HYCU R-Cloud Hybrid Cloud Edition v5.3.0

User Guide

July 2026

Legal notices

Copyright notice

© 2026 HYCU. All rights reserved.

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

Trademarks

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

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

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

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

CentOS is a trademark of Red Hat, Inc. (“Red Hat”).

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

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

Google Chrome™, Google Cloud™, and Google Cloud Storage™ are trademarks of Google LLC.

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

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

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

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

"SUSE" and the SUSE logo are trademarks of SUSE LLC or its subsidiaries or affiliates.

Ubuntu is a registered trademark of Canonical Ltd.

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

XenServer and XenCenter are trademarks of Citrix Systems, Inc. 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 17

HYCU key features and benefits 18

Data protection environment overview 20

HYCU data protection 22

2 Deploying the HYCU virtual appliance 23

HYCU instances 25

Sizing resources for your HYCU backup infrastructure 27

Adjusting firewall configuration 28

Targets 30

Virtual machine and volume group protection 32

Restore of individual files and application awareness 34

File share protection 36

Bucket protection 41

Infrastructural services 42

User interaction and administration 43

SpinUp 44

Adjusting antivirus configuration 44

Deploying HYCU to a Nutanix AHV cluster 45

Uploading the HYCU virtual appliance image to a Nutanix AHV cluster 46

Creating a virtual machine for HYCU deployment on a Nutanix AHV cluster 47

Configuring HYCU on the virtual machine 48

Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment 49

Deploying HYCU to a XenServer environment 52

Creating a virtual machine for HYCU deployment in a XenServer environment 53
Attaching a disk to the created virtual machine 53
Configuring HYCU on the created virtual machine 54
Deploying HYCU to an Azure Local environment 55
Uploading the HYCU virtual machine image to an Azure Local environment 56
Creating a virtual machine for HYCU deployment in an Azure Local environment 56
Configuring HYCU on the virtual machine 57
Deploying HYCU to a Hyper-V cluster 59
Uploading the HYCU virtual machine image to a Hyper-V cluster 60
Preparing a virtual machine for HYCU deployment on a Hyper-V cluster 60
Configuring HYCU on the virtual machine 62
Deploying HYCU to an AWS GovCloud (US) environment 64
Deploying HYCU to an Azure environment 65
Deploying HYCU to an Azure Government environment 67
Signing in to HYCU 68
Setting the language 70
3 Establishing a data protection environment 72
Adding sources 73
Adding a Nutanix cluster 74
Adding a vCenter Server 77
Adding a XenServer 78
Adding an Azure Local environment 80
Adding a Hyper-V cluster 80
Adding a cloud environment 81

Adding a file server 85
Adding an object server 98
Adding a server 99
Setting up targets 100
Setting up an NFS target 101
Setting up an SMB target 104
Setting up a Nutanix target 107
Setting up a Nutanix Objects target 109
Setting up an iSCSI target 112
Setting up an Azure target 115
Setting up an Amazon S3 / S3 Compatible target 117
Setting up a Google Cloud target 121
Setting up a tape target 124
Setting up a Data Domain target 128
Defining your backup strategy 130
Taking advantage of predefined policies 131
Creating a custom policy 132
Setting a default policy 152
4 Protecting virtual machines 154
Planning virtual machine protection 155
Preparing your data protection environment 155
Preparing for disaster recovery of your data 161
Preparing for the restore to a different source 163
Server specifics 168
SpinUp specifics 170
Enabling access to data 174
Setting up virtual machine backup configuration options 178

Backing up virtual machines ... 182
Restoring virtual machines ... 184
Restore options ... 186
Restoring a virtual machine ... 188
Cloning a virtual machine ... 205
Validating the virtual machine backup ... 228
Restoring virtual disks ... 233
Cloning virtual disks ... 235
Exporting virtual disks ... 238
Restoring individual files ... 243
5 Protecting applications ... 252
Enabling access to application data ... 252
Planning application protection ... 256
Preparing for SQL Server application protection ... 257
Preparing for Active Directory application protection ... 262
Preparing for Exchange Server application protection ... 262
Preparing for Oracle application protection ... 265
Preparing for SAP HANA application protection ... 267
Preparing for PostgreSQL application protection ... 268
Backing up applications ... 270
Restoring whole applications ... 272
Restore options ... 273
Restoring a virtual machine ... 274
Cloning a virtual machine ... 293
Restoring SQL Server databases ... 313
Restoring Exchange Server databases, mailboxes, and public folders ... 317
Restoring Oracle database instances and tablespaces ... 321

Restoring PostgreSQL database clusters ... 324

6 Protecting file shares ... 327

Configuring file share backup options ... 327

Backing up file shares ... 329

Restoring file share data ... 332

Restoring a whole file share or individual files ... 335

Restoring specific files ... 338

7 Protecting volume groups ... 343

Backing up volume groups ... 343

Restoring volume groups ... 344

Restoring a volume group ... 345

Cloning a volume group ... 347

Exporting virtual disks ... 348

8 Protecting buckets ... 350

Configuring bucket backup options ... 350

Backing up buckets ... 351

Restoring bucket data ... 353

Restoring a whole bucket or individual objects ... 356

Restoring specific objects ... 358

9 Recovering your data protection environment ... 362

Preparing for disaster recovery ... 363

Deploying a recovery HYCU backup controller ... 364

Importing targets ... 366

Performing disaster recovery ... 368

Restoring the HYCU backup controller to the original source ... 369

Restoring the HYCU backup controller to a different source ... 374

10 Performing daily tasks 381
Using the HYCU dashboard 382
Managing HYCU jobs 385
Managing HYCU events 387
Configuring event notifications 388
Setting up email notifications 388
Setting up webhook notifications 389
Enabling the purge of events and jobs 391
Using HYCU reports 392
Getting started with reporting 393
Viewing reports 395
Generating reports 396
Scheduling reports 397
Exporting and importing reports 398
Viewing entity details 399
Viewing the backup status of entities 401
Viewing the R-Shield status 403
Viewing the backup data size change rate 404
Viewing tier statuses 405
Filtering and sorting data 406
Filtering data in panels 406
Sorting data in panels 415
Exporting the contents of the panel 416
Managing targets 417
Viewing target information 417
Editing a target 419
Activating or deactivating a target 421

Removing a target ... 421
Managing policies ... 422
Viewing policy information ... 422
Editing a policy ... 423
Deleting a policy ... 423
Performing a manual backup ... 424
Setting up a validation policy ... 425
Running malware scans manually ... 431
Overriding the R-Shield status ... 433
Archiving data manually ... 435
Recreating snapshots ... 436
Adjusting the HYCU backup controller resources ... 438
Setting up the appearance of your HYCU web user interface ... 438
11 Managing users ... 440
HYCU groups ... 440
User roles ... 442
Setting up a user environment ... 444
Creating a user ... 445
Adding a user to a group ... 450
Creating a self-service group ... 451
Setting ownership ... 452
Activating or deactivating users or self-service groups ... 457
Switching to another group ... 458
Updating your user profile ... 459
12 Administering ... 461
Adding a cloud account ... 462

Adding a cloud provider account 462
Adding a HYCU account 470
Configuring target encryption 471
Exporting an encryption key 471
Importing an encryption key 471
Integrating HYCU with identity providers 471
Adding an identity provider to HYCU 472
Managing HYCU instances 484
Creating a HYCU instance by using the HYCU web user interface 484
Viewing HYCU instance information 486
Deleting a HYCU instance 487
Setting the iSCSI initiator secret 487
Licensing 488
Creating a license request 490
Requesting and retrieving licenses 491
Activating licenses 492
Setting up logging 493
Configuring your network 496
Changing network settings 496
Limiting network bandwidth 498
Enabling external monitoring of HYCU 500
Managing data retention 501
Filtering restore point tiers 506
Expiring data automatically 507
Expiring data from the entity panel 507
Setting power options 509
Managing secrets 509

Adding a Conjur configuration 511
Editing a Conjur configuration 512
Removing a Conjur configuration 512
Configuring email delivery 512
Configuring email delivery using SMTP 513
Configuring email delivery for Microsoft 365 514
Upgrading HYCU 515
Creating the JSON index file 518
Applying HYCU updates 519
Applying an update by using the HYCU web user interface 521
Applying an update by using the shell script 523
Configuring SSL certificates 524
Creating a self-signed certificate 525
Creating a certificate signing request 526
Importing a custom certificate 527
Sharing telemetry data with HYCU 530
Removing HYCU 531
13 Tuning your data protection environment 534
Accessing the HYCU backup controller virtual machine by using SSH 535
Enabling HTTPS for WinRM connections 538
Using Kerberos authentication 538
Configuring FIPS mode for HYCU 540
Enabling FIPS mode for HYCU 540
Disabling FIPS mode for HYCU 541
Setting up LDAPS authentication 541
Setting up two-factor authentication 542
Managing API keys 543

Generating an API key ... 543
Revoking an API key ... 544
Managing FIDO authenticators ... 544
Adding a new FIDO authenticator ... 544
Revoking a FIDO authenticator ... 545
Securing SMTP connections ... 545
Setting up HYCU to use multiple networks ... 545
Setting up HYCU to use multiple networks on a Nutanix AHV or Nutanix ESXi cluster ... 547
Setting up HYCU to use multiple networks in a vSphere environment ... 548
Setting up HYCU to use multiple networks in a XenServer environment ... 549
Setting up HYCU to use multiple networks in an Azure Local environment ... 550
Setting up HYCU to use multiple networks on a Hyper-V cluster ... 551
Increasing the size of the HYCU virtual disks ... 552
Assigning privileges to a vSphere user ... 552
Configuring Prism Central user permissions ... 556
Setting permissions in AWS GovCloud (US) ... 558
Using the HYCU REST API Explorer ... 560
Using the command-line interface ... 560
Using the pre and post scripts ... 561
14 Monitoring data protection environments ... 563
Managing HYCU controllers ... 564
Adding a HYCU controller ... 564
Viewing information about HYCU controllers ... 567
Editing a HYCU controller ... 569

Removing a HYCU controller ... 569
Using the HYCU Manager console ... 569
Viewing entity data ... 571
Viewing events ... 573
Using HYCU Manager reports ... 574
Creating reports ... 575
Scheduling reports ... 575
Generating reports ... 577
Downloading and viewing reports ... 577
Performing administration tasks ... 577
15 Protecting data across on-premises and cloud environments ... 586
Protecting data across on-premises and AWS environments ... 586
Migrating virtual machines across different environments ... 587
Performing disaster recovery of data to AWS ... 597
Protecting data across on-premises and Google Cloud environments ... 599
Migrating virtual machines across different environments ... 600
Performing disaster recovery of data to Google Cloud ... 607
Protecting data across on-premises and Azure environments ... 609
Migrating virtual machines across different environments ... 609
Performing disaster recovery of data to Azure ... 617
Protecting data across on-premises and Azure Government environments ... 619
Migrating virtual machines to cloud ... 620
Performing disaster recovery of data to Azure Government ... 625
A Customizing HYCU configuration settings ... 628
Snapshot settings ... 630

Utilization threshold settings 630
SQL Server application settings 631
Settings for aborting jobs 631
File server settings 631
Object server settings 633
Data rehydration settings 633
Disaster recovery settings 634
User management settings 634
Upgrade settings 635
B After restoring a virtual machine to a different source 636
After restoring a virtual machine to a Nutanix AHV cluster 637
After restoring a virtual machine to a Nutanix ESXi cluster 638

Chapter 1

About HYCU

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

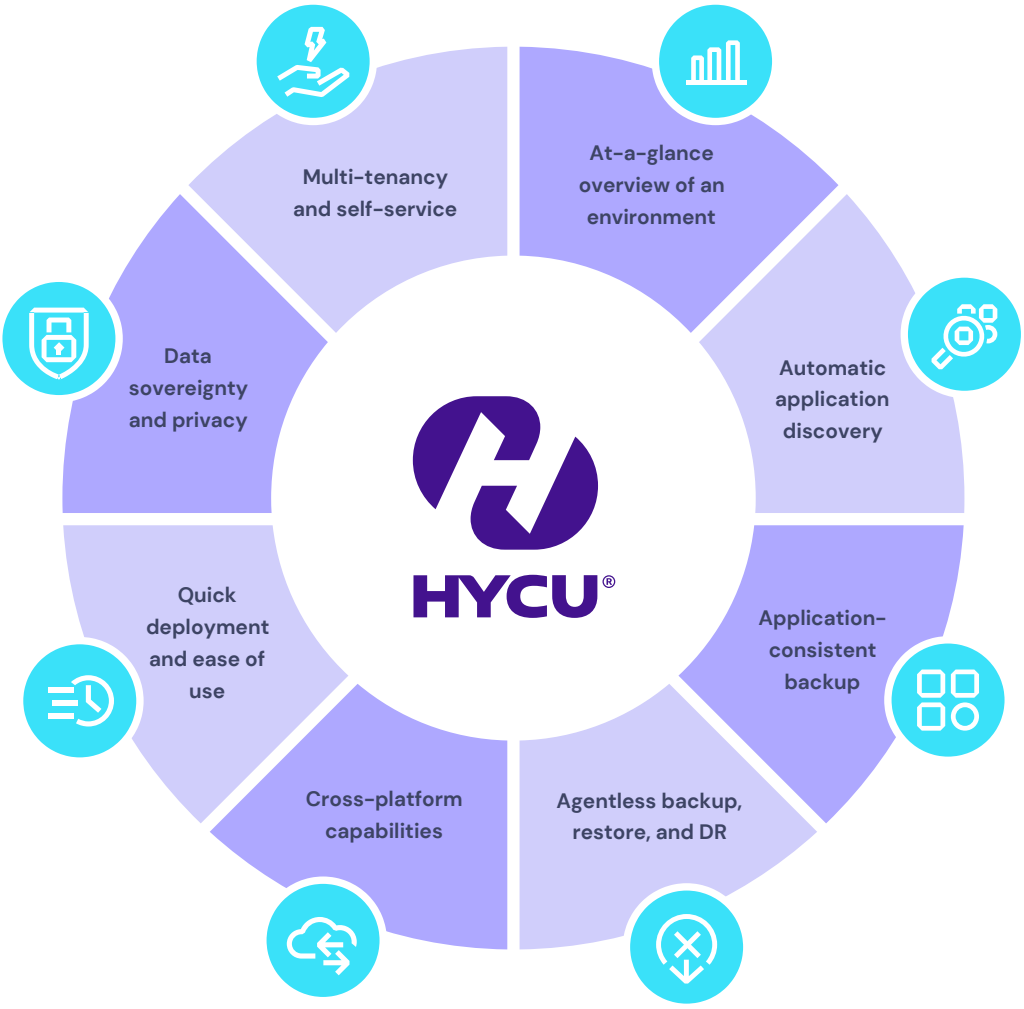

Figure 1-1: Introduction to HYCU

HYCU key features and benefits

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

- **Protects against data loss**

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

- **Simplifies deployment**

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

- **Provides new-found visibility**

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

- **Protects data in a few minutes**

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

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

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

- **Schedules backups based on RPOs**

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

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

- **Lets you choose targets and sources**

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

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

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

- **Offers a scalable backup for file servers and object servers**

  Cuts down the time it takes to back up file shares and buckets, saves a significant amount of computing resources, and allows you to take more frequent backups, reducing the amount of data loss in case of a failure.

