prerequisite



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


Accessing the Nutanix Clusters dialog box

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

To add a Nutanix cluster, follow these steps:

1. In the Nutanix Clusters dialog box, click **+ New**. The New dialog box appears.
2. Enter the name of the Nutanix cluster in URL format: `https://<server_name>:<port>`.
3. Enter the user name and password of the user with cluster administrative rights.
4. 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**).

 **Important** In Nutanix environments that use VMware ESXi hypervisors, 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.

Setting up backup targets

Backup targets are locations where the protected data is stored. HYCU allows you to store your data to AWS S3 or S3-compatible storage, Azure storage, an NFS share, an SMB share, and an iSCSI storage device.

The approach to set up backup targets is common for different target types. However, there are specific prerequisites and steps that are required for each target type. Depending on which backup target you want to set up, see one of the following sections:

- [“How to set up an AWS S3/Compatible target” below](#)
- [“How to set up an Azure target” on the next page](#)
- [“How to set up an NFS target” on page 23](#)
- [“How to set up an SMB target” on page 24](#)
- [“How to set up an iSCSI target” on page 26](#)

How to set up an AWS S3/Compatible target


Prerequisites

- The target is configured and accessible.
- The bucket is created.

Recommendation

It is recommended that the backup target is dedicated only to a single HYCU backup controller.

Accessing the Targets panel


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

To set up an AWS S3/Compatible target, follow these steps:

1. In the Targets panel, click **+ New**. The New Target dialog box appears.
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 MB, GB, or TB).
 - 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.

3. In the Target section, do the following:
 - a. From the Type drop-down menu, select **AWS S3/Compatible**.
 - 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.
4. Click **Save**.

The backup target is added to the list of targets. For details on managing backup targets, see [“Managing backup targets” on page 75](#).

How to set up an Azure target


Prerequisites

- The target is configured and accessible.
- *Only if using an Azure Government account.* The `target.azure.government.cloud` configuration setting is enabled. For details, see [“Azure account settings” on page 110](#).

Recommendation

It is recommended that the backup target is dedicated only to a single HYCU backup controller.

Accessing the Targets panel


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

To set up an Azure backup target, follow these steps:

1. In the Targets panel, click **+ New**. The New Target dialog box appears.
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 MB, GB, or TB).
 - 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 the Azure cloud, see [“Archiving data to the Azure archive storage tier” on page 35](#).

3. In the Target section, do the following:
 - a. From the Type drop-down menu, select **Azure**.
 - 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.

4. Click **Save**.

The backup target is added to the list of targets. For details on managing backup targets, see [“Managing backup targets” on page 75](#).

How to set up an NFS target


Prerequisites

- The target is configured and accessible.
- There is enough free space to store the backup data in the target location.
- If deduplication is enabled on the backup target, the backup target is dedicated exclusively to HYCU backups. By dedicating a backup 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.

Recommendation

It is recommended that the backup target is dedicated only to a single HYCU backup controller.


Accessing the Targets panel

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

To set up an NFS backup target, follow these steps:

1. In the Targets panel, click **+ New**. The New Target dialog box appears.
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 MB, GB, or TB). If your backup 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 backup 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.

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.

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

- The deduplication ratio may be affected by it (in cases where the backup target has deduplication enabled).
- You cannot use this backup target for an internal backup.

4. Click **Save**.

The backup target is added to the list of targets. For details on managing backup targets, see [“Managing backup targets” on page 75](#).

How to set up an SMB target

Prerequisites


- The target is configured and accessible.
- There is enough free space to store the backup data in the target location.
- If deduplication is enabled on the backup target, the backup target is dedicated exclusively to HYCU backups. By dedicating a backup 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*.

Recommendation


It is recommended that the backup target is dedicated only to a single HYCU backup controller.


Accessing the Targets panel

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

To set up an SMB backup target, follow these steps:

1. In the Targets panel, click **+ New**. The New Target dialog box appears.
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 MB, GB, or TB). If your backup 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 backup 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.
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.

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

- The deduplication ratio may be affected by it (in cases where the backup target has deduplication enabled).
- You cannot use this backup target for an internal backup.

4. Click **Save**.


The backup target is added to the list of targets. For details on managing backup targets, see [“Managing backup targets” on page 75](#).

How to set up an iSCSI target


Prerequisites

- The target is configured and accessible.
- The target has not been initialized yet.
- The iSCSI storage device is dedicated only to a single HYCU backup controller.

⚠ Caution Disregarding this prerequisite may result in data loss or corruption. Therefore, make sure to avoid the following scenarios:

- Several HYCU backup controllers are using the same backup target simultaneously.
 - Any appliance other than HYCU and HYCU itself are using the same backup target simultaneously.
- The HYCU iSCSI Initiator secret is added on the iSCSI server if you want to enable mutual authentication between HYCU and the iSCSI server.
 - For improved backup and restore performance, the iSCSI Data Service IP address is specified on the Nutanix cluster by using the Prism console ( > **Cluster Details**). This automatically enables the Nutanix load balancing feature during backup and restore, which eliminates heavy I/O load on the Nutanix cluster and containers. For details, see Nutanix documentation.

Accessing the Targets panel

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

To set up an iSCSI backup target, follow these steps:


1. In the Targets panel, click **+ New**. The New Target dialog box appears.
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 MB, GB, or TB).

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

3. In the Target section, do the following:

- a. From the Type drop-down menu, select **iSCSI**.
- b. Enter the target portal IP address (the External Data Service IP address if it is specified on Nutanix, otherwise, the Nutanix cluster IP) and the target name (the iSCSI Qualified Name (IQN) or Extended Unique Identifier (EUI) of the iSCSI storage device that you can acquire on the iSCSI server).

 **Note** If data from sources other than HYCU resides on the storage device, such 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.

 **Note** If you enable target encryption, you cannot use this backup target for an internal backup.

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 98](#).

5. Click **Save**.

The backup target is added to the list of targets. For details on managing backup targets, see [“Managing backup targets” on page 75](#).

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

- **Recovery Point Objective (RPO)**
RPO is the maximum tolerable data loss interval (in months, weeks, days, hours, or minutes). This is a time period in which data can be lost after an event that causes a virtual machine or an application to go offline occurs.
- **Recovery Time Objective (RTO),**
RTO is the maximum tolerable time of disruption (in months, weeks, days, hours, or minutes). This is a time period that is needed before a virtual machine or an application is brought back online.

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

- **Applying a predefined backup policy**
You can use any of the predefined backup policies (Gold, Silver, or Bronze) to simplify the data protection implementation. For details, see [“Applying a predefined backup policy” below](#).
- **Creating a custom backup policy**
If none of the predefined backup policies meets the needs of your environment, you can create a new backup policy and tailor it to your needs. For details, see [“Creating a custom backup policy” on the next page](#).

If you consider one of the predefined or custom backup policies satisfies all data protection goals of your environment, you can set such a policy as default. For details, see [“Setting a default backup policy” on page 36](#).

Applying a predefined backup policy

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

HYCU comes with the following predefined backup policies:

Type of predefined backup 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 backup sources from being backed up, you can use the Excluded backup policy.

Creating a custom backup policy

If the needs of your environment are not covered with any of the predefined backup policies, you can create a new backup policy and tailor it to your needs. While tailoring a backup policy to your needs and setting the desired RPO, RTO, and backup 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 run all backup jobs within a specific time frame to improve effectiveness and avoid overload of your environment.
Backup from replica	<p>Allows you to back up your virtual machines from their replicas in remote office/branch office (ROBO) environments. To be able to do so, you must create a protection domain that will include the virtual machines that you want to protect and specify the schedule, retention, and remote sites for replicating the virtual machines.</p> <p>Make sure that the replication retention on the Nutanix cluster is adjusted to the backup policy's RPO. 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. For example, if the Nutanix schedule interval is two hours and the RPO of the HYCU backup policy is eight hours, the retention policy for the remote site must be set to 4 or more snapshots (that is, at least the last four snapshots must be kept).</p> <p>For details on protecting virtual machines through the Nutanix Prism web console, see Nutanix documentation.</p>
Archiving	Allows you to preserve your data for future reference.
Copy	Allows you to create a copy of backup data.
Fast restore	<p>Allows you to restore a virtual machine or an application to the original container in a fast way by keeping local snapshots on the Nutanix cluster for the specified retention time.</p> <p>By default, HYCU keeps one snapshot on the Nutanix cluster. With this option enabled, HYCU will keep more than one snapshot on the Nutanix cluster (depending on your retention settings), which allows you to restore a virtual machine or an application in a fast way, reducing downtime.</p>

How to create a custom backup policy

You can create a custom backup 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 [“How to create a backup window” on page 32](#).
- 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 [“How to create a data archive” on page 34](#).
- *For backing up from replicas only.* A protection domain that includes the virtual machines that you want to protect is created and the schedule, retention, and remote sites for replicating the virtual machines are specified. For details on protecting virtual machines through the Nutanix Prism web console, see Nutanix documentation.

Limitation

For backing up from replicas only. If your Nutanix cluster is running AHV, you can back up virtual machines from both AHV and ESXi branch office environments. However, if your Nutanix cluster is running ESXi, you can back up virtual machines only from AHV branch office environments.


Accessing the Policies panel

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

To create a custom backup policy, follow these steps:

1. In the Policies panel, click **+ New**. The New Policy dialog box appears.
2. Enter a name and, optionally, a description of your backup policy.
3. Add any of the following policy options to the list of the enabled options by clicking it:
 - **Backup** (*enabled by default*)
 - **Backup window**
 - **Backup from replica** (*available only if the Fast restore option is disabled*)
 - **Archiving**
 - **Copy**
 - **Fast restore** (*available only if the Backup from replica option is disabled*)
4. Depending on which policy options you have enabled, do the following:

Enabled option	Procedure
Backup	To back up data, in the Backup section, do the following:

Enabled option	Procedure
	<ul style="list-style-type: none"> 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. d. From the Targets drop-down menu, select one or more backup 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 targets that can guarantee compliance with the RPO and RTO policy settings. Backup 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 backup source 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 fastest. <p> 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.</p> <ul style="list-style-type: none"> e. In the Backup threshold field, specify a backup threshold value of 0 through 100 percent (the default value is 25%). When the amount of changes on a virtual machine since the last full backup reaches the specified value, a new full backup is performed instead of an incremental one.
Backup window	<p>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 “How to create a backup window” on the next page.</p> <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>
Backup from replica	<p>To back up virtual machines from their replicas, in the Backup from replica section, from the Central site cluster drop-down menu, select</p>

Enabled option	Procedure
	the cluster on which the replicas of your virtual machines reside.
Archiving	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 “How to create a data archive” on page 34.
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 backup targets that you want to use for storing the copy of backup data. If you want your target to be selected automatically, make sure the Automatically selected option is selected. The backup copy target will be different from the backup target for data safety reasons. <p>Note When there are several backup 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 backup targets based on the estimated size of queued and running backups on them.</p>
Fast restore	<p>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>

5. Click **Save**.


The custom backup policy is created and added to the list of backup policies. For details on managing backup policies, see [“Managing backup policies” on page 78.](#)

How to create a backup window


HYCU enables you to define a time frame when your backup jobs are allowed to run. You can use this backup window to improve effectiveness and avoid 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 backup policies and custom backup policies.


 **Note** If you use a backup window, the backup jobs run only during the specified hours. Make sure that the RPO specified in the affected policy can be achieved within this backup window. If the RPO is shorter than the largest time frame in which backups do not run, such RPO cannot be achieved during the hours out of the backup window.


Accessing the Policies panel

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



To create a backup window, follow these steps:

1. In the Policies panel, click  **Backup Window**.
2. In the Backup Window dialog box that appears, click **+ New**.
3. In the New dialog box that appears, enter a name for your backup window and, optionally, a description.
4. Select the week days and hours during which you want backups to run. You can click and drag to quickly select a time frame that includes the days and hours you want to add.


 **Important** All scheduled backup jobs are run based on the HYCU backup controller time zone.

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.

5. Click **Save**.
6. 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 [“How to create a custom backup policy” on page 30](#).
- Assign a backup window to the existing backup policy. To do so, select the backup policy, click  **Edit**, and then make the required modifications.

Example

You have selected the Bronze policy and specified the time frame for the backup jobs to be from Monday to Friday from 6 PM to 6 AM, and from Saturday to Sunday all day long.