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

  The SpinUp functionality allows you to migrate protected data between the on-premises and cloud infrastructures (AWS, Google Cloud, Azure, or Azure Government environments). In the event of a disaster, the SpinUp functionality provides disaster recovery of mission-critical data to cloud.

- **Provides an efficient ROBO data protection solution for Nutanix**

  Backs up remote office/branch office (ROBO) data from data center replicas and enables a one-click restore within the data center or at any remote location.

Data protection environment overview

The data protection environment consists of the following components:

HYCU backup controller	A virtual machine that processes data collected from sources and presents it in the web user interface.
HYCU web user interface	An interface for protecting entities and administering the data protection environment.
Targets	Storage locations that HYCU uses for storing the protected data. Protected data can also be stored as snapshots.
Sources	Environments for which HYCU provides data protection—Nutanix clusters, vSphere environments, XenServer environments, Azure Local environments, Hyper-V clusters, AWS GovCloud (US) environments, Azure environments, Azure Government environments, file servers, object servers, and servers.
Entities	Objects to which you can assign a policy and for which you therefore provide data protection—virtual machines, applications, file shares, servers, Nutanix volume groups, buckets, and vSphere virtual machine templates. Data is always protected at a granular level, allowing you to restore either the whole entities or their parts (disks and application items).

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

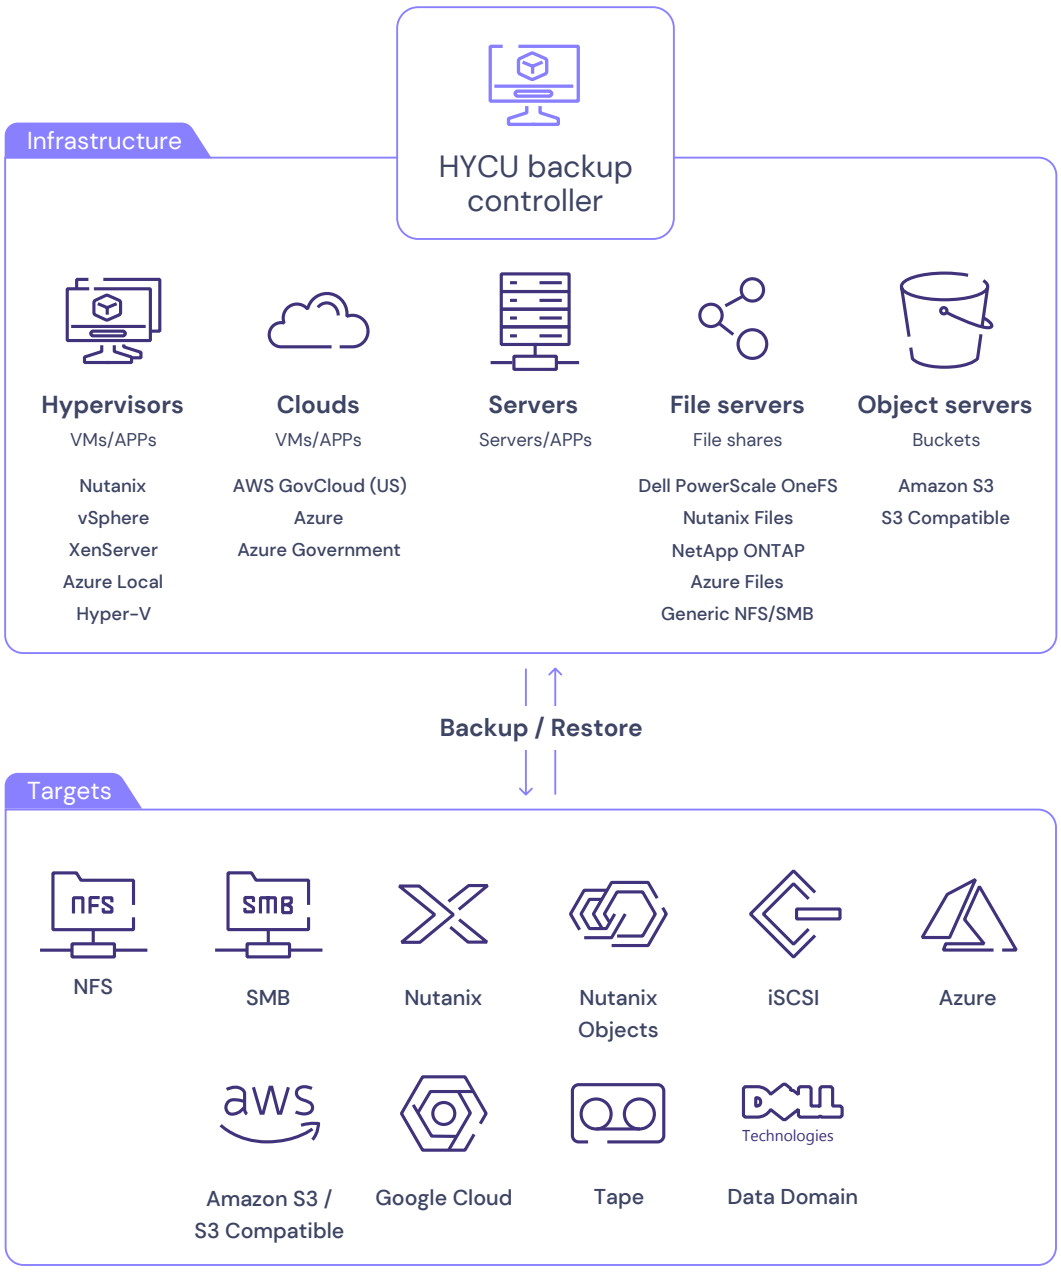

Figure 1-2: HYCU architecture

HYCU data protection

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

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

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

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

- “Protecting virtual machines” on page 154
- “Protecting applications” on page 252
- “Protecting file shares” on page 327
- “Protecting volume groups” on page 343
- “Protecting buckets” on page 350

Chapter 2

Deploying the HYCU virtual appliance

The HYCU virtual appliance is a preconfigured software solution that you can easily deploy to an environment for which you want to provide data protection.

Deployment modes

Mode	Description
HYCU Backup Controller	Enables you to protect virtual machines (including virtual machine templates), applications, file shares, servers, volume groups, and buckets. A HYCU backup controller is a virtual machine that processes data collected from the sources and presents it in its web user interface. **Note** Protecting file shares and buckets requires also a HYCU instance.
HYCU Instance	Enables you to protect file shares and buckets, speed up data protection operations related to vSphere virtual machines, and to optimize the use of resources when scanning for malware. For details on HYCU instances, see "HYCU instances" on page 25.
HYCU Manager	Enables you to manage HYCU controllers. HYCU Manager is a virtual machine residing in the source environment that collects data from all HYCU controllers in your on-premises and cloud data protection environments, and presents it in the web user interface.

For details on the supported environments for each deployment mode, see the *HYCU Compatibility Matrix*.

> **Important** Installing any third-party software on the HYCU backup controller, the HYCU instance, or HYCU Manager is not supported.

Deployment steps

Step	Instructions
1. Size the backup infrastructure for HYCU.	"Sizing resources for your HYCU backup infrastructure" on page 27
2. *Only if firewalls are configured on your network.* Open relevant ports in each involved firewall.	"Adjusting firewall configuration" on page 28
3. Customize antivirus settings.	"Adjusting antivirus configuration" on page 44
4. Deploy the HYCU virtual appliance to a source.	• "Deploying HYCU to a Nutanix AHV cluster" on page 45 • "Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment" on page 49 • "Deploying HYCU to a XenServer environment" on page 52 • "Deploying HYCU to an Azure Local environment" on page 55 • "Deploying HYCU to a Hyper-V cluster" on page 59 • "Deploying HYCU to an AWS GovCloud (US) environment" on page 64 • "Deploying HYCU to an Azure environment" on page 65 • "Deploying HYCU to an Azure Government environment" on page 67

After you successfully deploy the HYCU virtual appliance, you can access HYCU by using a supported web browser. For details on how to sign in to HYCU, see “Signing in to HYCU” on page 68.

HYCU instances

The HYCU instance is a virtual machine that HYCU uses for performing file share, bucket, and vSphere virtual machine data protection operations, and, optionally, for running the malware scans of backup data as part of the R-Shield policies.

Using HYCU instances for file share or bucket protection

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

Using HYCU instances for vSphere virtual machine protection

The HYCU instance can be used to protect vSphere virtual machine data if you want HYCU to use the HotAdd transport when backing up and restoring data, and when archiving data from snapshots. This speeds up the data protection operations. If you do not create a HYCU instance, the data protection is performed by the HYCU backup controller by using the NBDSSL transport. For details on VMware transport methods, see VMware documentation.

Using HYCU instances for malware scanning

If available, malware scanning of backup data is automatically performed by the HYCU instance. This helps optimizing the use of HYCU backup controller capacity. If you do not create a HYCU instance, malware scanning is performed by the HYCU backup controller.

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

- By deploying the HYCU virtual appliance and selecting the HYCU Instance mode. For details, see “Deploying the HYCU virtual appliance” on page 23.

- By using the HYCU web user interface. For details, see “Creating a HYCU instance by using the HYCU web user interface” on page 484.

  > ⓘ **Important** *For Azure Local environments and Hyper-V clusters:* You cannot use this method for creating a HYCU instance.

If you later decide to remove any HYCU instance from your data protection environment, you can do so as described in “Deleting a HYCU instance” on page 487.

Considerations

- You can create a HYCU instance before or after adding a source to HYCU.
- The created HYCU instance connects automatically to the corresponding HYCU backup controller.
- Each HYCU instance is by default created with 16 GiB of RAM, 1 CPU, 8 CPU cores, and the data disk size of 128 GiB. However, this can be overridden by setting the following configuration settings to the preferred values:
  - `afs.instance.memory.mb`
  - `afs.instance.cpu`
  - `afs.instance.cores.per.cpu`
  - `afs.instance.datadisk.size.gb`

  For details on how to customize the HYCU configuration settings, see “Customizing HYCU configuration settings” on page 628.
- If you change the host name or IP address of the HYCU backup controller, you must also change it on all connected HYCU instances. On each connected HYCU instance, update the `catalog.master.url` configuration setting in the `/hycudata/opt/grizzly/config.properties` file.
- *For HYCU instances running on Nutanix clusters:* When distributing the load among multiple HYCU instances, HYCU automatically prioritizes the HYCU instances that are running on the same Nutanix cluster as the HYCU backup controller and the file server. However, by changing the `afs.instance.afs.cluster.priority` or `afs.instance.bc.cluster.priority` configuration setting, you can adjust the load distribution process to your needs. For details on how to customize the HYCU configuration settings, see “Customizing HYCU configuration settings” on page 628.

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

Sizing resources for your HYCU backup infrastructure

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

HYCU virtual machine (HYCU backup controller, HYCU instance, HYCU Manager)

- Network connection:

  Make sure that you reserve an IP address for your virtual machine.
- System requirements:
  - Minimum requirements are 4 CPU cores and 4 GiB of RAM.
  - *For deploying in the HYCU Backup Controller mode:* Keep in mind that aspects beyond the size of your data protection environment affect the system requirements. Performance of the sources, target efficiency, the chosen backup strategy, backup data compression, and malware scanning may all increase or decrease the need for specific resources.

    For example, if you plan to copy and archive backup data, the number of required targets increases. Similarly, if you specify a short RPO or a small backup threshold, the load on your backup infrastructure increases and HYCU requires more storage and compute resources.

    Consider the following recommendations:

Number of VMs in the environment	System requirements				
	vCPU	Cores	Memory (in GiB)	OS disk (in GiB)	Data disk (in GiB)
Fewer than 50	8	1	8	20	128
50–200	8	2	16	20	128
200–500	16	2	32	20	128

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

System requirements for malware scanning

If you plan to enable malware scanning in your R-Shield policy, consider the following recommendations:

Number of VMs in the environment	System requirements				
	vCPU	Cores	Memory (in GiB)	OS disk (in GiB)	Data disk (in GiB)
Fewer than 50	8	2	16	20	128
50–200	16	2	32	20	128
200–500	16	4	64	20	128
More than 500	The figures vary. Contact HYCU Support.				

HYCU web user interface

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

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

Targets

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

Adjusting firewall configuration

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

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

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

- "Targets" on the next page
- "Virtual machine and volume group protection" on page 32
- "Restore of individual files and application awareness" on page 34
- "File share protection" on page 36
- "Bucket protection" on page 41
- "Infrastructural services" on page 42
- "User interaction and administration" on page 43
- "SpinUp" on page 44

Targets

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup of data to an NFS v4 target	HYCU backup controller HYCU instance	NFS v4 server	2049	TCP UDP
Backup of data to an NFS v3 target	HYCU backup controller HYCU instance	NFS v3 server	111 2049 mountd[a]	TCP UDP
Backup of data to an SMB target	HYCU backup controller HYCU instance	SMB server	445	TCP
Backup of data to an iSCSI target	HYCU backup controller	iSCSI server	3260	TCP
Backup of data to a cloud target	HYCU backup controller HYCU instance	Cloud server	443[b]	TCP
Archive of data to a QStar NFS target	HYCU backup controller HYCU instance	QStar server	111 2049 mountd[a] 18082[c]	TCP
Archive of data to a QStar SMB target	HYCU backup controller HYCU instance	QStar server	445 18082[c]	TCP
Backup of data to a Data Domain target	HYCU backup controller HYCU instance	Data Domain server	2049[d]	TCP
Backup of data to an SMB target when using Kerberos authentication	HYCU backup controller HYCU instance	Kerberos Domain Controller (KDC)	88	TCP UDP

[a] For details on the port number, see NFS server documentation.
[b] Cloud targets may utilize multiple IP addresses. For details on IP ranges used by public clouds, see respective cloud documentation.
[c] This is the default port for HTTPS connection, but other ports can also be used. HTTP connection is also supported, but it is not recommended.
[d] Data Domain servers by default use port 2049, but other ports can also be used. For instructions on how to change the port, see Dell documentation.

Virtual machine and volume group protection

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Data protection of VMs on a Nutanix cluster or volume groups[a]	HYCU backup controller	Cluster virtual server (cluster virtual IP address)[b]	3205 3260	TCP
		iSCSI target discovery portal (iSCSI Data Services IP address)[c]		
Backup of entities in a vSphere environment	HYCU backup controller	ESXi hosts	902	TCP
		vCenter Server	443	
Data protection of VMs in a vSphere environment by using the HotAdd transport	HYCU backup controller[d]	HYCU instance[d]	8443	TCP
	HYCU instance	HYCU backup controller		
		ESXi hosts	902	
		vCenter Server	443	
Data protection of VMs in a XenServer environment	HYCU backup controller	XenServer	443	TCP
Data protection of VMs in an Azure Local environment	HYCU backup controller	WinRM	5985 5986	HTTP HTTPS
Data protection of VMs on a Hyper-V cluster	HYCU backup controller	WinRM	5985 5986	HTTP HTTPS

[a] HYCU accesses Nutanix Volumes.

[b] Only if a cluster virtual IP address is specified for the Target Portal option in the iSCSI target

configuration in HYCU.
[c] Only if an iSCSI Data Services IP address is specified for the Target Portal option in the iSCSI target configuration in HYCU.
[d] The connection from the HYCU backup controller to the HYCU instance uses the HTTP Upgrade request to upgrade the connection to WebSocket. Make sure that your firewall is not configured to block such requests.

Restore of individual files and application awareness

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

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
		share		
Restore of files to an NFS file share	HYCU backup controller	System with an NFS file share	NFS4: 2049 NFS3: 111, mountd[b]	TCP
Restore of files to the local machine	System where the HYCU interface is accessed	HYCU backup controller	8443	TCP
Restore of individual files and application awareness for Windows VMs when using Kerberos authentication	HYCU backup controller	Kerberos Domain Controller (KDC)	88	TCP UDP
	HYCU instance			

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

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

File share protection

General

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup and restore of file shares	HYCU backup controller	HYCU instance	8443	TCP
	HYCU instance	HYCU backup controller		
Backup and restore of SMB file shares when using Kerberos authentication	HYCU backup controller	Kerberos Domain Controller (KDC)	88	TCP UDP
	HYCU instance			

Nutanix Files

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

[a] Only if HYCU accesses file shares by using the SMB protocol.
[b] Only if HYCU accesses file shares by using the NFS protocol.
[c] *NFSv3 only.* The actual port numbers vary from one vendor to another. The port numbers are customizable. You can run the `rpcinfo -p <NFSServerIP/Hostname>` command on the HYCU backup controller to determine the currently configured ports. For details, see NFS server documentation.
[d] HYCU uses the Nutanix REST API v3.

Dell PowerScale OneFS

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

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
HYCU instances on a Nutanix ESXi cluster or in a vSphere environment				

[a] Only if HYCU accesses file shares by using the SMB protocol.
[b] Only if HYCU accesses file shares by using the NFS protocol.
[c] *NFSv3 only.* The actual port numbers vary from one vendor to another. The port numbers are customizable. You can run the `rpcinfo -p <NFSServerIP/Hostname>` command on the HYCU backup controller to determine the currently configured ports. For details, see NFS server documentation.

NetApp ONTAP

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

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
HYCU instances on a Nutanix ESXi cluster or in a vSphere environment				

[a] Only if HYCU accesses file shares by using the SMB protocol.
[b] Only if HYCU accesses file shares by using the NFS protocol.
[c] *NFSv3 only.* The actual port numbers vary from one vendor to another. The port numbers are customizable. You can run the `rpcinfo -p <NFSServerIP/Hostname>` command on the HYCU backup controller to determine the currently configured ports. For details, see NFS server documentation.

Azure Files

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup and restore of Azure Files file shares	HYCU backup controller	Azure server	443 445	TCP
	HYCU instance			

Generic file shares

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

[a] Only if HYCU accesses file shares by using the SMB protocol.
[b] Only if HYCU accesses file shares by using the NFS protocol.
[c] *NFSv3 only.* The actual port numbers vary from one vendor to another. The port numbers are customizable. You can run the `rpcinfo -p <NFSServerIP/Hostname>` command on the HYCU backup controller to determine the currently configured ports. For details, see NFS server documentation.

Bucket protection

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Backup and restore of buckets	HYCU backup controller	HYCU instance	8443	TCP
	HYCU instance	HYCU backup controller		
	HYCU backup controller	Object server	80 443	
	HYCU instance			

Infrastructural services

Purpose	Communication endpoints		Ports at destination	Protocol
	Source	Destination		
Use of a DNS server	HYCU backup controller	DNS server	53	TCP UDP
	HYCU instance			
Use of an NTP server	HYCU backup controller	NTP server	123	UDP
	HYCU instance			
Use of an LDAP server	HYCU backup controller	LDAP server	LDAP: 389 LDAPS: 636	TCP
Use of the SMTP protocol for sending email notifications	HYCU backup controller	SMTP server	25[a]	TCP
Sharing telemetry data with HYCU	HYCU backup controller	Telemetry host: callhome.hycu.com[b]	443	TCP
		Data host: telemetry-production-bucket-us.s3.us-east-1.amazonaws.com[c]		
Monitoring HYCU	Monitoring tool	HYCU backup controller	8443	TCP
Upgrading HYCU from the HYCU repository	HYCU backup controller	Download host: hdn.hycu.com	443	TCP

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

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

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

User interaction and administration

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

SpinUp

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

[a] The following host names are used: authentication.r-cloud.hycu.com, registry.r-cloud.hycu.com, r-cloud.hycu.com, and the host name of the HYCU R-Cloud manager. For details, contact HYCU Support.
[b] The following host names are used: registry.azure.hycu.com and the host name of the HYCU for Azure manager. For details, contact HYCU Support.

Adjusting antivirus configuration

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

For details on the data protection scenarios when HYCU must be given access to data, see “Enabling access to data” on page 174.

Considerations

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

Deploying HYCU to a Nutanix AHV cluster

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

Prerequisite

The backup infrastructure must be sized according to the requirements described in “Sizing resources for your HYCU backup infrastructure” on page 27.

Limitation

HYCU does not support UEFI Secure Boot.

Deployment tasks

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

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

Task	Instructions
	cluster” on page 47
3. Configure HYCU on the created virtual machine.	“Configuring HYCU on the virtual machine” on page 48

Uploading the HYCU virtual appliance image to a Nutanix AHV cluster

Procedure

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

      > ⓘ Important The HYCU virtual appliance image must be uploaded to the Nutanix AHV cluster in the following format: `hycu-<Version>-<Revision>`
      >
      > For example: `hycu-5.3.0-3634`
      >
      > If you enter the HYCU image name in a different format, you will not be able to use this image for an upgrade.

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

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

Procedure

1. In the menu bar in the Nutanix Prism web console, click **Home**, and then select **VM**.
2. Click **Create VM** at the upper right of the screen.
3. In the Create VM dialog box, provide the following information:
   a. In the General Configuration section, do the following:
      i. Enter a virtual machine name and, optionally, its description.
      ii. Make sure that the time zone is set to UTC (the default value).
      iii. Leave the Use this VM as an agent VM check box clear.
   b. In the Compute Details section, enter the number of virtual CPUs and cores per virtual CPU, and the amount of memory to allocate to this virtual machine.
   c. In the Disks section, click **Add New Disk**, and then, in the Add Disk dialog box, specify a system disk:
      i. From the Type drop-down menu, select **DISK**.
      ii. From the Operation drop-down menu, select **Clone from Image Service**.
      iii. From the Bus Type drop-down menu, select **SCSI**.
      iv. From the Image drop-down menu, select the image you uploaded.
      v. In the Size (GiB) field, leave the default size of the system disk (20 GiB).

      > **Note** You can later increase the size of the system disk if needed. For details, see "Increasing the size of the HYCU virtual disks" on page 552.

      vi. Click **Add**.
   d. In the Boot Configuration section, select the preferred virtual machine boot mode.
   e. In the Disks section, click **Add New Disk**, and then, in the Add Disk dialog box, specify a data disk:
      i. Leave the default values for the type of storage device, the device contents, and the bus type.

ii. From the Storage Container drop-down menu, select a storage container for the image to be uploaded.

iii. In the Size (GiB) field, enter 128.

> **Note** You can later increase the size of the data disk if needed. For details, see "Increasing the size of the HYCU virtual disks" on page 552.

iv. Click **Add**.

4. In the Network Adapters (NIC) section, click **Add New NIC**, and then select a VLAN and click **Add**.
5. Click **Save**.

Configuring HYCU on the virtual machine

Procedure

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

   For details on deployment modes, see "Deployment modes" on page 23.
4. Tab to **OK** and press **Enter**.
5. In the Network Configuration dialog box that opens, do the following:
   a. Enter the values for the following:
      - *Optional.* Host name for the virtual machine

        The default host name is generated automatically during the HYCU virtual appliance deployment. If you want to use a custom host name, keep in mind the following:
        - *Only if you selected the HYCU Backup Controller or HYCU Manager mode.* The host name should begin with a letter and may contain only letters, numbers, and hyphens (-).

- *Only if you selected the HYCU Instance mode.* For host name naming conventions, see "Managing HYCU instances" on page 484.

- IPv4 address (for example, `10.1.100.1`)
- Subnet mask (for example, `255.0.0.0`)
- Default gateway (for example, `10.1.1.1`)
- *Optional.* DNS server (for example, `10.1.1.5`)
- *Optional.* Search domain (for example, `domain.com`)

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

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

The progress of the HYCU configuration displays.

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

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

The progress of the HYCU backup controller assignment displays.

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

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see "Licensing" on page 488.

Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment

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

Prerequisite

The backup infrastructure must be sized according to the requirements described in “Sizing resources for your HYCU backup infrastructure” on page 27.

Consideration

The HYCU backup controller uses the time zone as configured on the vCenter Server.

Limitation

HYCU does not support UEFI Secure Boot.

Procedure

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

URL	Specify a URL to the HYCU OVF package.
Local file	Browse your file system for the HYCU OVF package. **Important** When you are browsing your file system, make sure to select both the `.ovf` file and the `.vmdk` file related to the OVF package.

   Click **Next**.
4. In the Select name and location section, enter a name for the HYCU virtual machine and specify a location where you want to deploy it, and then click **Next**.
5. In the Select a resource section, select where to run the deployed package, and then click **Next**.
6. In the Review details section, verify the package details, and then click **Next**.
7. In the Select Configuration section, do the following:
   a. Select a deployment configuration:
      - **HYCU Backup Controller**
      - **HYCU Instance**
      - **HYCU Manager**

For details on deployment modes, see "Deployment modes" on page 23.

b. Click **Next**.