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

How to create 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).
- *Only for archiving data to the Azure archive storage tier.* Data archives are stored in the Azure cloud with the Blob Storage or General Purpose v2 (GPv2) accounts.


Azure archive storage tier limitations

- General Purpose v1 (GPv1) accounts do not support moving data archives to the archive storage tier.
- 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**.

To create a data archive, follow these steps:


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.
6. Provide information about when to archive data, the retention period to be used, and the archive target.
For details on how HYCU manages archiving data to the Azure cloud, see [“Archiving data to the Azure archive storage tier”](#) below.
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 [“How to create a custom backup policy”](#) on page 30.
- Assign a data archive to the existing backup policy. To do so, select the backup policy, click  **Edit**, and then make the required modifications.

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. 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.


If you are using the hot storage tier for archiving data, you get a warning that the cool storage tier is recommended for archiving.

 **Important** When restoring data archives, HYCU performs a data rehydration task 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 task can take a few hours to complete.


Setting a default backup policy

You can select one of the predefined or custom backup policies to be the default backup policy for your HYCU environment. After you set a default backup policy, it is assigned to all existing backup sources 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**.

To set a default backup policy, follow these steps:

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 backup policy to all backup sources that do not have an assigned policy (that is, existing and newly discovered ones).
 - Click **No** if you want to assign the default backup policy only to newly discovered backup sources.

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

Chapter 4

Protecting data

With the HYCU backup and recovery 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 back up virtual machines residing on Nutanix clusters and applications running on virtual machines. After you establish your backup environment (that is, add Nutanix clusters, set up backup targets, and, optionally, create backup policies), you can enable data protection. When you complete the first backup, you can restore the data that is backed up if it becomes damaged or corrupted.

Because HYCU is application-aware, when you set credentials for virtual 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. To ensure application consistency, HYCU provides the application-aware backup and restore.

The approach you choose for backing up your data largely depends on the type of restore you want to perform. You may want to perform the restore at the virtual machine or application level, or be able to restore only specific files inside the virtual machines. HYCU provides different levels of data consistency depending on your restore strategy.

The following flowchart explains the process of enabling data protection:

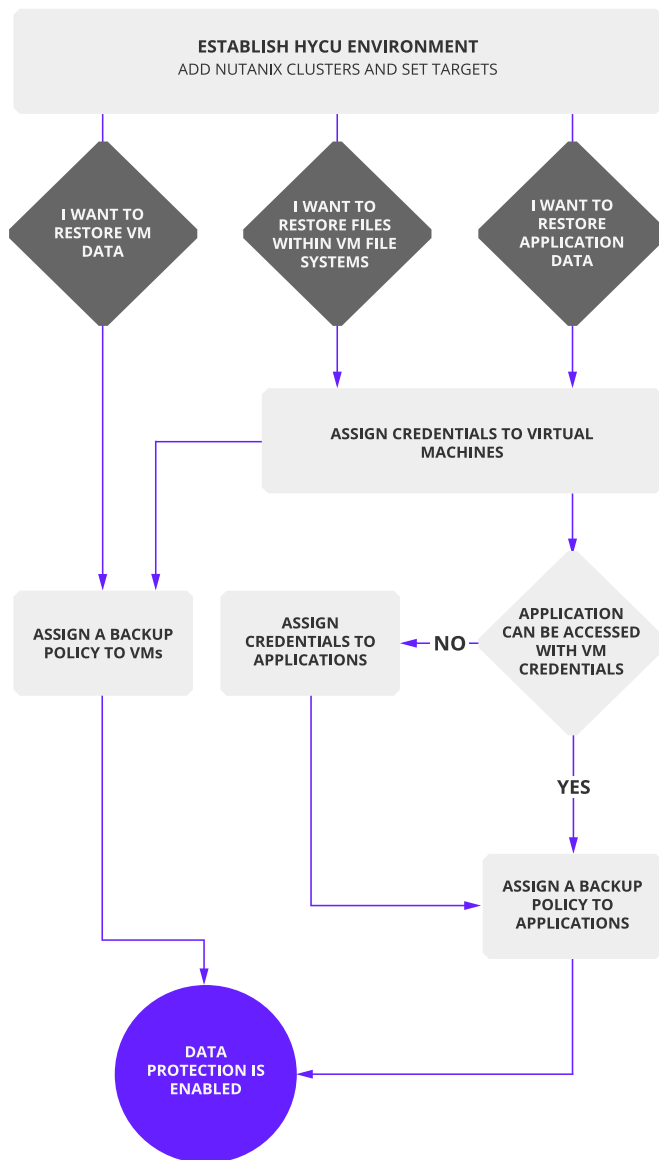


Figure 4-1: Enabling data protection

The following table explains which approach you should use for protecting your data and provides quick access to the backup instructions:

What do I want to restore?	Which type of backup should I select?	Which tasks do I need to perform?
Virtual machine data	Basic backup	1. Back up virtual machines. For instructions, see “Backing up virtual machines” on page 42.

What do I want to restore?	Which type of backup should I select?	Which tasks do I need to perform?
Individual files within virtual machine file systems	File-level backup	<ol style="list-style-type: none"> 1. Enable access to data. For instructions, see “Enabling access to data” below. 2. Back up virtual machines. For instructions, see “Backing up virtual machines ” on page 42.
Application data	Application-aware backup	<ol style="list-style-type: none"> 1. Enable access to data. For instructions, see “Enabling access to data” below. 2. Back up applications. For instructions, see “Backing up applications” on page 43.

Enabling access to data

When the recovery goals of your environment require backing up data inside the virtual machine file systems, you must enable HYCU to access the files and applications residing on the virtual machines.

Prerequisites

Before you start enabling access to data, make sure the following prerequisites are met:

- On Windows 7 and 2008 R2, Windows PowerShell 3.0 is installed. For an application-aware backup, the Windows PowerShell Script Execution Policy (Set-ExecutionPolicy) is set to RemoteSigned.
- On Windows 7, 8, and 10, and Windows Server 2008 R2, WinRM is enabled and configured by using the `winrm quickconfig` command.
- Windows user account with WinRM permissions exists. For an application-aware backup, this account should have access to the application.
- On Linux, port 22 is open and the SSH daemon is running.
- For the Nutanix cluster running on ESXi: VMware Tools and the Nutanix Guest Tools (NGT) software bundle 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.

For detailed information, see Microsoft documentation.

Oracle-specific considerations

- 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.
- The OS user must have sudo privileges.


To enable access to files and applications residing on the virtual machines, complete these tasks:

1. For a file-level backup and an application-aware backup, provide access to files inside virtual machines or application data. For details, see [“Assigning credentials to virtual machines” below](#).
2. *Not applicable for Active Directory.* For an application-aware backup, provide access to application data if the discovered applications do not use virtual machine credentials. For details, see [“Assigning credentials to applications” on the next page](#).


Assigning credentials to virtual machines



A file-level backup enables you to back up virtual machines in environments where you foresee that you will want to restore only specific files inside the virtual machines.


Accessing the Virtual Machines panel

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

To be able to perform a file-level backup, assign credentials to virtual machines:

1. Select which virtual machines you want to back up.
2. Click  **Credentials**. The Credential Groups dialog box appears.
3. Click **+ New**.
4. Enter the credentials that are required for accessing the virtual machines.
5. Click **Save**.
6. Click **Assign**.


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

Application discovery

The process of application discovery starts automatically after you assign credentials to virtual machines. HYCU can discover the following applications that are running on virtual machines:

- SQL Server
- Active Directory

 **Note** The following roles are supported for Active Directory: Active Directory Domain Services, Active Directory Lightweight Directory Services, Active Directory Certificate Services, Active Directory Federation Services, and Active Directory Rights Management Services.

- Exchange Server
- Oracle


For a list of supported application versions, see the *HYCU Compatibility Matrix*.

When the application discovery job completes, the discovered applications are listed in the Applications panel.


Assigning credentials to applications

If you can access the discovered application you want to protect with the virtual machine credentials, the Discovery status for these applications in the Application panel is green and you can start protecting such applications. For instructions, see [“Backing up applications” on page 43](#).


If the virtual machine credentials do not have proper permissions to access any of the discovered applications that you want to protect, the Discovery status for these applications in the Application panel is red. For such applications, you must assign specific credentials with permissions to access them.

 **Note** Because access to Active Directory is always granted with the virtual machine credentials, you do not need to assign credentials to it.

Accessing the Applications panel

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

To assign credentials to applications, follow these steps:

1. Select which applications you want to back up.
2. Click  **Credentials**. The Credentials dialog box appears.
3. Use the switch to enable using operating system credentials with enough permissions to access the application.
4. Enter credentials for the user account with required permissions and access to the application. 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 or Backup Operators group.
- *SQL Server only.* 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.

5. Click **Save**.


Backing up virtual machines

HYCU provides you with two types of virtual machine backup, a basic backup and a file-level backup. Both these types enable you to back up your virtual machine data in a fast and efficient way. The only difference is that for the file-level backup, where the recovery goals of your environment require backing up data inside the virtual machine file systems, you must make sure that access to data is provided. For detailed information, see [“Enabling access to data” on page 39](#).

Limitation



- Only a backup of local fixed disks 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.


Accessing the Virtual Machines panel


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

To back up virtual machines, follow these steps:

1. 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 in panels” on page 71](#).

2. Click  **Policies**. The Policies dialog box appears.
3. From the list of policies, select the desired backup policy.

 **Note** If the backup policy is already assigned to the applications that are running on the selected virtual machines, a dialog box appears that enables you to automatically unassign the backup policy from the applications and assign it to the selected virtual machines. In this case, application consistency will not be mandatory for a successful virtual machine backup.

4. Click **Assign** to assign the backup policy to the selected virtual machines.

After you assign the backup policy, the backup is scheduled according to the values that you defined for your backup policy.

 **Note** If required, you can also perform a manual backup at any time. For details, see [“Performing a manual backup” on page 45](#).

Backing up applications

An application-aware backup allows a consistent backup of the following applications:

- SQL Server
- Active Directory
- Exchange Server
- Oracle

Prerequisites

Before you start backing up applications, the following prerequisites must be met:


SQL Server	<ul style="list-style-type: none"> • Databases reside on the local disks in the Nutanix environment. • Credentials are assigned to applications. For detailed information about assigning credentials to applications, see “Enabling access to data” on page 39. • <i>If you plan to restore 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>If you plan to restore databases that are part of an AlwaysOn Availability Group.</i> Either all nodes in the AlwaysOn Availability Group are protected by HYCU or only the node with the synchronized databases of the AlwaysOn 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.
Active Directory	<p>The Nutanix Guest Tools (NGT) software bundle is installed on the client virtual machine. For detailed information about installing NGT, see Nutanix documentation.</p>
Exchange Server	<ul style="list-style-type: none"> • The Nutanix Guest Tools (NGT) software bundle is installed on the client virtual machine. For detailed information about installing NGT, see Nutanix documentation. • All databases are mounted. • Credentials are assigned to applications. For detailed information about assigning credentials to applications, see “Enabling access to data” on page 39. • 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</p>

	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 is deleted, you need to restore only the Exchange Server application.
Oracle	The SSH service is enabled on the Oracle server and listening on port 22, which is open for incoming connections.

Limitations


Before you start backing up applications, keep in mind the following application-specific limitations:

SQL Server	<ul style="list-style-type: none"> • Backing up SQL Servers with failover clusters is not supported. Consequently, assigning policies to such applications is not possible. • 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.
Oracle	You can protect data only for single-instance Oracle databases. Backing up Oracle Real Application Clusters (RAC) databases is not supported. Consequently, assigning policies to such databases is not possible.

 **Note** If Active Directory and Exchange Server applications are running on the same virtual machine and you plan to use the same approach for protecting both applications, you can assign a backup policy only to the Exchange Server application. In this case, the state of Active Directory application is backup consistent state and it is backed up together with Exchange Server automatically.


After you make sure that all the prerequisites are met and that you are familiar with all the limitations, you can continue with backing up applications.


Accessing the Applications panel


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

To back up applications, follow these steps:

1. In the Applications panel, select the 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 in panels](#)” on page 71.

2. Click  **Policies**. The Policies dialog box appears.
3. From the list of policies, select the desired backup policy.

 **Note** If the backup policy is already assigned to the virtual machines on which the selected applications are running, a dialog box appears that enables you to automatically unassign the backup policy from the virtual machines and assign it to the selected applications. In this case, the virtual machines on which the applications are running are also protected by the same backups.

4. Click **Assign** to assign the backup policy to the selected applications.


After you assign the backup policy to the selected applications, the backup is scheduled according to the values that you defined for your backup policy.


 **Note** If required, you can also perform a manual backup of any application at any time. For details, see [“Performing a manual backup” below](#).

Performing a manual backup

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

To perform a manual backup, follow these steps:

1. In the Virtual Machines or Applications panel, select which backup sources you want to back up.
2. Click  **Backup** to perform the backup of the selected backup sources.
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.

Chapter 5

Restoring data

You can start restoring data when at least one successful backup is performed. HYCU enables you to restore virtual machines, virtual machine disk files, individual files within virtual machine file systems, applications and application items to any available restore point. Depending on what kind of data you want to restore, see one of the following sections:

- [“Restoring virtual machine data” below](#)
- [“Restoring application data” on page 52](#)

Restoring virtual machine data


You can perform a restore of virtual machine data on a virtual machine level or a file level. This means that you can restore an entire virtual machine or only specific files within virtual machine file systems. Depending on what kind of virtual machine data restore you want to perform, see one of the following sections:

- [“Restoring an entire virtual machine” below](#)
- [“Restoring individual files” on page 50](#)

Restoring an entire virtual machine


When restoring an entire virtual machine, you can select among the following restore options:

Restore option	Description
Restore VM^a	Enables you to restore a virtual machine to its original or a new location with original settings. Select this option if you want to replace the original virtual machine with the restored one. In this case, the original virtual machine is deleted automatically before the restore. For instructions, see “How to restore a virtual machine” on the next page .
Clone VM	Enables you to create a clone of the original machine by restoring it to a new location with custom settings. Select this option if you want to keep the original virtual machine. For instructions, see “How to clone a virtual machine” on page 48 .

Restore option	Description
	<p> Note By using the Clone VM option, you can also restore a Nutanix ESXi virtual machine to a Nutanix AHV cluster. For details, see “Restoring to a different hypervisor” on page 111.</p>
Restore VM Disk Files	<p>Enables you to restore virtual machine disk files to an NFS or SMB shared location. Select this option if you want to use the virtual machine disk files later to create a new virtual machine. For instructions, see “How to restore virtual machine disk files” on the next page.</p>

^aNot supported for Nutanix ESXi clusters.

Accessing the Virtual Machines panel


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


How to restore a virtual machine

Limitation


Restoring a virtual machine running on a Nutanix ESXi cluster is not supported.

To restore a virtual machine to its original location or a new location with original settings, follow these steps:

1. In the Virtual Machines panel, click the virtual machine that you want to restore to open the Details section. The Details section appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Details section.
2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Click  **Restore VM**. The VM Restore Options dialog box opens.
4. Select **Restore VM**, and then click **Next**.
5. From the drop-down menu, select where you want to restore the virtual machine.

 **Note** By default, the original location is selected. If you decide to restore the virtual machine to another container, it will not be restored from the snapshot, but from the target. Therefore, no fast restore will be performed.

6. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore. The original virtual machine is deleted automatically.
7. Click **Restore**.

 **Note** *Nutanix ESXi clusters only.* Because the minimum RAM required for restoring a virtual machine is 256 MB, any virtual machine with less RAM is automatically set to 256 MB during the restore.

How to clone a virtual machine

To create a clone of the original virtual machine, restore it to a new location with custom settings. To do so, follow these steps:

1. In the Virtual Machines panel, click the virtual machine that you want to restore to open the Details section. The Details section appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Details section.
2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Click **Restore VM**. The VM Restore Options dialog box opens.
4. Select **Clone VM**, and then click **Next**.
5. From the drop-down menu, select where you want to restore the virtual machine.

Note By default, the original location is selected. If you decide to restore the virtual machine to another container, it will not be restored from the snapshot, but from the target. Therefore, no fast restore will be performed.
6. Specify a new name for the virtual machine.
7. 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.
8. Click **Restore**.

After cloning a virtual machine


If you restored a virtual machine to a new location on another Nutanix cluster, make sure to add a network adapter to the virtual machine afterward. For details on how to do this, see Nutanix documentation.

How to restore virtual machine disk files

You can restore virtual machine disk files to a shared location. Access to the shared location can be limited to specific accounts in your organization that can later use these disk files to create a new virtual machine. By doing so, you provide greater flexibility when restoring data and optimize restore data management. To restore virtual machine disk files, follow these steps:

1. In the Virtual Machines panel, click the virtual machine that you want to restore to open the Details section. The Details section appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Details section.
2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Click **Restore VM**. The VM Restore Options dialog box opens.

4. Select **Restore VM disk files**, and then click **Next**.


 **Important** During the restore of virtual machine disk files, you cannot perform additional restores or expire backups for this virtual machine.

5. From the drop-down menu, select where you want to restore the disk files, and then provide the required information:
 - **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).
 - **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).
6. Click **Restore**.

After restoring virtual machine disk files


After the restore of the virtual machine disk files is complete, you can upload them to the Nutanix image service (for Nutanix AHV clusters) or to the VMware datastore (for Nutanix ESXi clusters) and create a new virtual machine by using them. You do not need HYCU to create a virtual machine from the restored disk files.

 **Note** Keep in mind that due to features such as data resiliency and provisioning, the disk files may require more space than on the shared location. Therefore, make sure to reserve enough space in the Nutanix container or the VMware datastore. For details on how to do this, see Nutanix or VMware documentation.

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.

 **Important** *Nutanix ESXi clusters only.* Because the format in which the data is restored is not a native VMware format, you must convert the disk image file to the VMDK format before uploading it to the VMware datastore. To do so, follow these steps:

1. Navigate to the folder where the disk image file is located.
2. As the root user, run the following command:

```
qemu-img convert -f raw -O vmdk <source_disk_image> <target_disk_image>.vmdk
```

Restoring individual files

You can restore individual files within virtual machine file systems to their original or a new location on the original Windows or Linux virtual machine, or to a shared SMB location. This alternative to restoring an entire virtual machine allows you to restore only one or more individual files (that have been deleted for some reason and are now missing on the virtual machine) without the need to restore all virtual machine data.


Prerequisites

- *When restoring to the original virtual machine.* You have enabled access to the virtual machine file system. For details, see [“Enabling access to data” on page 39](#).
- One of the following file systems is used:
 - **On Windows:** NTFS, FAT, or FAT32
 - **On Linux:** xfs, ext4/ext3/ext2, reiserfs, or btrfs



Limitations

- Performing a file-level restore on dual-boot systems is not supported.
- On Linux, a file-level restore of symbolic links and soft links is performed only when restoring to the original location.

Accessing the Virtual Machines panel

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

To restore individual files, follow these steps:

1. In the Virtual Machines panel, click the virtual machine that contains the files that you want to restore to open the Details section. The Details section appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Details section.
2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Depending on whether the snapshot for the selected restore point is online, do one of the following:
 - If the snapshot is online, the  **Restore Files** option is available and you can start the procedure for restoring the files by clicking it.
 - If the snapshot is not online, you first need to prepare the files for the restore:
 - a. Click  **Prepare for Restore Files**. The Prepare for Restore Files dialog box appears.
 - b. Use the **Restore with original settings** switch if you want to restore data with original settings.

If you choose to use different settings for restoring data, select a container.

c. Click **Prepare**.

The **Restore Files** option becomes available and you can start the procedure for restoring the files by clicking it.

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.

The Single File Restore dialog box appears.

5. Select where you want to restore the individual files:

Original virtual machine	<p>To restore the individual files to the original virtual machine, follow these steps:</p> <ol style="list-style-type: none"> Click Virtual machines, and then click Next. Select the location on the virtual machine where you want to restore the individual files, and provide the required information: <ul style="list-style-type: none"> Original location Select how the restore should save the files when there is a file with the same name and location on the virtual machine (overwrite the file, rename the original file, or rename the restored file). Alternate location (on the same VM) Specify the path to an alternate location on the same virtual machine in the following format: <code>C:\<path></code> Use the Restore ACL switch if you want to restore the original access control list. <p>Important If the original 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 Nutanix cluster), you cannot select it for restoring the individual files.</p>
SMB share	<p>To restore the individual files to an SMB share, follow these steps:</p>

	<ol style="list-style-type: none"> a. Click Fileshares, and then click Next. b. Specify the path to a shared folder in the following format: <div style="background-color: #f0f0f0; padding: 2px; margin: 5px 0;"><code>\\server\<i><path></i></code></div> c. <i>Optional</i>. Provide user credentials to access the SMB share.
--	--

6. Click **Restore**.


Restoring application data

You can perform a restore of application data on an application level or an application item level. This means that you can restore a whole application or only specific application items for which you have performed an application-aware backup. Depending on what kind of application data restore you want to perform, see one of the following sections:

- [“Restoring a whole application” below](#)
- [“Restoring application items” on page 56](#)

Restoring a whole application

With HYCU, you can restore a whole application to its original or a new location by restoring the virtual machine on which the application is running. In addition, you can restore the disk files of the virtual machine on which the application is running to a shared location.

 **Note** *Active Directory only*. HYCU does not perform an authoritative restore.

When restoring a whole application, you can select among the following restore options:

Restore option	Description
Restore VM^a	Enables you to restore an application by restoring the virtual machine on which it is running to its original or a new location with original settings. Select this option if you want to replace the original virtual machine on which your application is running with the restored one. In this case, the original virtual machine is automatically deleted before the restore. For instructions, see “How to restore a virtual machine” on the next page .
Clone VM	Enables you to create a clone of the original machine by restoring it to a new location with custom settings. Select this option if you want to keep the original virtual machine on which your application is running. For instructions, see “How to clone a virtual machine” on page 54 .
Restore VM disk files	Enables you to restore an application by restoring the disk files of the virtual machine on which it is running to an NFS or SMB shared location. Select this option if you want to use these virtual machine disk

Restore option	Description
	files later to create a new virtual machine. For instructions, see “How to restore virtual machine disk files” on page 55.

^aNot supported for Nutanix ESXi clusters.

Accessing the Applications panel

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

How to restore a virtual machine

⚠ 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 45.](#)

Limitation

Restoring a virtual machine running on a Nutanix ESXi cluster is not supported.

To restore a virtual machine to the original location or a new location with original settings, follow these steps:

1. In the Applications panel, click the application that you want to restore to open the Details section. The Details section appears only if you click an application. Selecting the check box before the name of the application will not open the Details section.
2. In the Details section that appears at the bottom of the screen, select the desired restore point, and then click **🔄 Restore**.

⚠ 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. From the drop-down menu, select where you want to restore the virtual machine.

📄 Note By default, the original location is selected. If you decide to restore the virtual machine to another container, it will not be restored from the snapshot, but from the target. Therefore, no fast restore will be performed.

6. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore. The original virtual machine is deleted automatically.
7. Click **Restore**.

📄 Note *Nutanix ESXi clusters only.* Because the minimum RAM required for restoring a virtual machine is 256 MB, any virtual machine with less RAM is automatically set to 256

MB 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 the Nutanix Guest Tools (NGT) software bundle to ensure the future successful backups of application data.


How to clone a virtual machine

To create a clone of the original virtual machine, restore it to a new location with custom settings. To do so, follow these steps:

1. In the Applications panel, click the application that you want to restore to open the Details section. The Details section appears only if you click an application. Selecting the check box before the name of the application will not open the Details section.
2. In the Details section that appears at the bottom of the screen, select the desired restore point, and then click **Restore**.

 **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. From the drop-down menu, select where you want to restore the virtual machine.

 **Note** By default, the original location is selected. If you decide to restore the virtual machine to another container, it will not be restored from the snapshot, but from the target. Therefore, no fast restore will be performed.

6. Specify a new name for the virtual machine.
7. 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.
8. Click **Restore**.

During the restore, the original application instance is offline and not accessible.

After cloning a virtual machine

If you restored a virtual machine to a new location on another Nutanix cluster, make sure to add a network adapter to the virtual machine afterward. For details on how to do this, see Nutanix documentation.