8. In the Select storage section, select where to store the files for the deployed package, and then click **Next**.
9. In the Select networks sections, leave the default values, and then click **Next**.
10. In the Customize template section, enter the values for the following:
    - *Optional.* Host name for the virtual machine

      The default host name is generated automatically during the HYCU virtual appliance deployment. If you want to use a custom host name, keep in mind the following:
      - *Only if you selected the HYCU backup controller or HYCU Manager mode.* The host name should begin with a letter and may contain only letters, numbers, and hyphens (-).
      - *Only if you selected the HYCU instance mode.* For host name naming conventions, see "Managing HYCU instances" on page 484.
    - IPv4 address (for example, `10.1.100.1`)
    - Subnet mask (for example, `255.0.0.0`)
    - Default gateway (for example, `10.1.1.1`)
    - *Optional.* DNS server (for example, `10.1.1.5`)
    - *Optional.* Search domain (for example, `domain.com`)

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

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

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

      - HYCU backup controller user
      - HYCU backup controller password

    Click **Next**.

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

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

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

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see "Licensing" on page 488.

Deploying HYCU to a XenServer environment

The HYCU virtual appliance is distributed as a virtual machine disk image file (a `.vmdk` file) that you can easily deploy to a XenServer environment by using XenCenter.

Prerequisite

The HYCU virtual machine disk image must be saved on your system.

Deployment tasks

When deploying HYCU to a XenServer environment, you must perform the following tasks:

Task	Instructions
1. Create a virtual machine for HYCU deployment from the virtual machine disk image.	"Creating a virtual machine for HYCU deployment in a XenServer environment" on the next page
2. Attach a new disk to the created virtual machine.	"Attaching a disk to the created virtual machine" on the next page
3. Configure HYCU on the created virtual machine.	"Configuring HYCU on the created virtual machine" on page 54

Creating a virtual machine for HYCU deployment in a XenServer environment

Procedure

1. Sign in to XenCenter.
2. Click **File**, and then click **Import...**.
3. Click **Browse**, search for the `.vmdk` file on your system, and then click **Next**.
4. Select the **Rocky Linux 8** virtual machine template, and then click **Next**.
5. Enter a name for the virtual machine, and then click **Next**.
6. Allocate the virtual CPU and memory resources to the virtual machine. For details on the system requirements, see "Sizing resources for your HYCU backup infrastructure" on page 27.
7. Click **Next**.
8. Specify the location for the virtual machine, and then click **Next**.
9. Select the **Place all imported virtual disks on this target SR** option, and then specify the local or shared storage repository where you want to place all imported virtual disks.
10. Click **Next**.
11. Configure networking for the virtual machine, and then click **Next**.
12. Select the **BIOS Boot** mode for the virtual machine, and then click **Next**.
13. Select the **Don't use Operating System Fixup** option, and then click **Next**.
14. Review all the settings, and then clear the **Start the new VM(s) automatically as soon as the import is complete** check box.
15. Click **Finish**.

The import operation may take some time to complete. To check the progress of the import, click the Logs tab.

Attaching a disk to the created virtual machine

Procedure

1. Select the virtual machine.
2. Click the **Storage** tab, and then click **Add...**.

3. Enter a name for the new virtual disk and, optionally, its description.
4. Specify the size for the new virtual disk. For details on the system requirements, see “Sizing resources for your HYCU backup infrastructure” on page 27.
5. Select the location where the new virtual disk will be stored. Make sure that you specify the same storage repository as for the imported virtual disks.
6. Click **Add**.

Configuring HYCU on the created virtual machine

Procedure

1. From the list of virtual machines in XenCenter, select the one you created and start it.
2. In the HYCU Mode Selection dialog box that opens, select one of the following deployment modes:
   - **HYCU Backup Controller**
   - **HYCU Instance**
   - **HYCU Manager**

   For details on deployment modes, see “Deployment modes” on page 23.
3. Tab to **OK** and press **Enter**.
4. In the Network Configuration dialog box that opens, do the following:
   a. Enter the values for the following:
      - *Optional.* Host name for the virtual machine

        The default host name is generated automatically during the HYCU deployment. If you want to use a custom host name, keep in mind the following:
        - *Only if you selected the HYCU Backup Controller or HYCU Manager mode.* The host name should begin with a letter and may contain only letters, numbers, and hyphens (-).
        - *Only if you selected the HYCU Instance mode.* For host name naming conventions, see “Managing HYCU instances” on page 484.
      - IPv4 address (for example, `10.1.100.1`)
      - Subnet mask (for example, `255.0.0.0`)

- Default gateway (for example, `10.1.1.1`)
- *Optional.* DNS server (for example, `10.1.1.5`)
- *Optional.* Search domain (for example, `domain.com`)

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

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

The progress of the HYCU configuration displays.

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

> **Important** If the HYCU backup controller host name cannot be resolved from the HYCU instance (for example, in environments that do not use DNS servers), make sure to use an IP address in the following format:
> `https://<IPAddress>:<Port>`

The progress of the HYCU backup controller assignment appears.

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

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see "Licensing" on page 488.

Deploying HYCU to an Azure Local environment

The HYCU virtual appliance is distributed as a virtual machine image that you can deploy to an Azure Local environment by using Windows Admin Center.

Limitation

HYCU does not support UEFI Secure Boot.

Deployment tasks

When deploying HYCU to an Azure Local environment, you must perform the following tasks:

Task	Instructions
1. Upload the HYCU virtual machine image to an Azure Local environment.	"Uploading the HYCU virtual machine image to an Azure Local environment" below
2. Create a virtual machine for HYCU deployment.	"Creating a virtual machine for HYCU deployment in an Azure Local environment" below
3. Configure HYCU on the created virtual machine.	"Configuring HYCU on the virtual machine" on the next page

Uploading the HYCU virtual machine image to an Azure Local environment

Procedure

1. Sign in to Windows Admin Center for your Azure Local environment.
2. Go to **Volumes**, and then click the **Inventory** tab.
3. Select the volume to which you want to upload the HYCU virtual machine image.
4. Click **New Folder** to create a new folder named `HycuImages`.
5. Click **Upload** to upload and store the HYCU virtual machine image.

Creating a virtual machine for HYCU deployment in an Azure Local environment

Procedure

1. In Windows Admin Center, go to **Virtual machines**, and then click the **Inventory** tab.
2. Click **Add**, and then click **New**.
3. Enter a name for your virtual machine.
4. From the Generation drop-down menu, select **Generation 1**.

5. Configure the virtual processors, the memory, and the network settings. For instructions, see Azure Local documentation.
6. Click **Create**.
7. Attach a system disk and a data disk to the virtual machine. To do so, follow these steps:
   a. Select the virtual machine that you created, and then click **Settings**.
   b. Select **Disks**, and then do the following:
      i. Click **Add disk**, and then specify a system disk:
         I. Click **New Virtual Hard Disk**, and then select **Copy a virtual hard disk**.
         II. Click **Browse** and select the path to the HYCU virtual machine image that you uploaded.
         III. Click **OK**.
         IV. Click **Create**.
      ii. Click **Add disk**, and then specify a data disk:
         I. Click **New Virtual Hard Disk**, and then select **Create an empty virtual hard disk.**
         II. Click **Create**.

         > Note The minimum size for the data disk is 128 GiB.

   c. Click **Save disk settings**, and then wait until the Settings page for your virtual machine reopens.
   d. Click **Close**.

Configuring HYCU on the virtual machine

Procedure

1. Use a remote desktop connection to connect to the Failover Cluster Manager of the Azure Local cluster where the virtual machine that you created is located.
2. Under Roles, access the console of your virtual machine.
3. In the HYCU Mode Selection dialog box that opens, select one of the following deployment modes:
   - **HYCU Backup Controller**
   - **HYCU Instance**

- **HYCU Manager**

For details on deployment modes, see "Deployment modes" on page 23.

4. Tab to **OK**, and then press **Enter**.
5. Specify the values for the following:
   - *Optional.* Host name for the virtual machine

     The default host name is generated automatically during the HYCU virtual appliance deployment. If you want to use a custom host name, keep in mind the following:
     - *Only if you selected the HYCU Backup Controller or HYCU Manager mode.* The host name should begin with a letter and may contain only letters, numbers, and hyphens (-).
     - *Only if you selected the HYCU Instance mode.* Make sure that you enter a unique host name for each HYCU instance that you create and follow these rules:
       - The host name should contain only letters, numbers, hyphens (-), and periods. The maximum number of characters must be 253 and at least one of the characters must be a letter.
       - The maximum number of characters in each host name segment must be 63. A host name segment may not begin or end with a hyphen.
       - The top-level domain may not begin or end with a number.
   - 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 (-).

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

   > **Important** If the HYCU backup controller host name cannot be resolved from the HYCU instance (for example, in environments that do

> not use DNS servers), make sure to use the IP address:
> `https://<IPAddress>:<Port>`

The progress of the HYCU backup controller assignment is shown.

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

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see "Licensing" on page 488.

Deploying HYCU to a Hyper-V cluster

The HYCU virtual appliance is distributed as a virtual machine image that you can deploy to a Hyper-V cluster.

Limitation

HYCU does not support UEFI Secure Boot.

Recommendation

You can deploy HYCU to a Hyper-V cluster by using Failover Cluster Manager or Windows Admin Center. However, it is highly recommended that you use Failover Cluster Manager. In this section, you are guided through the steps that you must perform if you are using Failover Cluster Manager. For instructions on how to use Windows Admin Center, see "Deploying HYCU to an Azure Local environment" on page 55.

Deployment tasks

When deploying HYCU to a Hyper-V cluster, you must perform the following tasks:

Task	Instructions
1. Upload the HYCU virtual machine image to a Hyper-V cluster.	"Uploading the HYCU virtual machine image to a Hyper-V cluster" on the next page
2. Prepare a virtual machine for HYCU deployment.	"Creating a virtual machine for HYCU deployment" on page 61

Task	Instructions
	"Creating and attaching a new data disk to the virtual machine" on page 62
3. Configure HYCU on the created virtual machine.	"Configuring HYCU on the virtual machine" on page 62

Uploading the HYCU virtual machine image to a Hyper-V cluster

Procedure

1. Use a remote desktop connection to connect to the Hyper-V cluster node.
2. On the Hyper-V cluster node, open Server Manager, and then, under Tools, open Failover Cluster Manager.
3. In Failover Cluster Manager, expand the relevant cluster, click **Storage**, and then click **Disks**.
4. Locate the CSV (Cluster Shared Volume) in the `ClusterStorage` folder. You can identify the CSV by the volume number. For example, the CSV location is `C:\ClusterStorage\Volume1`.
5. Create a dedicated folder named `HYCU_Controller` in the CSV under the default location for the Hard Disks: `C:\ClusterStorage\Volume1\<DefaultHardDiskLocation>\HYCU_Controller`
6. Upload the HYCU virtual machine image file to the dedicated folder. This folder will serve as the custom storage location for the HYCU virtual machine.

Preparing a virtual machine for HYCU deployment on a Hyper-V cluster

When preparing a virtual machine for HYCU deployment on a Hyper-V cluster, you must perform the following tasks:

Task	Instructions
1. Create a virtual machine.	"Creating a virtual machine for HYCU

Task	Instructions
	deployment" on the next page
2. Create and attach a new data disk to the virtual machine.	"Creating and attaching a new data disk to the virtual machine" on the next page

Creating a virtual machine for HYCU deployment

Procedure

1. In the Failover Cluster Manager navigation pane, right-click **Roles**, and then under Virtual Machines, click **New Virtual Machine**.
2. Select the target cluster node on which you want to create the virtual machine. The new virtual machine wizard opens.
3. In the Specify Name and Location step, do the following:
   a. Enter a name for your HYCU virtual machine.
   b. Enable the **Store the virtual machine in a different location** option, and the select the location for the HYCU virtual machine image.
   c. Click **Next**.
4. From the Generation drop-down menu, select **Generation 1**, and then click **Next**.
5. Define the virtual machine startup memory size. The minimum required size is 8 GiB.
6. Click **Next**.
7. Select an external switch to connect to the network, and then click **Next**.
8. Select the **Use an existing virtual hard disk** option, and then enter the location of the `HYCU_Controller` folder as you specified in "Uploading the HYCU virtual machine image to a Hyper-V cluster" on the previous page.
9. Click **Next**.
10. Verify the configuration.
11. Click **Finish**.

The HYCU virtual machine is created as a new Role and is by default powered off.

Creating and attaching a new data disk to the virtual machine

Procedure

1. In the Failover Cluster Manager navigation pane, click **Roles**.
2. From the listed roles, select the HYCU virtual machine, and then, in the Action pane, click **Settings**.
3. In the left pane of the virtual machine settings dialog box, under the name of the virtual machine, expand **IDE Controller 1**, and then click **Hard Drive**.
4. Click **Add**, and then click **New**. The new virtual hard disk wizard opens.
   a. In the Choose Disk Format step, select **VHD**, and then click **Next**.
   b. In the Choose Disk Type step, select **Fixed size**, and then click **Next**.
   c. In the Specify Name and Location step, enter the data disk name and location. When specifying the data disk location, it is recommended to select the same location as for the uploaded HYCU image `.vhd` file. For the uploaded `.vhd` file location details, see "Uploading the HYCU virtual machine image to a Hyper-V cluster" on page 60. Click **Next**.
   d. In the Configure Disk step, select **Create a new blank virtual hard disk**. The minimum virtual hard disk size is 128 GB.
   e. Verify the configuration and click **Finish**. After the new hard disk is created, you will be redirected to the virtual machine settings dialog box. The new data disk is listed under IDE Controller 1.
5. Click **Apply**, and the then click **OK**.