How to restore virtual machine disk files

You can restore virtual machine disk files to a shared location. Access to the shared location can be limited to specific accounts in your organization that can later use these disk files to create a new virtual machine. By doing so, you provide greater flexibility when restoring data and optimize restore data management. To restore virtual machine disk files, follow these steps:

1. In the Applications panel, click the application that you want to restore to open the Details section. The Details section appears only if you click an application. Selecting the check box before the name of the application will not open the Details section.
2. In the Details section that appears at the bottom of the screen, select the desired restore point, and then click **Restore**.

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 disk files**, and then click **Next**.

Important During the restore of virtual machine disk files, you cannot perform additional restores or expire backups for this virtual machine.

5. From the drop-down menu, select where you want to restore the disk files, and then provide the required information:
 - **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).
 - **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).
6. Click **Restore**.

During the restore, the original application instance is offline and not accessible.

After restoring virtual machine disk files

After the restore of the virtual machine disk files is complete, you can upload them to the Nutanix image service (for Nutanix AHV clusters) or to the VMware datastore (for Nutanix ESXi clusters) and create a new virtual machine by using them. You do not need HYCU to create a virtual machine from the restored disk files.

Note Keep in mind that due to features such as data resiliency and provisioning, the disk files may require more space than on the shared location. Therefore, make sure to reserve enough space in the Nutanix container or the VMware datastore. For

details on how to do this, see Nutanix or VMware documentation.

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.

⚠ Important *Nutanix ESXi clusters only.* Because the format in which the data is restored is not a native VMware format, you must convert the disk image file to the VMDK format before uploading it to the VMware datastore. To do so, follow these steps:

1. Navigate to the folder where the disk image file is located.
2. As the root user, run the following command:

```
qemu-img convert -f raw -O vmdk <source_disk_image> <target_disk_image>.vmdk
```

Restoring application items

HYCU enables you to restore SQL Server, Exchange Server, and Oracle application items. Depending on which application items you want to restore, see one of the following sections:

- [“How to restore SQL Server databases” below](#)
- [“How to restore Exchange Server databases, mailboxes, and public folders” on page 58](#)
- [“How to restore Oracle database instances and tablespaces” on page 60](#)

⚠ Important If a virtual machine is deleted from the Nutanix cluster, 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.

How to restore SQL Server databases


With HYCU, you can restore SQL Server databases to the original or a different SQL Server instance.

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 higher version of the application.


- Databases that are part of an AlwaysOn Availability Group can be restored only to a primary node (from a secondary or primary node) and cannot be restored to a point in time.

Accessing the Applications panel


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


To restore SQL Server databases, follow these steps:


1. In the Applications panel, click the application whose databases you want to restore to open the Details section.

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


2. In the Details section that appears at the bottom of the screen, select the desired restore point.

 **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.
4. Select **Restore databases**, and then click **Next**.
5. From the Target instance drop-down menu, select where you want to restore the databases.

 **Note** 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. However, if you are restoring the databases that are part of the AlwaysOn Availability Group to a different SQL Server instance, they will not be renamed.

6. 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.
7. *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.

 **Important** For a successful point-in-time restore, make sure that the database recovery model is set to full or bulk-logged.

8. Click **Restore**.
9. *SQL Server 2012 and 2014 AlwaysOn Availability Groups.* Join the restored databases to an AlwaysOn 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 AlwaysOn Availability Group, it is recommended to perform a new backup of your AlwaysOn Availability Group.

How to restore 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, 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.


Prerequisite

Only when 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.

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.
- Restoring multiple databases, mailboxes, and/or public folders at the same time is not supported.

Accessing the Applications panel


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


To restore Exchange Server application items, follow these steps:

1. In the Applications panel, click the application whose application items you want to restore to open the Details section.





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

2. In the Details section that appears at the bottom of the screen, select the desired restore point.

 **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 items.

3. Click  **Restore**. The Restore MS Exchange Server dialog box appears.
4. Select which application items you want to restore, and then click **Next**:

Restore databases	To restore the databases, do the following: <ol style="list-style-type: none"> a. From the Target server drop-down menu, select where you want to restore data.
--------------------------	--

	<p>You can select a restore location 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.</p> <p> Important <i>Only if your mailbox server is a member of a DAG.</i> Make sure to select the target server on which the databases are currently active.</p> <p>b. From the list of databases that are available for a restore, select the database that you want to restore.</p> <p> Important <i>Only if you plan to restore multiple databases.</i> Make sure that you restore databases one by one.</p> <p>c. 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.</p> <p>d. Click Restore.</p> <p>After you restore one database, you can continue with restoring another one, as long as you do it one by one.</p>
<p>Restore mailboxes and/or public folders</p>	<p>To restore the mailboxes and/or public folders, do the following:</p> <p>a. In the Search field, enter the name of a mailbox and/or public folder that you want to restore, or, from the list of mailboxes and/or public folders that are available for a restore, select the one that you want to restore, and then click Next.</p> <p> Important <i>Only if you plan to restore multiple mailboxes and/or public folders.</i> Make sure that you restore mailboxes and/or public folders one by one.</p> <p> 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.</p> <p>b. Select where you want to restore data:</p> <ul style="list-style-type: none"> • Original mailbox • Alternate mailbox, and then enter an alternate mailbox name.

	<p>c. Select the mode for restoring data:</p> <ul style="list-style-type: none"> • Restore in place Enables you to restore data to the original location. • Restore to subfolder <i>(not supported for public folders)</i> Enables you to restore data to the hycu subfolder that is created automatically. <p>d. <i>Only if restoring data to the original location.</i> 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.</p> <p>e. Enter a temporary recovery database path. The default one is C:\ProgramData\Hycu.</p> <p>f. Click Restore.</p>
--	--

How to restore Oracle database instances and tablespaces

With HYCU, you can restore the whole Oracle database instance or the selected tablespaces to the original location.

Limitation

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

Restore 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 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 to the latest backup state, you must select a backup that was performed before the specified point in time so that the database instance can be brought to the last backup state by applying the archive log files from the last 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.

Accessing the Applications panel


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


To restore an Oracle database instance or tablespaces, follow these steps:

1. In the Applications panel, click the application whose database you want to restore to open the Details section.

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

2. In the Details section that appears at the bottom of the screen, select the desired restore point.

 **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.
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. *Whole database instance restore only*. 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**.

Chapter 6

Protecting the HYCU backup controller

It is crucial for your HYCU backup environment that the strategy you choose to protect your HYCU backup controller is highly reliable and ensures security and recoverability.

HYCU provides an internal backup as a disaster protection strategy for your HYCU backup controller virtual machine. If a disaster with the HYCU backup controller occurs (for example, if it is deleted by accident or if the Nutanix cluster on which it is running goes down), use this strategy to successfully protect and recover the HYCU backup controller.

 **Note** HYCU uses synthetic full backups for backing up the HYCU backup controller. This means that each backup represents a consolidation of the full backup and a number of incremental backups. After a new backup is created, all old backups are marked as expired.

To improve protection of your HYCU backup controller and increase its invulnerability, it is highly recommended to combine the internal backup strategy with the native Nutanix data protection strategy (for example, by including the HYCU backup controller to the Nutanix protection domain). For detailed instructions on how to implement data protection for virtual machines in Nutanix, see Nutanix documentation.

Limitation

For an internal backup, you can use only NFS, SMB, and iSCSI targets that are not encrypted.

Backing up the HYCU backup controller

After you deploy the HYCU virtual appliance, perform an internal backup of your HYCU backup controller by assigning a backup policy to it.



Internal backup considerations

- Make sure that your backup policy has only one backup target selected (NFS, SMB, or iSCSI) and that it does not have the Archiving and Copy options enabled. Otherwise, assigning the backup policy to the HYCU backup controller will not be possible. For detailed information about backup policies, see [“Defining your backup policy strategy”](#)

on page 27.

- If you use an iSCSI backup target for the internal backup, take a note which one you set up. You will need this information for recovery purposes.
- It is recommended that the RPO value in the backup policy that you intend to assign to the HYCU backup controller, is lower than any RPO already set for other virtual machines in the HYCU backup environment.

To assign a backup policy, perform these steps:


1. In the navigation pane, click  **Virtual Machines**.
2. From the list of virtual machines, select your HYCU backup controller, and then click  **Policies**. The Policies dialog box appears.
3. Select the backup policy for your HYCU backup controller.
4. Click **Assign** to assign the backup policy to the HYCU backup controller.


If you change the backup policy in any of the following ways after assigning it to the HYCU backup controller, keep in mind the following:


- If you add multiple backup targets to the policy (or automatic selection is enabled), a new full backup will be performed every time the target is changed.
- If you add one or more backup copy targets, every backup copy size will be equal to the size of the full backup.

Location of backup files

- Backup files are located on the backup target (as specified in the backup policy) in the following folder:
 - On SMB and NFS targets: `/bkcptrl/<controllerUuid>`
 - On an iSCSI target: `/mnt/i_scsi-<iscsi_ID>/bkcptrl/<controllerUuid>`

 **Note** You can use the console or a file transfer client (for example, WinSCP) to browse the target.

In this instance, `<controllerUuid>` is the HYCU backup controller UUID, which you can see in the Licensing dialog box, in the Controller string. To access the Licensing dialog box, click  **Administration**, and then select **Licensing**.

 **Note** Take a note of the target and the backup location.

- Every time the target is changed, the information in the Events panel is updated and shows a message about the new location of the backup files.


Recovering the HYCU backup controller

Prerequisite

iSCSI target only. Make sure to establish a connection to the iSCSI target you used for backing up the HYCU backup controller.

To recover the HYCU backup controller, follow these steps:

1. *iSCSI target only*. Perform the following steps to retrieve the disk image files from the backup target:
 - a. Log on to the Nutanix Prism web console (for Nutanix AHV clusters) or the VMware vSphere Web Client (for Nutanix ESXi clusters).
 - b. Deploy a temporary HYCU backup controller virtual machine that will be used for retrieving files from the target. For details on how to do this, see one of the following sections:
 - [“Deploying HYCU on a Nutanix AHV cluster” on page 12](#)
 - [“Deploying HYCU to a Nutanix ESXi cluster” on page 16](#)
 - c. Log on to HYCU.
 - d. Set up the same iSCSI target as you used for the HYCU backup controller you want to recover. For details, see [“How to set up an iSCSI target” on page 26](#).
2. Retrieve two HYCU backup controller disk image files from the backup target and save them to a temporary location. Access the HYCU logical volume that is located on the backup target (you can identify the HYCU logical volume by its `hycu_lv` prefix). The location of the disk image files is as follows:
 - On SMB and NFS targets: `/bkcptrl/<controllerUuid>`
 - On iSCSI target: `/mnt/i_scsi-<iscsi_ID>/bkcptrl/<controllerUuid>`

 **Note** You can use the console or a file transfer client (for example, WinSCP) to browse the target.

The `bkcptrl` folder contains the folders of the both newly created HYCU backup controller and the HYCU backup controller you want to recover. To determine in which folder the required disk image files reside, check the HYCU backup controller ID or search both folders—the required disk image file is the one whose size is approximately 32 GB.

To locate the disk image files in the `/bkcptrl/<controllerUuid>`, consider their size. The size of the HYCU backup controller data disk image file is equal to or greater than 32 GB, whereas the size of the HYCU backup controller OS disk image file is approximately 10 GB. Besides the HYCU backup controller disk image files, this folder may also contain subfolders with other virtual machine backups. Ignore these subfolders when recovering the HYCU backup controller.

3. *Nutanix ESXi clusters only*. Convert the disk image files from RAW format to VMDK format:
 - a. Navigate to the folder where the disk image file is located.
 - b. As the root user, run the following command:

```
qemu-img convert -f raw -O vmdk <source_disk_image> <target_disk_image>.vmdk
```


4. Upload the disk image files to the image service by using the Prism web console (for Nutanix AHV clusters) or to the datastore by using the VMware vSphere Web Client (for Nutanix ESXi clusters).
5. After you upload the disk image files, depending on whether the HYCU backup controller is present in the list of virtual machines, select one of the following procedures:
 - HYCU backup controller is present:

Shut the HYCU backup controller down, remove the old disks from it, and then add new disks based on the disk image files you uploaded. After the new disks are added, make sure to turn the HYCU backup controller on.