Configuring HYCU on the virtual machine

Procedure

1. Use a remote desktop connection to connect to the Failover Cluster Manager of the Hyper-V cluster where the virtual machine that you created is located.
2. Under Roles, access the console of your virtual machine.
3. In the HYCU Mode Selection dialog box that opens, select one of the following deployment modes:
   - **HYCU Backup Controller**
   - **HYCU Instance**
   - **HYCU Manager**

For details on deployment modes, see "Deployment modes" on page 23.

4. Tab to **OK**, and then press **Enter**.
5. Specify the values for the following:
   - *Optional.* Host name for the virtual machine

     The default host name is generated automatically during the HYCU virtual appliance deployment. If you want to use a custom host name, keep in mind the following:
     - *Only if you selected the HYCU Backup Controller or HYCU Manager mode.* The host name should begin with a letter and may contain only letters, numbers, and hyphens (-).
     - *Only if you selected the HYCU Instance mode.* Make sure that you enter a unique host name for each HYCU instance that you create and follow these rules:
       - The host name should contain only letters, numbers, hyphens (-), and periods. The maximum number of characters must be 253 and at least one of the characters must be a letter.
       - The maximum number of characters in each host name segment must be 63. A host name segment may not begin or end with a hyphen.
       - The top-level domain may not begin or end with a number.
   - 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 (-).
6. Tab to **OK**, and then press **Enter**.
7. *Only if deploying HYCU in the HYCU Instance mode.* In the HYCU Backup Controller dialog box that opens, enter the HYCU backup controller URL and the user name and password that you use to access HYCU.

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

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

The progress of the HYCU backup controller assignment is shown.

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

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see "Licensing" on page 488.

Deploying HYCU to an AWS GovCloud (US) environment

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

Prerequisite

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

Limitations

- You can deploy HYCU only in the HYCU Backup Controller mode.
- HYCU does not support UEFI Secure Boot.

Consideration

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

Procedure

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

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

> ⓘ Important Make sure that you use the following format for the HYCU virtual appliance image:
> `hycu-<Version>-<Revision>`
>
> For example: `hycu-5.3.0-363`.
>
> If you enter the HYCU image name in a different format, you will not be able to use this image for an upgrade.

4. Navigate to **EC2**, select **Instances**, and then click **Launch Instances**.
5. On the Choose AMI page, select the HYCU AMI that you created, and then click **Select**.
6. Follow the wizard to create a HYCU backup controller virtual machine from the AMI, making sure that you do the following:
   a. Add a new volume with the size of at least 128 GiB.
   b. Create a new firewall rule to allow ingress network traffic on TCP port 8443 from the entire subnetwork to which the HYCU backup controller belongs.

   You can leave the default values for the remaining options or adjust them to your needs. For details on all options, see AWS documentation.
7. Review the virtual machine details, and then click **Launch**.
8. Specify a key pair that will be used to connect to your HYCU backup controller. Click **Launch Instances**.

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see “Licensing” on page 488.

Deploying HYCU to an Azure environment

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

Prerequisite

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

Limitations

- You can deploy HYCU in the HYCU Backup Controller mode. If you want to create a HYCU instance or HYCU Manager virtual machine in Azure, contact HYCU Support.
- HYCU does not support UEFI Secure Boot.

Consideration

After deploying HYCU, a warning message stating that the virtual machine agent status is not ready may be displayed in Azure. You can safely ignore this message.

Recommendation

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

Procedure

1. Sign in to Azure.
2. Create a managed image from the HYCU virtual appliance image:
   a. In the Images navigation pane, click **Create**. In the Create an image menu that opens, make sure you specify the following:
      - In the Instance details section, in the Name field, enter the name of the HYCU virtual appliance image in the following format:

        `hycu-<Version>-<Revision>`

        For example, `hycu-5.3.0-3634`.
      - In the OS disk section, select the following:
        - OS type: **Linux**
        - VM Generation: **Gen 1**

      You can leave the default values for the remaining options, or adjust them to your needs.
   b. Click **Review + Create** to review the information, and then click **Create** to create the managed image.
3. Create a virtual machine from the managed image. Make sure the virtual machine is configured with an additional disk of 128 GiB in size. For instructions, see Azure documentation.
4. *Only if you use a network security group.* Create a new firewall rule to allow ingress network traffic on TCP port 8443 from the entire subnetwork to

which the HYCU backup controller belongs. For details, see Azure documentation.

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see “Licensing” on page 488.

Deploying HYCU to an Azure Government environment

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

Prerequisite

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

Limitations

- You can deploy HYCU in the HYCU Backup Controller mode. If you want to create a HYCU instance or HYCU Manager virtual machine in Azure Government, contact HYCU Support.
- HYCU does not support UEFI Secure Boot.

Considerations

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

Recommendation

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

Procedure

1. Sign in to Azure Government.
2. Create a managed image from the HYCU virtual appliance image:
   a. In the Images navigation pane, click **Create**. In the Create an image menu that opens, make sure you specify the following:
      - In the Instance details section, in the Name field, enter the name of the HYCU virtual appliance image in the following format:

        `hycu-<Version>-<Revision>`

        For example, hycu-5.3.0-3634.
      - In the OS disk section, select the following:
        - OS type: **Linux**
        - VM Generation: **Gen 1**

      You can leave the default values for the remaining options, or adjust them to your needs.
   b. Click **Review + Create** to review the information, and then click **Create** to create the managed image.
3. Create a virtual machine from the managed image. Make sure the virtual machine is configured with an additional disk of 128 GiB in size. For instructions, see Azure documentation.
4. *Only if you use a network security group.* Create a new firewall rule to allow ingress network traffic on TCP port 8443 from the entire subnetwork to which the HYCU backup controller belongs. For details, see Azure documentation.

You can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period. For instructions, see “Licensing” on page 488.

Signing in to HYCU

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

Consideration

If you enter your credentials incorrectly too many times when signing in to HYCU, your account will be locked for a short period of time.

Procedure

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

Environments	URL to enter

• Nutanix AHV
• Nutanix ESXi
• vSphere
• XenServer
• Azure Local
• Hyper-V
	`https://<ServerName>:8443`
	In this instance, `<ServerName>` is the fully qualified domain name of the HYCU server.
	For example:
	`https://hycu.example.com:8443`

• AWS GovCloud (US)[a] • Azure • Azure Government	`https://<IPAddress>:8443` In this instance, `<IPAddress>` is the external IP address of the HYCU backup controller.

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

2. On the sign-in page, depending on how you want to sign in to HYCU, do one of the following:
   - *By using dedicated sign-in credentials for HYCU.* Enter your sign-in user name and password.

     You can use the default user name (admin) and password (admin) for initial access to HYCU. For security purposes, it is highly recommended that you change the default password. To change the password, click at the upper right of the screen, and then select **Change Password**.
   - *By using an identity provider.* Depending on the type of your identity provider, do one of the following:
     - *For Active Directory:* Enter your Active Directory user name in the format `<UserName>@<Domain>` or `<Domain>\<UserName>` and your Active Directory password, and then click **Sign In**.
     - *For identity providers that support the OpenID Connect authentication protocol:* Click the preferred identity provider, and then, if required,

enter your credentials.

- *For identity providers that use LDAP:* Enter your LDAP user name and password, and then click **Sign In**. You must enter the user name in the format `<UserName>@<Domain>` or `<Domain>\<UserName>`, where <*Domain*> is the domain that was configured when the identity provider was added to HYCU.

For details on how to integrate HYCU with identity providers, see "Integrating HYCU with identity providers" on page 471.

3. *Only if two-factor authentication is enabled for your account.* Enter the appropriate two-factor credentials:
   - *For using time-based one-time passwords (OTP)*: Enter the six-digit authentication code generated by your authentication application (for example, Google Authenticator or a compatible application).

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

     When you sign in for the first time after two-factor authentication was enabled for your account, a security dialog box opens, prompting you to set up an authenticator (for example a security key or a fingerprint reader). The procedure depends on the selected authenticator and operating system. Follow the instructions to set up the authenticator. For details, see "Managing FIDO authenticators" on page 544.

> **Note** Keep in mind that the level of access depends on your user permissions. For details, see "Managing users" on page 440.

Setting the language

When you access the HYCU web user interface or the HYCU Manager console, the current browser language is detected and if it is one of the supported languages, the user interface is displayed in that language. If the browser

language is not one of the supported languages, the user interface is displayed in English. For a list of supported languages, see the *HYCU Compatibility Matrix*.

Consideration

The HYCU REST API Explorer is available only in English.

Procedures

- If you are an infrastructure or a self-service group administrator, you can set the preferred language for a user. For instructions, see "Creating a user" on page 445.
- You can set your preferred language by using the Update Profile option. For instructions, see "Updating your user profile" on page 459.
- You can set the preferred language for notifications that are sent when events occur. For instructions, see "Configuring event notifications" on page 388.

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

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

Chapter 3

Establishing a data protection environment

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

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

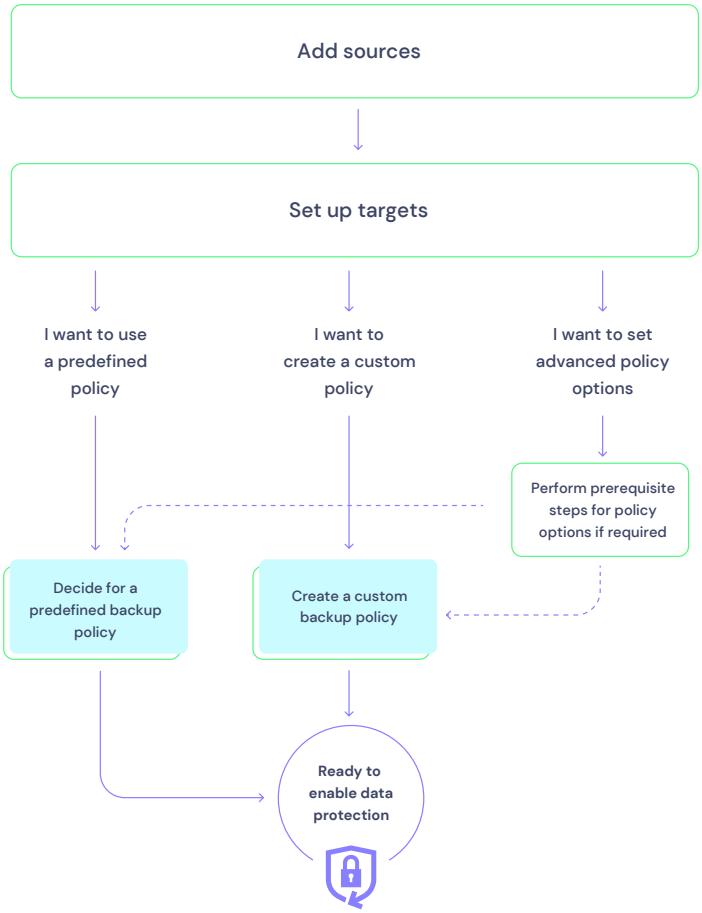

Figure 3-1: Establishing a data protection environment

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

- “Adding sources” below
- “Setting up targets” on page 100

You can enable data protection by using predefined policies that come with HYCU. If you do not want to use any of them, make sure that you create your own policies. For details, see “Creating a policy” on page 134.

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

> **Note** Before you start protecting your data protection environment, make sure that the HYCU backup controller is protected. This way, you can quickly recover and resume your data protection activities in case of a disaster. For more information, see “Preparing for disaster recovery of your data” on page 161.

Adding sources

An environment for which HYCU provides data protection consists of one or more sources that you add to HYCU depending on what kind of data you want to protect. For instructions on how to add a specific source, see one of the following sections:

I want to protect …	Source	Instructions
Virtual machines Applications running on virtual machines	Nutanix cluster	“Adding a Nutanix cluster” on the next page
	vCenter Server[a]	“Adding a vCenter Server” on page 77
	XenServer	“Adding a XenServer” on page 78
	Azure Local environment	“Adding an Azure Local environment” on page 80
	Hyper-V cluster	“Adding a Hyper-V cluster” on page 80

	AWS GovCloud (US) region Azure subscription Azure Government subscription	"Adding a cloud environment" on page 81
Servers[b] Applications running on servers	Server	"Adding a server" on page 99
File shares	File server	"Adding a file server" on page 85
Volume groups	Nutanix cluster	"Adding a Nutanix cluster" below
Buckets	Object server	"Adding an object server" on page 98

[a] Including virtual machine templates in vSphere environments.
[b] For details on protecting server data, see "Protecting virtual machines" on page 154.

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

Adding a Nutanix cluster

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

Prerequisites

- *For Nutanix ESXi clusters:*
  - Your cluster must be registered to the vCenter Server through the Prism web console. For details on how to do this, see Nutanix documentation.
  - A user with specific privileges for vCenter Servers must be specified. For details on which privileges must be assigned to a vSphere user, see "Assigning privileges to a vSphere user" on page 552.

- If you plan to perform any of the data protection tasks described in “Configuring Prism Central user permissions” on page 556, the Nutanix cluster that hosts the virtual machines must be registered with Prism Central and your Prism Central user must have a role with sufficient permissions assigned.

Considerations

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

Recommendations

- For better performance, it is recommended that an iSCSI Data Service IP address is specified on the Nutanix cluster that you plan to add to HYCU. This automatically enables the Nutanix load balancing feature during data protection operations, which eliminates heavy I/O load on the Nutanix cluster and storage containers. For details on how to specify an iSCSI Data Service IP address, see Nutanix documentation.
- It is recommended that you create and use a separate local service account with full administrator permissions to access your Nutanix cluster. By doing so, you will avoid potentially locking out the main admin account.

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

Procedure

1. In the Sources dialog box, click the **Hypervisor** tab, and then click **Add**.
2. Enter the Prism Element cluster host name or virtual IP address in the following format:

`https://<PrismElementClusterHostNameorVirtualIPAddress>:<Port>`

3. Enter the user name and password of a user with cluster administrative rights.

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

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

   > **Note** If you use Conjur for managing your HYCU secrets, you can enable the **Retrieve values from secrets manager** switch if you want to provide the secret instead of browsing for the file. For details on managing secrets, see "Managing secrets" on page 509.

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

5. Click **Next**.
6. *Only if you are adding a Nutanix ESXi cluster.* In the New vSphere Credentials dialog box, assign the vSphere credentials to the Nutanix ESXi cluster by specifying the URL of the vCenter Server to which the Nutanix ESXi cluster is registered, and the user name and password of a user with specific privileges for vCenter Servers. Click **Next**.

   > **Note** After you add a Nutanix ESXi cluster, the vC icon next to its type shows that it has the required vCenter Server permissions.

7. If you plan to perform any of the data protection tasks described in "Configuring Prism Central user permissions" on page 556, in the New Prism Central Credentials dialog box, specify the URL of Prism Central with which your Nutanix cluster is registered, and the user name and password of a Prism Central user that has a role with sufficient permissions assigned. Otherwise, leave all the fields blank. Click **Next**.
8. *Only if you use Nutanix Database Service (NDB).* In the New NDB Credentials dialog box, specify the URL of the NDB, and the user name and password of a user with administrative rights. Click **Next**.
9. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

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

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

Adding a vCenter Server

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

Prerequisite

A user with specific privileges for vCenter Servers must be specified. For details on which privileges must be assigned to a vSphere user, see “Assigning privileges to a vSphere user” on page 552.

Limitation

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

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

Procedure

1. In the Sources dialog box, click the **Hypervisor** tab, and then click **Add**.
2. Enter the vCenter Server host name or IP address in the following format:

   `https://<vCenterServerFQDNorIPAddress>:<Port>`

The default port for the vCenter Server is 443.

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

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

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

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

You can later do the following:

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

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

  3. Click **Remove**.

Adding a XenServer

A XenServer environment consists of a standalone XenServer or a XenServer pool that hosts virtual machines for which HYCU provides data protection. Adding a XenServer to HYCU is the first step to protecting your data.

Prerequisite

For pools: The XenServer that you add to HYCU must be a pool coordinator.

Considerations

For pools:

- HYCU automatically detects if the XenServer that you are adding is the pool coordinator. If you try to add a XenServer other than the pool coordinator to HYCU, HYCU automatically switches to adding the pool coordinator and informs you about the action.
- If the pool coordinator changes after you add it to HYCU, all the data will be transferred to the new pool coordinator if the original pool coordinator gets back online. Otherwise, you must update the XenServer information manually by editing the source in HYCU.

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

Procedure

1. In the Sources dialog box, click the **Hypervisor** tab, and then click **Add**.
2. Enter the XenServer host name or IP address in the following format:

   `https://<XenServerFQDNorIPAddress>:<Port>`

   The default port for the XenServer is 443.

   > **Important** Make sure you configure the HYCU DNS settings in a way that allows HYCU to resolve this FQDN and, consequently, connect to the XenServer on which the virtual machines that you want to protect are running.

3. Enter the user name and password of a user with specific privileges for the XenServer, and then click **Next**.
4. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

You can later do the following:

- Edit any of the existing XenServers. To do so, select the XenServer, click **Edit**, make the required modifications, and then click **Save**.
- Remove the XenServers that you do not need anymore as follows:

1. Select the XenServers that you want to remove, and then click **Remove**.
2. *Only if you want to delete snapshots created by HYCU.* Enable the **Delete snapshots** switch.

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

3. Click **Remove**.

Adding an Azure Local environment

Adding an Azure Local environment to HYCU is the first step to protecting your data.

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

Procedure

1. In the Sources dialog box, on the **Hypervisor** tab, click **Add**.
2. Enter the cluster name of your Azure Local environment.

   > **Note** The cluster name must include the fully qualified domain name (FQDN). To get the cluster name, go to Windows Admin Center, select **Cluster Manager** from the drop-down menu in the top navigation bar. From the cluster connections, copy the relevant cluster name.

3. Enter the user name and password of a user with the Administrator role for your cluster, and then click **Next**.
4. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

You can also edit any of the existing Azure Local environments (click **Edit** and make the required modifications) or remove the ones that you do not need anymore (click **Remove**).

Adding a Hyper-V cluster

Adding a Hyper-V cluster to HYCU is the first step to protecting your data.

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

Procedure

1. In the Sources dialog box, on the **Hypervisor** tab, click **Add**.
2. Enter the name of your Hyper-V cluster.

   > **Note** The cluster name must include the fully qualified domain name (FQDN). To get the cluster name, go to Windows Admin Center, select **Cluster Manager** from the drop-down menu in the top navigation bar. From the cluster connections, copy the relevant cluster name.

3. Enter the user name and password of a user with the Administrator role for your cluster, and then click **Next**.
4. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

You can also edit any of the existing Hyper-V cluster (click **Edit** and make the required modifications) or remove the ones that you do not need anymore (click **Remove**).

Adding a cloud environment

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

Cloud source	I plan to...	Instructions
AWS GovCloud (US) region	Protect virtual machines and applications in an AWS GovCloud (US) environment.	“Adding an AWS GovCloud (US) region” on the next page
Azure subscription	Protect virtual machines and applications in an Azure environment.	“Adding an Azure subscription” on page 83
Azure Government subscription	Protect virtual machines and applications in an Azure Government environment.	“Adding an Azure Government subscription” on page 84

Adding an AWS GovCloud (US) region

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

Prerequisites

- An AWS GovCloud (US) account must be added to HYCU. For instructions, see "Adding an AWS GovCloud (US) account" on page 465.
- Your user within the AWS GovCloud (US) account that you add to HYCU must have a policy with the required permissions attached. For a list of the required permissions, see "Setting permissions in AWS GovCloud (US)" on page 558.

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

Procedure

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

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

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

You can also do the following:

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

snapshots switch.

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

Adding an Azure subscription

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

Prerequisite

An Azure service principal must be added to HYCU. For instructions, see "Adding an Azure service principal" on page 468.

Limitation

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

> #### Accessing the Sources dialog box
>
> To access the Sources dialog box, click ⚙ **Administration**, and then select **Sources**.

Procedure

1. In the Sources dialog box, click the **Cloud** tab, and then click ✚ **Add**.
2. Select **Azure Subscription**, and then click **Next**.
3. From the Service Principal drop-down menu, select the service principal that has access to the virtual machines and applications that you want to protect.
4. The Application ID field is already filled in.
5. From the Subscription drop-down menu, select the Azure subscription that you want to add to HYCU, and then click **Next**.
6. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

You can also:

- View any of the existing Azure subscriptions. Click 👁 **View**.
- Remove the ones that you do not need anymore. Click 🗑 **Remove**.

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

Adding an Azure Government subscription

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

Prerequisite

An Azure Government service principal must be added to HYCU. For instructions, see "Adding an Azure Government service principal" on page 469.

Limitation

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

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

Procedure

1. In the Sources dialog box, click the **Cloud** tab, and then click **Add**.
2. Select **Azure Government Subscription**, and then click **Next**.
3. From the Service Principal drop-down menu, select the service principal that has access to the virtual machines and applications that you want to protect.
4. The Application ID field is automatically filled in.
5. From the Subscription drop-down menu, select the Azure Government Subscription that you want to add to HYCU, and then click **Next**.
6. In the Summary dialog box, verify that the validation was successful, and then click **Save**.

You can also:

- View any of the existing Azure subscriptions. Click **View**.
- Remove the ones that you do not need anymore. Click **Remove**.

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

Adding a file server

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

- The storage provider APIs enable HYCU to perform incremental backups.
- HYCU can utilize snapshot capabilities to create consistent backups.

For protecting file shares, a HYCU instance is required. For details, see "HYCU instances" on page 25.

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

- "Adding a Nutanix Files file server" below
- "Adding a Dell PowerScale OneFS file server" on page 87
- "Adding a NetApp ONTAP file server" on page 92
- "Adding an Azure Files file server" on page 96
- "Adding a generic file server" on page 97

Adding a Nutanix Files file server

Prerequisites

- HYCU must have access to the file server that you are adding. For details, see "Enabling HYCU to access a Nutanix Files server" on page 87.
- To ensure that HYCU can use Kerberos authentication, your DNS server must be configured to perform reverse DNS lookups. Otherwise, NTLM authentication will be used.

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

Procedure

1. In the Sources dialog box, click the **File Servers** tab, and then click **Add**.
2. Select **Nutanix Files Server**, and then click **Next**.
3. *Optional.* Enter a custom display name for the file server.
4. Enter the host name or IP address of the file server in the following format:

   `https://<Hostname/IP>:<Port>`

   Entering the port is optional if the default value is used (9440).

   > **Important** If you are providing the host name, make sure the name is unique.

5. Specify the user name and password of a user with administrative rights for REST API access on the file server.

   > **Important** To ensure that HYCU can use Kerberos authentication, the user name must include a fully qualified domain name. For example, `<UserName>@<DomainName>.<DomainSuffix>` or `<DomainName>.<DomainSuffix>\<UserName>`.

6. Click **Next**.
7. Enable the **Use SMB protocol for accessing shares** switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually.
8. Enable the **Use NFSv4/NFSv3 protocol for accessing shares** switch if you plan to protect NFS file shares.
9. Click **Save**.

You can later do the following:

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

> **Important** *Only if you plan to use a socket-based or a core-based File server license for protecting file shares.* The Nutanix cluster on which the file server that you are adding resides must also be added to HYCU as a source. For instructions, see "Adding a Nutanix cluster" on page 74.

Enabling HYCU to access a Nutanix Files server

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

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

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

1. Establish a connection to the Nutanix cluster:

```
ssh @<NutanixClusterHostname>
```

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

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

Adding a Dell PowerScale OneFS file server

Prerequisites

- HYCU must have access to the file server that you are adding. For details, see "Dell PowerScale OneFS user permissions" on page 91.
- *Only if you plan to protect SMB file shares:*
  - The SMB user must have the run as root permissions in each file share, and must also have the Backup and Restore privileges assigned. If the user belongs to an Active Directory domain, the user must be a member of the Backup Operators group.
  - The SMB `Dot Snap Accessible Child` and `Dot Snap Accessible Root` options must be set to `Yes`.
- *Only if you plan to protect NFS file shares by using the NFSv4 protocol.*
  - Root squashing must be disabled on the file server.
  - If you plan to protect only the data within the exported file share subdirectory (and not the root directory), the Enable mount access to subdirectories option must be enabled on the file server.
- To manage multiple access zones as a single source, the file shares in all access zones must be accessible by using a common SMB account.

Considerations

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

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

- *Only if you plan to protect SMB file shares.* As part of adding a file server, HYCU tests the SMB credentials using a random SMB file share on the server. If a user does not have permissions for all file shares, the test may fail. The failed test causes the file server adding process to end with a reported error. You can skip the test by changing the `afs.skip.smb.test` configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.

Accessing the Sources dialog box

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

Procedure

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

I want to...	Instructions
Use customized network and authentication	a. Make sure the **Manage multiple zones** switch is disabled. b. *Optional.* Enter a custom display name for the file

I want to...	Instructions
configurations for each access zone.	server. c. In the Zone URL field, enter the host name or IP address of the Dell PowerScale OneFS access zone in the following format: `https://<HostName/IP>:<Port>` Entering the port is optional if the default value is used (8080). > ⓘ Important When providing the host name, make sure to consider the following: > • If you are providing the host name, the name must be unique. > • To ensure that HYCU can use Kerberos authentication, the host name of the file server must include the Kerberos domain. For example, `https://<HostName>.<DomainName>.<DomainSuffix>:<Port>`. d. Specify the user name and password of a user with administrative rights for REST API access on the file server, and then click **Next**. > ⓘ Important To ensure that HYCU can use Kerberos authentication, the user name must include a fully qualified domain name. For example, `<UserName>@<DomainName>.<DomainSuffix>` or `<DomainName>.<DomainSuffix>\<UserName>`. e. *Only if the current access zone is not the System zone.* Under System zone credentials, enter the System zone URL, and the System zone user name and password. f. Enable the **Use SMB protocol for accessing shares** switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup

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

4. Click **Save**.

You can later do the following:

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

Dell PowerScale OneFS user permissions

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

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

Adding a NetApp ONTAP file server

Prerequisites

- HYCU must have access to the file server that you are adding. For details, see "NetApp ONTAP user permissions" on page 96.
- To ensure that HYCU can use Kerberos authentication, your DNS server must be configured to perform reverse DNS lookups. Otherwise, NTLM authentication will be used.
- *Only if you plan to protect SMB file shares.* The SMB user must have full read and write access to all the file shares that you plan to back up. If the user belongs to an Active Directory domain, the user must be a member of the Backup Operators group.
- *For NetApp ONTAP version 9.10.1 or later:* Client access to the Snapshot copy directory on volumes and shares must be enabled. For details, see NetApp ONTAP documentation.
- *Only if you plan to protect NFS file shares using the NFSv4 protocol.*
  - *For generic NFS file shares:* To be able to list exports for the NFSv4 servers, generic file shares must also support the NFSv3 protocol.