For details, see Nutanix or VMware documentation.
 - HYCU backup controller is not present:
 - a. Log on to the Nutanix Prism web console (for Nutanix AHV clusters) or the VMware vSphere Web Client (for Nutanix ESXi clusters).
 - b. Create a new virtual machine. For details on how to do this, see one of the following sections:
 - [“Creating a virtual machine for HYCU deployment on a Nutanix AHV cluster” on page 14](#)
 - [“Deploying HYCU to a Nutanix ESXi cluster” on page 16](#)
 - c. Click **Power on** to turn on the HYCU backup controller.
 - d. Configure the network settings for the new virtual machine because a new network card is assigned to it.

iSCSI target only. After you recovered the HYCU backup controller, you can delete the temporary HYCU backup controller virtual machine you created for retrieving the disk image files. For details, see Nutanix or VMware documentation.

Chapter 7

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 HYCU environment.	“Using the HYCU dashboard” below
Track tasks that are running in my environment and get an insight into the specific task status.	“Checking the status of jobs” on page 68
View all events that occurred in my environment.	“Viewing events” on page 69
View the backup status of backup sources.	“Viewing backup source details” on page 69
Narrow down the list of displayed elements in panels.	“Filtering data in panels” on page 71
View backup target information, activate or deactivate a backup target, increase the size of an iSCSI backup target, or edit or delete a backup target.	“Managing backup targets” on page 75
View backup policy information, or edit or delete a backup policy.	“Managing backup policies” on page 78
Mark restore points as expired.	“Expiring backups manually” on page 80

In case of the 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 81](#).

Using the HYCU dashboard

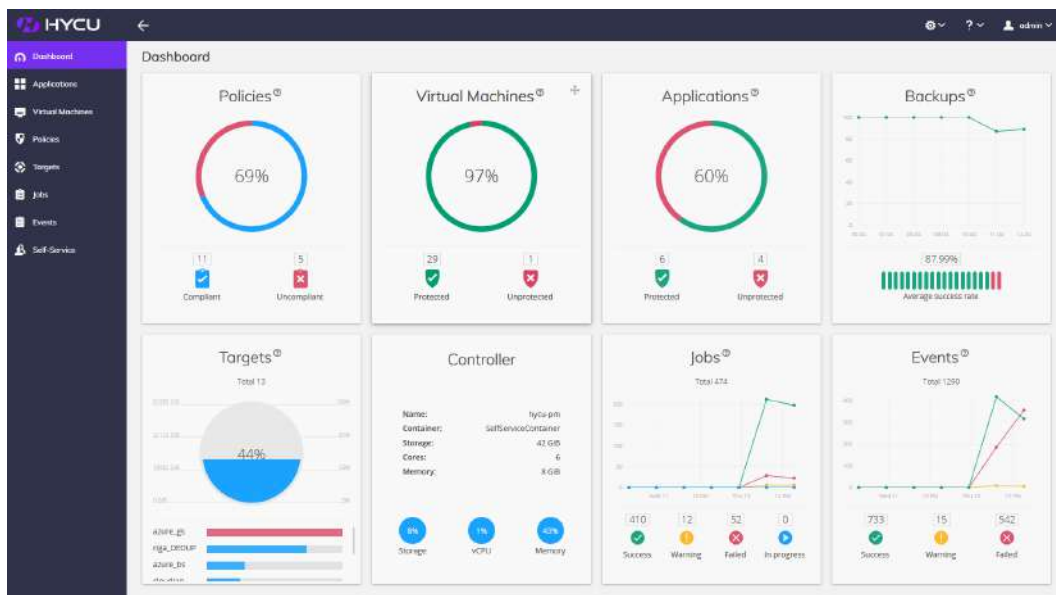
The HYCU dashboard provides you with an at-a-glance overview of the data protection status in your environment. This intuitive dashboard enables you to monitor all data

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 permissions define 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
Policies	Shows the percentage of policies that are compliant, and the exact number of compliant and incompliant policies. A policy is considered compliant if all virtual machines and applications within this policy are compliant with the policy settings. For detailed information about policies, see “Defining your backup policy strategy” on page 27 .
Virtual Machines	Shows the percentage of protected virtual machines in your environment, and the exact number of protected and unprotected virtual machines. A virtual machine is considered protected if it has an assigned policy and if there is at least one valid backup within the defined retention period. For detailed information about protecting virtual machines, see “Backing up virtual machines” on page 42 .
Applications	Shows the percentage of protected applications, and the exact number of protected and unprotected applications. An application


Dashboard widget	Description
	is considered protected if it has an assigned policy and if there is at least one valid backup within the defined retention period. For detailed information about protecting applications, see “Backing up applications” on page 43 .
Backups	Shows the backup job success rate for the last seven days.
Targets*	Shows the number of existing backup targets, overall capacity utilization, and the utilization per target type. For detailed information about setting up backup targets, see “Setting up backup targets” on page 21 .
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 81 .
Jobs	Shows the number of jobs in the HYCU 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 HYCU environment in the last 48 hours. It also shows the number of the events according to their severity. For details, see “Viewing events” on the next page .

* An 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 the status of processes that are currently running.
- Check the status of completed and stopped processes.
- Check more details about a specific job in the Details section that appears at the bottom of the screen after you select the job.

 **Tip** To minimize the Details section, click  **Minimize** or press **Spacebar**. To

- return it to its original size, click ▲ **Maximize** or press **Spacebar**.
- 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**.

The following table shows the job information:

Job information	Description
Name	Name of the job that was performed (for example, adding a cluster, 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:

Level	Severity level of the event (Info, Warning, or Error)
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 Details section where you can find the event summary and more details about the event, click the desired event.


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

Viewing backup source details

You can view the details about each virtual machine and discovered application in the Details section of the Virtual Machines or Applications panel. The following details are

available:

Summary	Shows detailed information about the selected backup source.
Restore point	<p>You can view the following information about each restore point:</p> <ul style="list-style-type: none"> • Date and time when the restore point was created. • Full or Incremental: Shows the type of backup. • Snapshot: Visible only if the Nutanix cluster contains a local snapshot that enables you to perform a fast restore. • Copy: Visible only if the Copy policy option is enabled. • Archive: Visible only if the Archiving policy option is enabled.
Compliance	<p>Shows the compliancy status of the backup source (Success, Failure, or Undefined).</p> <p>By pausing on a compliancy status indicated by a respective icon, additional information about the backup is available. You can see backup frequency, when the last successful backup was performed, the time limit you set for a restore, the estimated time required for a restore, and the number of copies.</p>
Backup status	For details, see “Viewing the backup status of backup sources” below.
Restore status	<p>Shows a progress bar indicating the progress of the backup source restore.</p> <p> Note If you double-click a progress bar, you are directed to the Jobs panel where you can check details about the related job.</p>



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

Viewing the backup status of backup sources

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

Limitation

The Success with errors backup status for virtual machines with attached volume groups is available only for a Nutanix AHV cluster.

Backup status of the backup source	Restore a VM?	Restore VM files?	Restore an APP?
 (Success)	✓	✓	✓
 (Success with warnings)	✓	✓	✓ ^a


Backup status of the backup source	Restore a VM?	Restore VM files?	Restore an APP?
○ (Success with errors)	✓ ^b	✓ ^c	✓ ^d
✘ (Error)	✘	✘	✘

^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).

 **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 backup 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 job.

Filtering data in panels

HYCU enables you to filter data in the panels so you can easily find what you need. Each panel contains different filtering options and displays only the information that meets the specified filter criteria (for example, filtering the data in the Virtual Machines panel helps you to focus only on the virtual machines that you are interested in or responsible for).

To filter data in panels, follow these steps:

1. Click  **Filters**. The Filters side panel opens.
2. Select your filter criteria.
3. Click **Apply Filters**.

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

- [“Filtering options in the Applications panel” on the next page](#)
- [“Filtering options in the Virtual Machines panel” on the next page](#)
- [“Filtering options in the Policies panel” on page 73](#)

- [“Filtering options in the Targets panel” on page 74](#)
- [“Filtering options in the Jobs panel” on page 74](#)
- [“Filtering options in the Events panel” on page 75](#)
- [“Filtering options in the Self-Service panel” on page 75](#)

Filtering options in the Applications panel

In the Filters 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.
Clusters	From the drop-down menu, select the clusters that host the virtual machines on which the applications are running.
Policies	From the drop-down menu, select the backup policies that are assigned to the virtual machines on which the applications are running.
Owners	From the drop-down menu, select the owners that are assigned to the virtual machines on which the applications are running.
Application types	From the drop-down menu, select the application types.
Compliance	Select one or more check boxes to filter by the compliance status: <ul style="list-style-type: none"> • Success: Discovered applications are compliant. • Failure: Discovered applications are not compliant. • Undefined: The compliance status cannot be retrieved.
Protection	Select one or more check boxes to filter by the protection status: <ul style="list-style-type: none"> • Yes • No • Deleted
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 machine is offline or not reachable.

Filtering options in the Virtual Machines panel

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

Filtering option	Action
Search	Enter a search term. You can filter by the virtual machine name, the HYCU UUID, or the Nutanix UUID.
Clusters	From the drop-down menu, select the clusters that host the virtual machines.
Credential groups	From the drop-down menu, select the credentials for the virtual machines.
Policies	From the drop-down menu, select the backup policies that are assigned to the virtual machines.
Owners	From the drop-down menu, select the owners that are assigned to the virtual machines.
Compliance	Select one or more check boxes to filter by the compliance status: <ul style="list-style-type: none"> • Success: Virtual machines are compliant. • Failure: Virtual machines are not compliant. • Undefined: The compliance status cannot be retrieved.
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 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: <ul style="list-style-type: none"> • Yes • No • Deleted

Filtering options in the Policies panel

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

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

Filtering option	Action
	<ul style="list-style-type: none"> • Success • Failure • Undefined

Filtering options in the Targets panel

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

Filtering option	Action
Search	Enter a search term. You can filter by the name of the backup target.
Target type	Select one or more check boxes to filter by the target type: <ul style="list-style-type: none"> • AWS S3/Compatible • Azure • NFS • SMB • iSCSI
Status	Select one or more check boxes to filter by the status of the target: <ul style="list-style-type: none"> • Ok • Warning • Error • Undefined

Filtering options in the Jobs panel

In the Filters 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: <ul style="list-style-type: none"> • Ok • Warning • Failed • Queued • Executing

Filtering option	Action
	<ul style="list-style-type: none"> • Aborted

Filtering options in the Events panel

In the Filters 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.
Username	From the drop-down menu, select the user name.
Category	Enter a text string to filter the list to include only the categories with the specified string.
Severity	Select one or more check boxes to filter by the severity of the event: <ul style="list-style-type: none"> • Success • Warning • Failed

Filtering options in the Self-Service panel

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


Filtering option	Action
User group name	Enter the user group name.
Protection service	Select one of the following to filter by the protection service: <ul style="list-style-type: none"> • Tenant • Administrator
Status	Select one of the following to filter by the status of the user group or user (that is, which user groups or users are allowed to log on to HYCU and which are not): <ul style="list-style-type: none"> • Active • Inactive

Managing backup targets

If you have the proper permissions, you can view backup target information, edit backup target properties, activate or deactivate a backup target, or delete a backup 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 backup target information

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

Backup target information	Description
Name	Name of the backup target.
Type	Type of backup target (AWS S3/Compatible, Azure, SMB, NFS, or iSCSI).
Health	<p>Health status of the backup target:</p> <ul style="list-style-type: none"> • Grey: Shows the initial backup target status before a health test. It also indicates an inactive backup target. • Green: The backup target is in a healthy state with backup target utilization of less than the configured value (by default, 90%). • Yellow: Backup target utilization is over the configured value (by default, 90%). • Red: Backup target utilization is over the configured value (by default, 95%). It also indicates a backup target error state after a test task (for example, an I/O error occurred, the backup target is not accessible, the permission is denied, and so on). <p>When HYCU checks if there is enough space on the backup target to perform a backup, it calculates it by using the total provisioned space of all disks on the virtual machine, regardless of whether the backup is full or incremental.</p>
Size	Estimation of the amount of storage space that should be reserved for the backup files (in MB, GB, or TB).
Utilization	Percentage of the specified backup target size that is already used for storing protected data.
Mode	<p>Mode of the backup target:</p> <ul style="list-style-type: none"> • Read/Write: You can use this backup target for backing up and restoring data. • Read Only: You can use this backup target only for restoring data.


Backup target information	Description
Status	<p>Status of the backup target:</p> <ul style="list-style-type: none"> • Active: You can use this backup target for backing up and restoring data. • Inactive: You cannot use this backup target for backing up and restoring data. This status indicates that the backup target is deactivated due to maintenance tasks (for example, adding new disks). <p>For details on how to change the status of the backup target, see “Activating or deactivating a backup target” below.</p>


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

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

Editing a backup target

To edit a backup target, follow these steps:



1. In the Targets panel, select the backup target that you want to edit, and then click  **Edit**. The Edit Target dialog box appears.
2. Edit the selected backup target as required. For detailed information about backup target properties, see [“Setting up backup targets”](#) on page 21.

 **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.

3. Click **Save**.

Activating or deactivating a backup target

To change the status of a backup target (that is, to activate or deactivate it), follow these steps:

1. In the Targets panel, select the backup target that you want to activate or deactivate.
2. Change the status of the selected backup target by clicking  **Activate** or  **Deactivate**.
3. *Only if deactivating the backup target.* Click **Yes** to confirm that you want to deactivate the selected backup target.

If you deactivate a backup target, this target will not be used for backup and restore operations anymore.


Increasing the size of an iSCSI backup target

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

Prerequisites

- The size of the backup target has been increased on the iSCSI server.
- No backup or restore job is in progress on the selected backup target.
- No other maintenance task is already running on the selected backup target (such as editing the backup target and updating the iSCSI Initiator secret or resetting mutual CHAP authentication sessions for the backup targets with CHAP authentication enabled).
- No other size increase of the selected backup target has already been started.

To increase the size of an iSCSI backup target, follow these steps:


1. In the Targets panel, select the backup target whose size you want to increase.
2. In the Details section that appears at the bottom of the screen, click  **Extend**.
3. Click **Yes** to confirm that you want to increase the size of the selected backup target.


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

Deleting a backup target

You can delete a backup target if it does not contain protected data. After deleting a backup target, no backup or restore actions including this backup target are possible anymore.

To delete a backup target, follow these steps:

1. In the Targets panel, select the backup 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 backup target.

Managing backup policies

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

Accessing the Policies panel

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

Viewing backup policy information

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


Backup policy information	Description
Name	Name of the backup policy.
Compliance	<p>Compliance status of the backup policy:</p> <ul style="list-style-type: none"> • Success • Failure • Undefined <p>A policy is considered compliant if all virtual machines and applications within this policy are compliant with the policy settings.</p>
VM Count	Total number of virtual machines that have the particular backup policy assigned to them.
APP Count	Total number of applications that have the particular backup policy assigned to them.
Description	Description of the backup policy (how often backup and restore jobs are performed).

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

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

Editing a backup policy

To edit a backup policy, follow these steps:

1. In the Policies panel, select the backup policy that you want to edit, and then click  **Edit**. The Edit Policy dialog box appears.
2. Edit the selected backup policy as required. For detailed information about backup policy properties, see [“How to create a custom backup policy” on page 30](#).
3. Click **Save**.

When editing a backup policy that is assigned to several virtual machines, one of which is the HYCU backup controller, make sure that the backup policy remains applicable for an

internal backup. For details on protecting the HYCU backup controller, see [“Protecting the HYCU backup controller” on page 62](#).

Deleting a backup policy

To delete a backup policy, follow these steps:

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

Expiring backups manually

If there is a restore point that you do not want to use for data restore anymore, you can mark it as expired. If the most recent restore point is marked as expired, the next backup will be a full backup.

Depending on which backup sources you want to expire, 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**.

To mark a restore point as expired, follow these steps:


1. In the Virtual Machines or Applications panel, select the backup source for which you want the old backups to be marked as expired.
2. In the Details section that appears at the bottom of the screen, select the restore point that you want to mark as expired.
3. Click **Expire**. Keep in mind that an expire action cannot be undone.
4. Click **Yes** to confirm that you want the selected restore point to be marked as expired.

⚠ Important Any subsequent incremental backups will also be marked as expired unless the status of the selected restore point is Error. In this case, only the selected restore point is expired and not the whole backup chain.

After you mark a recovery point as expired, the HYCU cleaning process removes expired backups from the backup target.

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 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 8

Managing HYCU users

The HYCU user management system provides security mechanisms that help prevent unauthorized users from accessing backup data. Only HYCU users have access to the backup environment and can benefit from all the features that HYCU offers.

Each HYCU user belongs to one of the two user groups with different protection services:


Protection service	Description
Administrator	<p>A user group with an administrator protection service (an administrator user group) is created by default and includes the following types of users:</p> <ul style="list-style-type: none">• Built-in administrator () Created during the deployment of the HYCU virtual appliance and cannot be edited or deleted.• Administrator Can be created, edited, and deleted by the built-in administrator or any other administrator. <p>Users belonging to the administrator user group have permissions to:</p> <ul style="list-style-type: none">• Perform all configuration tasks in HYCU and all the tasks related to the backup and restore of virtual machines that are not owned by any tenant user group.• Create new user groups and users, edit or delete the existing ones, set ownership for virtual machines, and enable or disable access to HYCU. <p>An administrator can monitor virtual machines owned by tenant user groups, but does not have access to backup resources.</p>
Tenant	<p>A user group with a tenant protection service (a tenant user group) represents a specific customer or department. An administrator assigns selected virtual machines to a tenant user group that gets access to backup resources and is authorized to perform backup and restore tasks. Can be created, edited, and deleted by an administrator.</p>

Protection service	Description
	 Important If a specific tenant user group is deleted, all data that is backed up by this user group is deleted from the database.

Setting up user groups and users


To give users specific rights to access only certain data protection areas within the HYCU environment, complete these tasks:

1. Create a new user group. For details, see [“Creating a new user group” below](#).


 **Note** The administrator user group is created by default. If you plan to add an administrator user, you can skip this task.

2. Add a new user to an already existing user group. For details, see [“Adding a new user” below](#).

Accessing the Self-Service panel


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



After you set up user groups and users, continue with setting ownership of virtual machines. For details, see [“Setting ownership of virtual machines” on the next page](#).

 **Note** Depending on the nature of your business, you can at any time enable or disable specific user groups or users from logging on to HYCU. For details, see [“Activating or deactivating a user group or a user” on the next page](#).

Creating a new user group

To create a new user group, follow these steps:

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

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

Adding a new user

Prerequisite

Only if you plan to use an Active Directory for authentication. Active Directory authentication is configured. For details on how to do this, see [“Configuring Active Directory authentication” on page 99](#).

To add a new user, follow these steps:

1. In the Self-Service panel, click a user group to which you want to add a new user. If you want to add a user to a new user group, first create the user group as described in [“Creating a new user group” on the previous page](#).
2. In the section that appears at the bottom of the panel and shows details about the selected user group, click **+ Add User**. The New Account dialog box opens.
3. Enter the user name.
4. Select one of the following user authentication types:
 - **HYCU**, and then enter the user password and, optionally, email address.

Note The minimum password length is six characters.

 - **AD**, and then, from the Active Directory drop-down menu, select the Active Directory the user account belongs to.
5. Click **Save**.

You can also edit any of the existing users (click **Edit User** and make the required modifications) or delete them (click **Remove User**). Keep in mind that the built-in administrator cannot be edited or deleted.

Activating or deactivating a user group or a user

To activate or deactivate a user group or a user (that is, enable or disable specific user groups or users from logging on to HYCU), follow these steps:

1. In the Self-Service panel, check the status of the user group or user that you want to activate or deactivate.
2. Do one of the following:
 - If the status of the selected user group or user is active and you want to deactivate it, click **Deactivate**.
 - If the status of the selected user group or user is inactive and you want to activate it, click **Activate**.

Setting ownership of virtual machines

As an administrator, you can set ownership of virtual machines, enabling specific user groups to access specific virtual machines and backup resources in the HYCU environment, and to perform backup and restore tasks.


Accessing the Virtual Machines panel

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

Assigning owners to virtual machines

To assign owners to virtual machines, follow these steps:


1. In the Virtual Machines panel, select virtual machines to which you want to assign owners, and then click **Owner**. The Accounts dialog box opens.
2. Select which user group you want to assign as an 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 user group to the relevant virtual machine.

Removing owners from virtual machines

To remove owners from virtual machines, follow these steps:

1. In the Virtual Machines panel, select virtual machines from which you want to remove owners, and then click **Owner**. The Accounts dialog box opens.
2. Select which user group you want to remove as an owner of the selected virtual machines, and then click **Unassign**.

 **Important** If you change an owner of virtual machines, all restore points for these virtual machines are deleted.

Chapter 9

Administering

After you deploy HYCU, you can perform various tasks to administer and customize HYCU for your data protection environment.

I want to...	Procedure
Obtain a permanent HYCU license.	"Licensing" on the next page
Upgrade HYCU to a new available version.	"Upgrading HYCU" on page 90
Change network settings.	"Changing network settings" on page 94
Change the HYCU listening port number.	"Changing the HYCU listening port number" on page 95
Configure the SSL certificate.	"Configuring the SSL certificate" on page 95
Configure FIPS-compliant mode for HYCU.	"Configuring FIPS-compliant mode for HYCU" on page 96
Set the iSCSI Initiator secret.	"Setting the iSCSI Initiator secret" on page 98
Configure encryption for backup targets.	"Configuring backup target encryption" on page 98
Set power options.	"Setting power options" on page 98
Configure Active Directory authentication.	"Configuring Active Directory authentication" on page 99
Configure log file settings to troubleshoot problems if HYCU does not perform as expected.	"Setting up logging" on page 99
Access the HYCU backup controller virtual machine by using SSH.	"Accessing the HYCU backup controller virtual machine by using SSH" on page 100
Use hyCLI.	"Using the command-line interface" on page 102
Use the HYCU REST API to automate tasks.	"Using the HYCU REST API Explorer" on

I want to...	Procedure
	page 102
Enable HTTPS for WinRM connections.	“Enabling HTTPS for WinRM connections” on page 102
Increase the size of the HYCU virtual disk.	“Increasing the size of the HYCU virtual disk” on page 103

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 105](#).

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 that best suits your environment. The following license types are available:

- Socket-based licenses

Licenses are based on the number of CPU sockets on all Nutanix clusters. You should determine the total number of CPU sockets on all the Nutanix clusters that you plan to protect by using HYCU, so that you purchase the required number of licenses.

- VM-based licenses

Licenses are based on the number of protected virtual machines on all Nutanix clusters. You should determine the total number of virtual machines on all the Nutanix clusters that you plan to protect by using HYCU, so that you purchase the required number of licenses.

Nutanix environment considerations

- *Remote office/branch office (ROBO) environment only.* The number of sockets and virtual machines residing in the ROBO environment is calculated and added to the total number of HYCU licenses.
- *Nutanix Community Edition (CE) environment only.* No HYCU licenses are required.

Perform the following tasks:

1. Buy a needed number of HYCU licenses. To discuss the options, contact your Sales representative.
2. Create a license request. For details, see [“Creating a license request” on the next page](#).
3. Request and obtain licenses from the web licensing portal. For details, see [“Requesting and retrieving licenses” on 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 Nutanix clusters that you want to protect to the HYCU environment. For instructions, see [“Adding Nutanix clusters” on page 20](#).

To create a license request, follow these steps:

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
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.

To do so:

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 (Socket based or VM based licenses) 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
EXP 02.08.2017
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. Activate the licenses in HYCU as follows:

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 information related to licensing at any time. The following information is displayed in the Licensing dialog box:

- License type
- Backup controller ID
- Status
- Actual number of sockets
- Licensed number of sockets

- Actual number of protected virtual machines
- Licensed number of protected virtual machines

Upgrading HYCU

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

Because the upgrade process aborts all currently running jobs, make sure that the jobs you do not want to be aborted are finished and targets are deactivated before you start the upgrade.

Besides a new version of the HYCU virtual appliance, which you will use for upgrading your current product version, a HYCU patch is provided to prepare your current version for an upgrade to version 3.0.0.

Depending on which hypervisor is running on your Nutanix cluster, follow the instructions in one of the following sections:

- [“Upgrading HYCU on a Nutanix AHV cluster” below](#)
- [“Upgrading HYCU on a Nutanix ESXi cluster” on page 92](#)


Upgrading HYCU on a Nutanix AHV cluster


Prerequisite

Create a snapshot of the HYCU backup controller. For instructions, see Nutanix documentation.