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

  - Root squashing must be disabled on the file server.
  - The export policy for the file shares that are going to be backed up must include a rule that allows super user access from the HYCU backup controller, and HYCU instance subnets or IP addresses.
- *Only if you plan to enable the NetApp SnapDiff feature for faster incremental backups of file shares.*
  - The version of NetApp ONTAP on your file server must be 9.8 or later.
  - SnapDiff v3 must be enabled on the storage virtual machine (SVM) on the NetApp ONTAP file server. For instructions, see SnapDiff Support in ONTAP.
  - Communication between HYCU and NetApp ONTAP is performed through the SVM's data LIF by using NFS RPC calls. Therefore, make sure that the LIF has NFS enabled in its policy and firewall rules. For instructions, see NetApp ONTAP documentation.

Considerations

- When adding a NetApp ONTAP file server, you can configure HYCU to manage the source in the following ways:
  - Manage all SVMs on your NetApp ONTAP file server as a single source. This is the basic option for environments where all the SVMs use the same network and authentication configuration.
  - Manage individual SVMs on your NetApp ONTAP file server as separate sources. This is the advanced option for environments where each SVM uses a customized network and authentication configuration.
- *Only if you plan to protect SMB file shares.* As part of adding a file server, HYCU tests the SMB credentials using a random SMB file share on the server. If a user does not have permissions for all file shares, the test may fail. The failed test causes the file server adding process to end with a reported error. You can skip the test by changing the `afs.skip.smb.test` configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.

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

Procedure

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

I want to...	Instructions
Use the same network and authentication configuration for all SVMs.	a. Make sure that the **Manage all SVMs** switch is enabled. b. *Optional.* Enter a custom display name for the file server. c. In the URL field, enter the host name or IP address of the NetApp ONTAP file server in the following format:

I want to...	Instructions
	`https://<Hostname/IP>:<Port>` Entering the port is optional if the default value is used (8080). > ⓘ Important If you enter the host name, make sure the name is unique. d. Specify the user name and password of a user with administrative rights for REST API access on the file server, and then click **Next**. > ⓘ Important To ensure that HYCU can use Kerberos authentication, the user name must include a fully qualified domain name. For example, `<UserName>@<DomainName>.<DomainSuffix>` or `<DomainName>.<DomainSuffix>\<UserName>`. e. Use the **Enable NetApp SnapDiff for incremental backups** switch if you want HYCU to utilize the NetApp SnapDiff feature to identify which files were modified between file share snapshots, resulting in faster incremental backups of file shares. f. Enable the **Use SMB protocol for accessing shares** switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually. g. Enable the **Use NFSv4/NFSv3 protocol for accessing shares** switch if you plan to protect NFS file shares.
Use customized network and authentication	a. Disable the **Manage all SVMs** switch. b. *Optional.* Enter a custom display name for the file

I want to...	Instructions

configurations for each SVM.	server.
	c. In the URL field, enter the host name or IP address of the NetApp ONTAP file server in the following format:
	`https://<Hostname/IP>:<Port>`
	Entering the port is optional if the default value is used (8080).
	> ⓘ Important If you enter the host name, make sure the name is unique.
	d. Specify the user name and password of a user with administrative rights for REST API access on the file server, and then click **Next**.
	> ⓘ Important To ensure that HYCU can use Kerberos authentication, the user name must include a fully qualified domain name. For example,
	`<UserName>@<DomainName>.<DomainSuffix>`
	or
	`<DomainName>.<DomainSuffix>\<UserName>.`
	e. From the Managed SVM drop-down menu, select the SVM that hosts the file shares that you want to protect.
	f. *Optional.* If you want to filter out certain SMB/NFS IP addresses due to your network environment needs, add the trusted IP addresses to the SVM IP whitelist. Enter a space after each IP address to add multiple IP addresses.
	g. Use the **Enable NetApp SnapDiff for incremental backups** switch if you want HYCU to utilize the NetApp SnapDiff feature to identify which files were modified between file share snapshots, resulting in faster incremental backups of file shares.

I want to...	Instructions
	h. Enable the **Use SMB protocol for accessing shares** switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually. i. Enable the **Use NFSv4/NFSv3 protocol for accessing shares** switch if you plan to protect NFS file shares.

4. Click **Save**.

You can later do the following:

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

NetApp ONTAP user permissions

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

Endpoint	Permissions
`/api/cluster`	Read/create
`/api/network/ip/interfaces`	Read-only
`/api/protocols/cifs/shares`	Read-only
`/api/storage/volumes`	Read/create/modify
`/api/svm/svms`	Read-only

Adding an Azure Files file server

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

Procedure

1. In the Sources dialog box, click the **File Servers** tab, and then click **Add**.
2. Select **Azure Files Server**, and then click **Next**.
3. *Optional.* Enter a custom display name for the file server.
4. In the Account Name field, enter the name of your Azure storage account.
5. In the Account Key field, enter your storage account access key.
6. Click **Save**.

You can later do the following:

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

Adding a generic file server

Prerequisites

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

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

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

Procedure

1. In the Sources dialog box, click the **File Servers** tab, and then click **Add**.
2. Select **Generic File Server**, and then click **Next**.
3. *Optional.* Enter a custom display name for the file server.

4. Enter the host name of the file server. Do not include `https://` or a port number in the host name.

   > **Important** To ensure that HYCU can use Kerberos authentication, the host name of the file server must include the Kerberos domain. For example, *<HostName>.<DomainName>.<DomainSuffix>*.

5. Enable the **Use SMB protocol for accessing shares** switch if you plan to protect SMB file shares. Enter the user name and password of a server or backup administrator with access to all SMB file shares within the file server. Keep in mind that you cannot assign credentials to each share individually.

   > **Important** To ensure that HYCU can use Kerberos authentication, the user name must include a fully qualified domain name. For example, *<UserName>@<DomainName>.<DomainSuffix>* or *<DomainName>.<DomainSuffix>\<UserName>*.

6. Enable the **SMB advanced settings** switch to specify a custom SMB port number. By default, the SMB port number is 445. If the destination port of the SMB server differs from the default one, you can specify an alternate port. Valid port numbers are from 0 through 65535.
7. Enable the **Use NFSv4/NFSv3 protocol for accessing shares** switch if you plan to protect NFS file shares.
8. Click **Save**.

You can later do the following:

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

Adding an object server

HYCU enables you to protect objects in buckets. You can add one or more object servers that host the buckets that you want to back up. For information on supported object server types, see the *HYCU Compatibility Matrix*.

For protecting buckets, a HYCU instance is required. For details, see "HYCU instances" on page 25.

Prerequisite

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

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

Procedure

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

   > **Important** Using the Path style access option to access the Cloudian buckets is mandatory.

6. Click **Save**.

You can later do the following:

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

Adding a server

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

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

Procedure

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

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

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

Setting up targets

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

Backup target type	Description
Target	Backup data can be stored on the following types of targets: NFS, SMB, Nutanix, Nutanix Objects, iSCSI, Azure (including Azure Government and Azure China), Amazon S3 / S3 Compatible (including AWS Government East and AWS Government West), Google Cloud, tape, and Data Domain. **Note** A file server can be used as an NFS or SMB target. If you plan to use a file server only as a target and not as a source, there is no need to add it to HYCU. The approach to setting up targets is common for different target types. However, there are specific prerequisites and steps that are required for each target type. Depending on which target you want to set up, see one of the following sections:

Backup target type	Description
	• "Setting up an NFS target" below • "Setting up an SMB target" on page 104 • "Setting up a Nutanix target" on page 107 • "Setting up a Nutanix Objects target" on page 109 • "Setting up an iSCSI target" on page 112 • "Setting up an Azure target" on page 115 • "Setting up an Amazon S3 / S3 Compatible target" on page 117 • "Setting up a Google Cloud target" on page 121 • "Setting up a tape target" on page 124 • "Setting up a Data Domain target" on page 128
Snapshot	*Not available for vSphere virtual machines residing on VMFS or NFS datastores.* Backup data is stored as a snapshot on the original location. **Important** If snapshots created by HYCU are corrupted or unavailable due to a disaster occurring in your data protection environment, you will not be able to restore backup data from this location. However, you can still restore your data from targets if data archives exist.

Setting up an NFS target

HYCU supports storing data to the file shares accessible by using the NFS protocol.

Prerequisites

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

- If the target resides on Windows, local permissions (security) must be set to `Full Control` for `Everyone`. If you want to limit access to this system only for HYCU, use the HYCU backup controller IP address for this purpose.
- The supported NFS version must be used. For a list of supported versions, see the *HYCU Compatibility Matrix*.
- *For protecting server data:* The target must be accessible from the server.

Limitations

For protecting server data:

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

Recommendation

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

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

Procedure

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

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

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

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

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

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

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

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

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

    > **Note** If you enable target encryption, keep in mind the following:
    > - The deduplication ratio may be affected by it (in cases where the target has deduplication enabled).
    > - To be able to import the encrypted target for restoring virtual machines, applications, file shares, buckets, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see "Exporting an encryption key" on page 471.

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

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

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

13. Click **Save**.

The target is added to the list of targets. For details on managing targets, see "Managing targets" on page 417.

Setting up an SMB target

HYCU supports storing data to the file shares accessible by using the SMB protocol.

Prerequisites

- The service must be configured and accessible for the HYCU backup controller and the HYCU instances.
- There must be enough free space on the target for storing the data.
- If deduplication is enabled on the target, the target must be dedicated exclusively to HYCU backups. By dedicating a target exclusively to HYCU backups, you ensure that accurate storage utilization reports are provided.
- The supported SMB version must be used. For a list of supported versions, see the *HYCU Compatibility Matrix*.
- *For protecting server data:* The target must be accessible from the server.

Limitations

For protecting server data:

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

Recommendation

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

Accessing the Targets panel

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

Procedure

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

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

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

   > **Important** Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, archiving of compressed data to targets with enabled compression may increase system requirements for the HYCU backup controller.
7. Click **Next**.
8. *Optional.* In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB). If your target is not dedicated exclusively to HYCU backups, you must leave this field empty.

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

   > **Note** If the target has deduplication enabled, HYCU's estimation of required storage space on the target may be higher than the actual

amount of space required on the storage media. Therefore, it is recommended to leave this field empty in such cases.

9. *Optional.* Enter the domain and user credentials.

   Important To ensure that HYCU can use Kerberos authentication, you must enter a DNS domain name. The legacy NetBIOS domain names are not supported by Kerberos.

10. Enter the SMB server host.

    Important To ensure that HYCU can use Kerberos authentication, you must enter the fully qualified domain name (FQDN).

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

    Note If you enable target encryption, keep in mind the following:
    - The deduplication ratio may be affected by it (in cases where the target has deduplication enabled).
    - To be able to import the encrypted target for restoring virtual machines, applications, file shares, buckets, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see "Exporting an encryption key" on page 471.

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

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

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

14. Click **Save**.

The target is added to the list of targets. For details on managing targets, see “Managing targets” on page 417.

Setting up a Nutanix target

HYCU supports using a Nutanix target for high-performance and resilient data storage on a Nutanix cluster.

For details on Nutanix storage, see Nutanix documentation.

Prerequisites

- The Nutanix cluster on which a Nutanix target will be created must be accessible to the HYCU backup controller.
- The Nutanix cluster must run a supported Nutanix AOS version. For a list of supported versions, see the *HYCU Compatibility Matrix*.

Limitations

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

Consideration

The storage container on a Nutanix cluster that HYCU creates automatically and uses as a Nutanix target must be dedicated exclusively to storing backup data. Because the names of such storage containers start with the `HYCU-` prefix, make sure not to create your own storage containers with the same prefix. Keep in mind that these storage containers are not available as destinations when restoring data, cloning data, and creating HYCU instances.

Recommendation

For better performance, it is recommended that an iSCSI Data Service IP address is specified on the Nutanix cluster on which a Nutanix target will be created. This automatically enables the Nutanix load balancing feature during data protection operations, which eliminates heavy I/O load on the Nutanix cluster and storage containers. For details on how to specify an iSCSI Data Service IP address, see Nutanix documentation.

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

Procedure

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

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

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

   > **Important** Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, archiving of compressed data to targets with enabled compression may increase system requirements for the HYCU backup controller.
7. Click **Next**.
8. *Optional.* In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).

   If you leave this field empty, HYCU retrieves the available amount of storage space from the target itself.
9. Enter the name of the Nutanix cluster in the following URL format:

   `https://<ServerName>:<Port>`
10. Enter the user name and password of a user with cluster administration rights.

    > **Important** When adding a Nutanix cluster that has client authentication enabled, make sure that you specify credentials of a local

user.

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

    For more information on these options, see Nutanix documentation.
12. Use the **Target encryption** switch if you want the data stored on this target to be encrypted.

    > **Note** If you enable target encryption, keep in mind the following:
    > - Enabling target encryption in combination with options intended to increase your cluster's effective storage capacity will prevent such options from taking effect.
    > - To be able to import the encrypted target for restoring virtual machines, applications, and volume groups, export the encryption key to a file and keep this file on safe. For instructions, see "Exporting an encryption key" on page 471.

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

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

The target is added to the list of targets. For details on managing targets, see "Managing targets" on page 417.

Setting up a Nutanix Objects target

HYCU supports using a Nutanix Objects target to store data in a secure S3 compatible object storage on top of the Nutanix AOS platform.