To upgrade HYCU on a Nutanix AHV cluster, follow these steps:

1. Log on to the Nutanix Prism web console, and then do the following:
 - a. In the menu bar, click **Home**, and then select **VM**.
 - b. Navigate to the HYCU backup controller virtual machine that you want to upgrade, and then double-click it. The Update VM dialog box opens.
 - c. In the Disks section, click **+Add New Disk**. This disk will be used by HYCU as a new data disk.
 - d. In the Add Disk dialog box, leave the default settings. In the Size (GiB) field, enter the disk size that is equal to or greater than the current HYCU disk size.



 **Note** You can later increase the size of the HYCU virtual disk if needed. For details, see [“Increasing the HYCU disk size in a Nutanix AHV cluster” on page 104](#).
 - e. Click **Add**.
 - f. In the Update VM dialog box, click **Save**.
2. Install the HYCU patch. To do so, perform these steps:


- a. Log on to HYCU. For details, see [“Logging on to HYCU” on page 17](#).
 - b. Click  **Administration**, and then select **Power Options**.
 - c. In the Power Options dialog, select **Suspend**.
 - d. Unpack the `hycu-2.0.1-<revision>.zip` patch file to the `/home/hycu` folder on the HYCU backup controller.
 - e. In the Nutanix Prism web console, select the HYCU backup controller virtual machine, and then click **Launch Console**.
 - f. Log on to the HYCU backup controller console. The following are the default credentials:

User name:	hycu
Password:	hycu/4u
 - g. Navigate to the `/home/hycu` folder, and then run the patch installation script:


```
sh /opt/grizzly/bin/HycuPatch.sh
```
 - h. Run the migration script:



```
sh /opt/grizzly/bin/migrate.sh
```


The HYCU backup controller restarts.
3. Upload the HYCU virtual appliance image that you want to use for an upgrade to your Nutanix cluster. To do so, follow these steps:
 - a. In the Nutanix Prism web console, 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 the HYCU image name in the format that should correspond to that of the HYCU image file you are uploading. Optionally, enter an annotation.

 **Important** The HYCU virtual appliance image must be uploaded to the Nutanix cluster in the following format:
`hycu-<version>-<revision>`

For example: `hycu-3.0.0-3634`

Make sure to leave out the `.qcow2` extension when entering the HYCU image name. If you enter the HYCU image name in a different format, you will not be able to use this image for an upgrade.
 - ii. From the Image Type drop-down menu, select **DISK**.
 - iii. From the Storage Container drop-down menu, select a container for the image to be uploaded.
 - iv. In the Image Source section, select one of the following:

- i. **From URL**
Specify the location of the image file by using a URL.
 - ii. **Upload a file**
Specify the location of the image file saved on your file system.
 - d. Click **Save**.
 - e. Click **Close** after the image is successfully uploaded.
4. Log on to HYCU.
 5. Click  **Administration**, and then select **Upgrade Software**.
 6. In the Upgrade Software dialog box, check the current version of HYCU and all available versions.
 7. From the list of the available versions, select the one to which you want to upgrade HYCU, and then click **Upgrade**.

 **Important** Before you first log on to HYCU, make sure to perform a hard reload of the HYCU webpage in your browser.

The old HYCU backup controller virtual machine will remain on the Nutanix AHV cluster and will be renamed to `<HYCU_BC_name>_version_<old_HYCU_version>`. After you make sure HYCU was upgraded successfully, you can safely delete it.


Upgrading HYCU on a Nutanix ESXi cluster

Prerequisites

- Create a snapshot of the HYCU backup controller using the Nutanix protection domain. For instructions, see Nutanix documentation.
- If there are any HYCU snapshots created by using VMware vSphere, make sure to remove them.

To upgrade HYCU on a Nutanix ESXi cluster, do the following:

1. Add a new disk to the HYCU backup controller. To do so, follow these steps:
 - a. Log on to the VMware vSphere Web Client.
 - b. Browse for and select the HYCU backup controller virtual machine.
 - c. Click the **Configure** tab, and then click **Edit**.
 - d. In the Edit Settings dialog box, click the **Virtual Hardware** tab.
 - e. From the New device drop-down menu, select **New Hard Disk**, and then click **Add** followed by **OK**. Make sure the disk size is equal to or greater than the current HYCU disk size.
2. Install the HYCU patch. To do so, perform these steps:


- a. Log on to HYCU. For details, see “Logging on to HYCU” on page 17.
 - b. Click  **Administration**, and then select **Power Options**.
 - c. In the Power Options dialog, select **Suspend**.
 - d. Unpack the `hycu-2.0.1-<revision>.zip` patch file to the `/home/hycu` folder on the HYCU backup controller.
 - e. In the Nutanix Prism web console, select the HYCU backup controller virtual machine, and then click **Launch Console**.
 - f. Log on to the HYCU backup controller console. The following are the default credentials:

User name:	hycu
Password:	hycu/4u
 - g. Navigate to the `/home/hycu` folder, and then run the patch installation script:


```
sh /opt/grizzly/bin/HycuPatch.sh
```
 - h. Run the migration script:



```
sh /opt/grizzly/bin/migrate.sh
```

The HYCU backup controller restarts.
3. In the VMware vSphere Web Client, follow these steps:
- a. Click the **VMs** tab, and then click **Deploy OVF Template...**
 - b. In the Select template context of the Deploy OVF Template dialog box, select **Local files**, and then click **Browse**.
 - c. Navigate to the HYCU virtual appliance package, select the HYCU disk image file (`.vmdk`) and the OVF template (`.ovf`), and then click **Open**.
 - d. In the Select template context, click **Next**.
 - e. In the Select name and location context, select the location where you want to deploy HYCU backup controller, and then click **Next**.
 - f. In the Select a resource context, select the Nutanix cluster where the HYCU backup controller will reside, and then click **Next**.
 - g. In the Review details context, verify the template details, and then click **Next**.
 - h. In the Select storage context, select the same Datastore as you are using for the HYCU backup controller virtual machine, and then click **Next**.
 - i. In the Select Network context, select the desired VM Network, and then click **Next**.
 - j. In the Customize template context, enter the values for the following:
 - *Optional*. Host name for the HYCU 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**.

- In the Ready to complete context, review the settings, and then click **Finish**. Leave the newly created virtual machine turned off.

 **Note** Creating a HYCU virtual machine may take a few moments.

- Log on to HYCU.
- Click  **Administration**, and then select **Upgrade Software**.
- In the Upgrade Software dialog box, check the current version of HYCU and all available virtual machines.
- From the list of the available virtual machines, select the one you created, and then click **Upgrade**.

After the upgrade process completes, the upgraded HYCU virtual machine is powered on automatically.

 **Important** Before you first log on to HYCU, make sure to perform a hard reload of the HYCU webpage in your browser.

The old HYCU backup controller virtual machine will remain on the Nutanix ESXi cluster and will be renamed to `<HYCU_BC_name>_version_<old_HYCU_version>`. After you make sure HYCU was upgraded successfully, you can safely delete it.

Changing network settings

This section describes the steps you must perform if you want to change your network settings.


Accessing the Network dialog box



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

To change network settings, follow these steps:

- In the General tab of the Network dialog box, enter a new host name, the IP address, the netmask associated with the chosen IP address, and the IP address of the gateway

that HYCU should use.


 **Important** Make sure that the IP address is set correctly so that it accurately reflects your network environment.

2. In the DNS Addresses tab, click **Add** to add a new DNS address.
If you want to delete any DNS address, click  **Delete** to the right of it.
3. In the Domains tab, click **Add** to add a new domain.
If you want to delete any existing domain, click  **Delete** to the right of it.
4. Click **Save**.

Changing the HYCU listening port number


This section describes the steps you must perform if you want to change the HYCU listening port number.

Accessing the Management Port dialog box

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

To change the listening port number of the HYCU web user interface, follow these steps:

1. In the Management Port dialog box, change the existing port number to the desired one.


 **Important** Make sure that the port you choose is not used by any other process.

2. Click **Save**.

Configuring the SSL certificate

This section describes the steps you must perform if you want to configure the SSL certificate.


Accessing the SSL Certificate dialog box

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

In the SSL Certificate dialog box that appears, you can view the information about your SSL certificate, such as the certificate holder's name, the certificate's expiry date, and the strength of the certificate keys.

Depending on whether you want to create a new self-signed certificate or import a CA certificate into HYCU, do one of the following:

- [“Creating a new self-signed certificate” on the next page](#)
- [“Importing a CA certificate” on the next page](#)

 **Note** It is recommended that you replace the default self-signed certificate with a CA signed certificate.


Creating a new self-signed certificate

To create a new self-signed certificate, do the following:

1. Click **Regenerate**.
2. Click **Yes** to confirm that you want to regenerate your SSL certificate.


Importing a CA certificate

You can import only CA certificates that are compliant with the PKCS#7 standard and encoded in the DER or PEM format. The HYCU backup controller holds only one custom SSL certificate. When you import a new certificate, the previous certificate is discarded.

 **Note** If the certificate that you want to import uses a wildcard for the Common Name (CN), make sure that the Certificate Subject Alt Name field exists and contains the list of all possible host names or FQDNs, and their corresponding IP addresses. Otherwise, the certificate may be recognized as invalid by your browser or hyCLI.

To import a CA certificate, do the following:

1. Click **Import**. The Import dialog box appears.
2. Browse for the following imported files:
 - Private key: Click **Browse** to select the private key associated with the certificate to be imported.
The private key should be created by using RSA or DSA algorithm and be compliant with the PKCS#1 or PKCS#8 standard.
 - Public certificate: Click **Browse** to select the signed public part of the server certificate corresponding to the private key.
 - CA certificate/chain: This field is optional, if the public certificate contains CA certificate/chain. Click **Browse** to select the certificate or chain of the signing authority for the public certificate.
3. Click **Import**.

 **Important** Any changes that you make to your SSL certificate will result in an automatic logoff.

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 web server:

```
systemctl stop grizzly.service
```

2. Enable FIPS-compliant mode:

```
/opt/grizzly/bin/enable_fips.sh
```

3. Start the HYCU web 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 web server:

```
systemctl stop grizzly.service
```

2. Disable FIPS-compliant mode:

```
/opt/grizzly/bin/enable_fips.sh -d
```

3. Start the HYCU web server:

```
systemctl start grizzly.service
```

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 backup targets” on page 21](#).

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**.

Configuring backup target encryption

If you enabled backup target encryption when setting up your backup target, you can get information about which algorithm is used, view a list of currently encrypted backup targets, and export the encryption key to a file.

Accessing the Encryption dialog box

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

To export the encryption key to a file, in the Encryption dialog box, click **Export**.

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	Pauses all HYCU backup controller activities. 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.
Resume	Allows HYCU backup controller activities to continue.

Configuring Active Directory authentication

In addition to standard HYCU authentication, you can also configure Active Directory authentication. This allows users to log on to HYCU with their Active Directory domain account.


Accessing the Active Directory dialog box

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

To configure Active Directory authentication, add one or more Active Directories as your authentication sources in HYCU. To do so, follow these steps:

1. In the Active Directory dialog box, click **+ New**. The New dialog box appears.
2. In the Name field, enter a name for the Active Directory.
3. In the Domain field, enter the domain name or the domain alias name.


For example, if you enter `mycompany.com` as the domain name or `mc` as the alias domain name, the user will be able to log on to HYCU with `<username>@mycompany.com` or `mc\<username>`.

 **Note** You can enter more than one domain or domain alias name. In this case, make sure to press **Enter** after entering each one.



4. In the Provider URL field, enter the Active Directory URL in the following format:

```
ldap://<hostname_or_IP_address>:<port>
```

Entering the port is optional if the default value is used, 389.

 **Note** You can enter more than one Active Directory URL. In this case, make sure to press **Enter** after entering each one.

5. 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 you configure Active Directory authentication, you can specify the AD user authentication type when creating a new user. For details on how to do this, see [“Adding a new user” on page 83](#).

Setting up logging

This section describes the steps you must perform if you want to set up logging to help you analyze and troubleshoot the entire HYCU operation and the backup and restore functionality.

Accessing the Logging dialog box

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

To set up logging, follow these steps:

1. In the Logging dialog box, set the maximum log file size and the number of log files to keep. The default log file size is 10 MB.
2. Select one of the following log levels:

Log level	Description
All	All activity is recorded to log files.
Severe	Errors that affect the immediate operation of HYCU are recorded to log files.
Warning	Potentially harmful situations that do not represent an immediate threat to the operation of HYCU are recorded to log files.
Informational	Informational messages about the operation of HYCU are recorded to log files.


3. When required, you can download log files by clicking **Download Logs**. After you extract the `log.zip` file, check the log files at the following location:

```
/opt/grizzly/logs/
```

4. Click **Save**.

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 web server (the Grizzly server) or the entire appliance.

 **Important** Using SSH to perform any tasks other than restarting the HYCU web 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@<HYCU_backup_controller_IP_address>
```

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@<HYCU_backup_controller_IP_address>
```

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 web server

To manage the HYCU web server, follow these steps:

1. Open a remote session to the HYCU backup controller virtual machine:

```
ssh hycu@<HYCU_backup_controller_IP_address>
```

When requested, enter the password for the hycu user.

2. Perform the desired operation on the HYCU web server:


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


When requested, enter the password for the hycu user.

Using the command-line interface

You can manage your backup 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 directory 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` and `log4j.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 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.

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.

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 -Type Custom -Subject
"CN=<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.

Increasing the size of the HYCU virtual disk

If you are running out of disk space on your HYCU backup controller, you can increase the size of the HYCU virtual disk as needed. Depending on which hypervisor is running on your


Nutanix cluster, follow the instructions in one of the following sections:

- [“Increasing the HYCU disk size in a Nutanix AHV cluster” below](#)
- [“Increasing the HYCU disk size in a Nutanix ESXi cluster” below](#)

Increasing the HYCU disk size in a Nutanix AHV cluster

To increase the size of the HYCU virtual 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 virtual disk.
 - b. In the Size (GiB) field, increase the size of the disk as required.
 - c. Click **Update**.
6. Click **Power on** to turn on the HYCU backup controller.

Increasing the HYCU disk size in a Nutanix ESXi cluster

To increase the size of the HYCU virtual disk in a Nutanix ESXi cluster, follow these steps:

1. Log on to the VMware 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. Use the **Hard disk** slider to increase the size of the HYCU virtual disk, 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.

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 backup sources 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**. The Policies dialog box appears.
 - c. Click **Unassign**. The Unassign Policy prompt appears.
 - 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**. The Policies dialog box appears.
 - c. Click **Unassign**. The Unassign Policy prompt appears.
 - d. Click **Yes** to confirm that you want to unassign the policies from the selected applications.
2. On the HYCU backup controller virtual machine, run the `/opt/grizzly/bin/HycuCleanup.pl` script as follows:

```
sudo perl HycuCleanup.pl -c <Nutanix_cluster> -u <user_name> -p
<password>
```

In this instance, `<Nutanix_cluster>` is the name of the Nutanix cluster in URL format: `https://<server_name>:<port>`.


 **Important** By running this command, you will also remove all third-party snapshots created by using Nutanix REST API v3, not only those created by HYCU.

3. Log on to the Nutanix Prism web console by using your Nutanix logon credentials. Delete the HYCU backup controller virtual machine as follows:
 - a. In the menu bar, click **Home**, and then select **VM**.
 - b. Click the **Table** tab, and then, from the list of virtual machines, select the HYCU backup controller virtual machine.
 - c. Click  **Delete**. In the dialog box that appears, click **Delete** to confirm that you want to delete the HYCU backup controller virtual machine.

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 backup 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:

- [“Email notification settings” on the next page](#)
- [“Snapshot settings” on page 109](#)
- [“Utilization threshold settings” on page 109](#)
- [“Display settings” on page 109](#)
- [“SQL Server application settings” on page 110](#)
- [“Settings for aborting jobs” on page 110](#)
- [“Azure account settings” on page 110](#)

How to customize HYCU configuration settings

To customize HYCU configuration settings, follow these steps:

1. Open a remote session to the HYCU backup controller virtual machine:


```
ssh hycu@<HYCU_backup_controller_IP_address>
```

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 100](#).

2. Use the following command to access and open the `config.properties` file:

```
sudo vi /opt/grizzly/config.properties
```

 **Note** Because you will use the vi console text editor to customize HYCU configuration settings, basic knowledge of using the editor is required.

3. Edit any of the existing configuration settings as required.
4. Save and exit the `config.properties` file. To do so, press the **Esc** key, and then type the following:

```
:wq!
```

5. Restart the HYCU web server (the Grizzly server) for the changes to take effect:

```
sudo service grizzly restart
```

Email notification settings

If you want HYCU to send email notifications based on the occurrence of a specific event, adjust the following settings:

Setting	Description
<code>event.notification.smtp.server</code>	Host name or IP address of the SMTP server from which email notifications are sent.
<code>event.notification.smtp.port</code>	Port number of the SMTP server from which email notifications are sent (usually set to 25).
<code>event.notification.smtp.tls</code>	If the server from which email notifications are sent uses TLS, set this setting to true. Otherwise, leave the default false value.
<code>event.notification.smtp.ssl</code>	If the server from which email notifications are sent uses SSL, set this setting to true. Otherwise, leave the default false value.
<code>event.notification.smtp.username</code>	User name of the account on the SMTP server from which email notifications are sent.
<code>event.notification.smtp.password</code>	Password of the account on the SMTP server from which email notifications are sent.

Setting	Description
event.notification.sender.email	Email from which email notifications are sent.
event.notification.sender.name	Display name of the email sender.
event.notification.sender.subject	Subject for the sent email.
event.notification.green.receiver.list	<p>List of email recipients that will get the Info notifications. To separate multiple email recipients, use a comma.</p> <p>If you leave this setting blank, HYCU will not send email notifications for this severity level of the event.</p>
event.notification.orange.receiver.list	<p>List of email recipients that will get the Warning notifications. To separate multiple email recipients, use a comma.</p> <p>If you leave this setting blank, HYCU will not send email notifications for this severity level of the event.</p>
event.notification.red.receiver.list	<p>List of email recipients that will get the Error notifications. To separate multiple email recipients, use a comma.</p> <p>If you leave this setting blank, HYCU will not send email notifications for this severity level of the event.</p>
event.notification.green.category.block.list	<p>List of blocked event categories for the Info notifications for which no email notifications will be sent. To separate multiple blocked categories, use a comma.</p> <p>If you leave this setting blank, HYCU will send email notifications for all categories of this severity level.</p>
event.notification.orange.category.block.list	<p>List of blocked event categories for the Warning notifications for which no email notifications will be sent. To separate multiple blocked categories, use a comma.</p> <p>If you leave this setting blank, HYCU will send email notifications for all categories of this severity level.</p>

Setting	Description
event.notification.red.category.block.list	<p>List of blocked event categories for the Error notifications for which no email notifications will be sent. To separate multiple blocked categories, use a comma.</p> <p>If you leave this setting blank, HYCU will send email notifications for all categories of this severity level.</p>

Snapshot settings

You can customize the maximum number of snapshots to be retained by adjusting these settings:

Setting	Description
max.snapshots.per.vm	Maximum number of snapshots that are retained per virtual machine. The default value is 24.
max.snapshots.per.cluster	Maximum number of snapshots that are retained per Nutanix cluster. The default value is 2400.

Utilization threshold settings

You can use the following settings to configure the backup target utilization thresholds:

Setting	Description
target.utilization.threshold.health.red	If the HYCU backup controller utilization of the backup target exceeds the specified value, its health status indicator becomes red. The default value is 0.95.
target.utilization.threshold.health.yellow	If the HYCU backup controller utilization of the backup 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 backup target, see [“Viewing backup target information” on page 76](#).

Display settings

You can use the following settings 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.
max.backups.displayed	Maximum number of backups that are displayed for a virtual machine. The default value is 100.

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 log compression is enabled by default (the default value is 1). If you want to disable it, make sure to set the value for this setting to 0.

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
jobs.abort.deadline.minutes	Time (in minutes) within which a job must be completed. The default value is 1440.
jobs.abort.interval.minutes	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 jobs.abort.deadline.minutes setting. The default value is 15.

Azure account settings

You can use the following setting to configure which Azure cloud will be used by backup targets:

Setting	Description
target.azure.government.cloud	If set to true, backup targets use the Azure Government Cloud. Otherwise, they use the standard Azure public cloud. Keep in mind that the value you set for this configuration setting will affect all Azure backup targets.

Appendix B

Restoring to a different hypervisor

When restoring a virtual machine, you can choose to restore it to a different hypervisor by using the **Clone VM** option:

- From Nutanix ESXi to Nutanix AHV

For details, see [“Restoring a Nutanix ESXi virtual machine to a Nutanix AHV cluster” below](#).

For details on how to restore a virtual machine by using the Clone VM option, see [“Restoring an entire virtual machine” on page 46](#).

Restoring a Nutanix ESXi virtual machine to a Nutanix AHV cluster

The additional steps that are required for a successful restore of a Nutanix ESXi virtual machine to a Nutanix AHV cluster differ based on whether VMware Tools was installed on the virtual machine.

Prerequisites

- A Nutanix AHV cluster is added to HYCU. For details on how to do this, see [“Adding Nutanix clusters” on page 20](#).
- *Only for a Windows virtual machine on which no VMware Tools was installed.* The Nutanix VirtIO package was installed on the virtual machine that you want to restore to a Nutanix AHV cluster before the virtual machine was backed up.

Depending on whether VMware Tools was installed on the virtual machine that you want to restore to a Nutanix AHV cluster, do one of the following:

- VMware Tools was installed on the virtual machine:
 1. Make sure vCenter Server is configured in a way that allows the virtual machine to be reachable also from the associated Nutanix ESXi cluster.
 2. Through the Nutanix Prism web console, enable, mount, and install the Nutanix Guest Tools software bundle on the virtual machine that you want to restore to a

Nutanix AHV cluster.

For details on Nutanix Guest Tools requirements and limitations, as well as how to install the software bundle on a virtual machine, see Nutanix documentation.

3. Perform a restore of the selected virtual machine by selecting the **Clone VM** option.

After the restore is performed, the virtual machine can be turned on on a Nutanix AHV cluster.

- No VMware Tools was installed on the virtual machine:

Windows virtual machine	<ol style="list-style-type: none"> 1. Perform a restore of the selected virtual machine to a Nutanix AHV cluster by selecting the Clone VM option. <p>Note Wait a few moments for the restored virtual machine to turn on.</p> <ol style="list-style-type: none"> 2. Through the Nutanix Prism web console, enable, mount, and install the Nutanix Guest Tools software bundle on the virtual machine. <p>For details on Nutanix Guest Tools requirements and limitations, as well as how to install the software bundle on a virtual machine, see Nutanix documentation.</p>
Linux virtual machine	<p>After you perform a restore of the selected virtual machine to a Nutanix AHV cluster by selecting the Clone VM option, the restored virtual machine might not turn on or goes into rescue mode. To be able to turn on the virtual machine, for each of the virtual machine disks, do the following:</p> <ol style="list-style-type: none"> 1. Make sure that the restored virtual machine is turned off. 2. As an administrator, log on to the Nutanix AHV cluster by using SSH. 3. List the virtual machine details: <pre>accli vm.get <VM_name></pre> <ol style="list-style-type: none"> 4. Take a note of the current bus and index values in the <code>disk_list</code> section. 5. Clone the existing disk to a new disk on the compatible bus: <pre>accli vm.disk_create <VM_name> bus=<bus_type> clone_from_vmdisk=vm:<VM_name>:<current_bus>.<current_index></pre> <p>In this instance, <code><VM_name></code> is the name of the restored virtual machine, <code><bus_type></code> is <code>scsi</code>, <code>ide</code>, or <code>sata</code>, <code><current_bus></code> is the bus value from the <code>disk_list</code> section, and <code><current_index></code> is the index value from the <code>disk_list</code> section.</p> <p>If the original virtual machine has the SATA or SCSI disks, clone them</p>

to the SATA disks. For example:

```
acli vm.disk_create test-vm bus=sata
clone_from_vmdisk=vm:test-vm:scsi.0
```

If the original virtual machine has the IDE disks, clone them to the IDE disks. For example:

```
acli vm.disk_create test-vm bus=ide
clone_from_vmdisk=vm:test-vm:scsi.0
```

After you perform the previous procedure for all the disks, do the following:

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. Select the boot disk and delete source disks, and then click **Save**.
6. Click **Power on** to turn on the restored virtual machine.
7. Enable, mount, and install the Nutanix Guest Tools software bundle on the virtual machine.

For details on how to update a virtual machine on a Nutanix cluster, see Nutanix documentation.

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!