For details on Nutanix Objects cloud storage, see Nutanix Objects documentation.

Prerequisites

- The service must be configured and accessible.
- If you want to provide secure HTTPS access, the CA certificate/chain must be imported to HYCU. For details, see "Importing a custom certificate" on page 527.
- The supported Nutanix Objects version must be used. For a list of supported versions, see the *HYCU Compatibility Matrix*.

Limitations

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

Consideration

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

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

Procedure

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

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

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

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

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

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

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

9. Provide the following information:

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

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

    > 📋 Note To be able to import the encrypted target for restoring virtual machines, applications, file shares, volume groups, and buckets, export the encryption key to a file and keep this file safe. For instructions, see “Exporting an encryption key” on page 471.

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

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

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

13. Click **Save**.

The target is added to the list of targets. For details on managing targets, see "Managing targets" on page 417.

Setting up an iSCSI target

HYCU supports using an iSCSI target to store data on the iSCSI block-level storage devices.

Prerequisites

- The service must be configured and accessible.
- Make sure that the target is not initialized.
- The HYCU iSCSI initiator secret must be added on the iSCSI server if you want to enable mutual authentication between HYCU and the iSCSI server.

Limitations

- Using the same iSCSI target to store data protected by more than one HYCU backup controller is not supported.
- An iSCSI target cannot be used for storing file share and bucket data.
- Storing server backup data on this type of target is not supported.

Considerations

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

Recommendation

If you set up a Nutanix volume group that resides on a container with deduplication, compression, or erasure coding enabled as an iSCSI target, the HYCU target used space and the Nutanix storage container used space will not match. Therefore, to avoid this, it is recommended that you set up a Nutanix target with the Deduplication, Erasure coding, and Hardware compression options enabled instead. For instructions on how to set up a Nutanix target, see "Setting up a Nutanix target" on page 107.

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

Procedure

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

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

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

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

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

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

9. Enter the target portal IP address and the target name.

   > **Note** If data from sources other than HYCU resides on the storage device, such a target cannot be set for HYCU backups.

10. If the iSCSI server requires CHAP authentication, in the CHAP section, do the following:

    a. Use the **CHAP** switch to enable CHAP authentication, and then provide a user name and the target secret (the security key) for the user's account to access the iSCSI server.

    b. Use the **Perform mutual authentication** switch if you want the iSCSI target to be authenticated by HYCU. In this case, the HYCU iSCSI initiator secret must be specified on the iSCSI server. For details about setting the iSCSI initiator secret, see "Setting the iSCSI initiator secret" on page 487.

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

    > **Important** To be able to import the encrypted target for restoring virtual machines, applications, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see "Exporting an encryption key" on page 471.

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

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

13. Click **Save**.

The target is added to the list of targets. For details on managing targets, see "Managing targets" on page 417.

Setting up an Azure target

HYCU supports using an Azure, Azure Government, or Azure China target to store data in the highly available, scalable, and secure Azure cloud storage.

For details on Azure storage, see Azure documentation.

Prerequisites

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

Limitations

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

Considerations

- Your data on the Azure target can be stored in the hot, cool, cold, or archive access tier. When restoring data archives, HYCU performs data rehydration during which the tier of the blobs is changed from the archive access tier to the hot access tier. Keep in mind that this can take a few hours to complete. HYCU moves data back to the archive access tier afterward.
- HYCU automatically moves a data archive that has a retention period set to at least 180 days from the Azure hot, cool, or cold access tier to the archive access tier during the next archive synchronization. By moving data archives to the archive access tier, HYCU ensures your data is stored most cost-efficiently because the archive access tier is optimized for storing data that is not accessed frequently and is stored for at least 180 days.
- An Azure target that has the immutability policy (WORM) set is represented by the icon in the list of targets.

Recommendations

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

the security settings to restrict access to data.

- To enable HYCU to delete the temporary containers that are created during backup and restore operations, it is recommended that you do not set up Azure targets that belong to storage accounts on which version-level immutability support is enabled.

Accessing the Targets panel

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

Procedure

1. In the Targets panel, click **Add**. The Add Target dialog box opens.
2. Depending on what kind of target you want to set up, do one of the following:
   - *To set up an Azure target:* Select **Azure**, and then click **Next**.
   - *To set up an Azure Government or Azure China target*: Click **Other**, select **Azure Government** or **Azure China**, and then click **Next**.
3. Enter a name for the target and, optionally, its description.
4. In the Concurrent Backups field, specify the maximum number of concurrent backups.

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

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

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

7. Click **Next**.
8. In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
9. Enter the storage account name, the secret access key, and the container name.

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

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

    > **Note** To be able to import the encrypted target for restoring virtual machines, applications, file shares, volume groups, and buckets, export the encryption key to a file and keep this file safe. For instructions, see "Exporting an encryption key" on page 471.

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

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

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

12. Click **Save**.

The target is added to the list of targets. For details on managing targets, see "Managing targets" on page 417.

Setting up an Amazon S3 / S3 Compatible target

HYCU supports using an Amazon S3 / S3 Compatible target to store data in the storage provided by Amazon S3 or any of the supported S3 compatible cloud storage solutions.

Prerequisites

- The service must be configured and accessible.
- The S3 bucket must be created and configured in Amazon S3 or any of the supported S3 compatible cloud storage solutions. For a list of supported cloud storage solutions, see the *HYCU Compatibility Matrix*.
- The following minimum required Amazon S3 permissions must be specified:
  - General permissions:

    `s3:GetObject`, `s3:GetObjectRetention`, `s3:DeleteObject`, `s3:PutObject`, `s3:ListBucket`, `s3:GetBucketAcl`, `s3:ListBucketMultipartUploads`, `s3:GetBucketLocation`, `s3:GetBucketObjectLockConfiguration`, `s3:DeleteObjectVersion`, `s3:ListBucketVersions`, and `s3:GetBucketVersioning`.
  - Additional permissions:
    - *For Amazon S3 targets:* `s3:GetBucketPublicAccessBlock`.
    - *For S3 compatible targets:* `s3:ListMultipartUploadParts` and `s3:AbortMultipartUpload`.
    - *For targets that have Object Lock (WORM) enabled:* `s3:PutObjectRetention`, `s3:PutObjectTagging`, and `s3:GetObjectTagging`.
    - *For Wasabi S3 compatible targets that have Object Lock (WORM) enabled:* `s3:ListMultipartUploadParts`, `s3:AbortMultipartUpload`, and `s3:GetObjectVersion`.
- *Only if you plan to store data to an Amazon S3 target in Amazon Virtual Private Cloud (VPC).* An interface VPC endpoint must be set up.
- *For S3 compatible targets:* If you want to provide secure HTTPS access, the CA certificate/chain must be imported to HYCU. For details, see “Importing a custom certificate” on page 527.
- *For setting up a Tencent Cloud target:* Make sure the service endpoint URL does not contain the bucket name. For example, if the Tencent Cloud access domain is `https://testbucket-1234567890.cos.ap-chengdu.myqcloud.com`, in the HYCU Service endpoint field, enter the URL without the bucket name:

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

Limitations

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

Considerations

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

Recommendations

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

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

Procedure

1. In the Targets panel, click **Add**. The Add Target dialog box opens.
2. Depending on what kind of target you want to set up, do one of the following:
   - *To set up an Amazon S3 or S3 compatible target*: Select **Amazon S3 / S3 Compatible**, and then click **Next**.
   - *To set up an AWS Government East or AWS Government West target:* Click **Other**, select **AWS Government East** or **AWS Government West**, and then click **Next**.

3. Enter a name for the target and, optionally, its description.
4. In the Concurrent Backups field, specify the maximum number of concurrent backups.

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

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

   > ⓘ Important Compression may cause degradation of HYCU performance if used with targets that are reserved for data archives, especially with backup chains that include numerous incremental backup images. Additionally, backing up, copying backup data, and archiving to targets with enabled compression and archiving of compressed data may increase system requirements for the HYCU backup controller.
7. Click **Next**.
8. In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
9. Enter the service endpoint URL.
10. From the Storage Class drop-down menu, select the storage class that you want to use for storing the data. If you leave the Default option selected, Amazon S3 selects the storage class autonomously.

    > ⓘ Important Different storage classes incur different data storage and retrieval costs. For details about Amazon S3 storage classes, see AWS documentation.
11. Enter the bucket name, the access key ID, and the secret access key. The access key and the secret access key are used to authenticate the Amazon API service calls.
12. Use the **Path style access** switch if you want HYCU to use a path-style URL (`https://s3.amazonaws.com/<BucketName>`) to access the bucket. HYCU

by default uses a virtual-hosted–style URL (`https://<BucketName>.s3.amazonaws.com`).

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

    > **Note** To be able to import the encrypted target for restoring virtual machines, applications, file shares, volume groups, and buckets, export the encryption key to a file and keep this file safe. For instructions, see "Exporting an encryption key" on page 471.

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

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

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

15. Click **Save**.

The target is added to the list of targets. For details on managing targets, see "Managing targets" on page 417.

Setting up a Google Cloud target

HYCU supports using a Google Cloud target to store data in Google Cloud Storage.

For details on Google Cloud Storage, see Google Cloud documentation.

Prerequisites

- A Google Cloud service account must be created and added to HYCU. For instructions on how to add a cloud account to HYCU, see "Adding a Google Cloud service account" on page 466.
- A Google Cloud Storage bucket must be created in the project that is linked to the created Google Cloud service account you added to HYCU.
- The service must be configured and accessible.

- *Only if Bucket Lock and Object Retention Lock, or Bucket Lock alone are enabled on the target (WORM).* The Google Cloud service account that you add to HYCU must have the following permissions granted in Google Cloud at the project level:
  - *For Bucket Lock:* `storage.buckets.create`, `storage.buckets.delete`, `storage.buckets.get`, `storage.buckets.update`, `storage.objects.create`, `storage.objects.delete`, `storage.objects.get`, `storage.objects.list`, and `storage.objects.update`.
  - *For Object Retention Lock:* The same permissions as for Bucket Lock with the additional `storage.objects.setRetention` permission.

Limitation

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

Considerations

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

Recommendations

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

Accessing the Targets panel

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

Procedure

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

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

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

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

   > **Note** The specified bucket should be created in a project that is linked to the Google Cloud service account you added to HYCU.
10. From the Cloud Account drop-down menu, select the Google Cloud service account you added to HYCU.

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

    > **Note** To be able to import the encrypted target for restoring virtual machines, applications, file shares, volume groups, and buckets, export the encryption key to a file and keep this file safe. For instructions, see "Exporting an encryption key" on page 471.

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

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

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

13. Click **Save**.

The target is added to the list of targets. For details on managing targets, see "Managing targets" on page 417.

Setting up a tape target

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

Prerequisites

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

For details, see QStar documentation.

Limitations

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

- A tape target cannot be used for storing individual file, application, file share, and bucket data.

Considerations

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

Procedure

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

   > ⓘ Important You must ensure that the QStar cache is large enough to support concurrent archive operations. Keep in mind that specifying several archive jobs to run concurrently may also increase system requirements for the HYCU backup controller.

5. Make sure the **Use for archiving** option is enabled.
6. Make sure the **Enable compression** option is disabled.
7. Click **Next**.
8. Follow the instructions that are relevant for your tape target type:

Target type	Instructions
QStar NFS	a. *Optional.* In the Size field, enter the maximum space that should be reserved for archive data (in MiB, GiB, or TiB). b. Provide user credentials that HYCU will use to access the shared folder and make web service calls. c. Enter the name of the Integral Volume set where you want to archive data. d. Provide the web service information. If the default port is used and HTTPS access to the QStar server is

Target type	Instructions
	configured, enter the host name of the QStar server. Otherwise, specify the URL that will be used to access the QStar server in the following format: `https://<QStarServer>:<Port>` e. *Optional.* Enter the path to the shared folder of the mounted Integral Volume set. If you leave this field empty, HYCU tries to retrieve the path to the shared folder. f. Use the **Target encryption** switch if you want the data stored on this target to be encrypted. Note If you enable target encryption, keep in mind the following: • The compression ratio may be affected by it (in cases where tape compression is enabled). • To be able to import the encrypted target for restoring virtual machines and volume groups, export the encryption key to a file and keep this file safe. For instructions, see "Exporting an encryption key" on page 471. g. *Only if charges for reading data from the target may apply.* Use the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges. With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.
QStar SMB	a. *Optional.* In the Size field, enter the maximum space that should be reserved for archive data (in MiB, GiB, or TiB). b. *Optional.* Specify the domain in which the account that has access permissions on the shared folder is registered.

Target type	Instructions
	c. Provide user credentials that HYCU will use to access the shared folder and make web service calls. d. Enter the name of the Integral Volume set where you want to archive data. e. Provide the web service endpoint information. If the default port is used and HTTPS access to the QStar server is configured, enter the host name of the QStar server. Otherwise, specify the URL that will be used to access the QStar server in the following format: `https://<QStarServer>:<Port>` f. *Optional.* Enter the path to the shared folder of the mounted Integral Volume set. If you leave this field empty, HYCU tries to retrieve the path to the shared folder. g. Use the **Target encryption** switch if you want the data stored on this target to be encrypted. Note If you enable target encryption, keep in mind the following: • The compression ratio may be affected by it (in cases where tape compression is enabled). • To be able to import the encrypted target for restoring virtual machines and volume groups, export the encryption key to a file and keep this file safe. For instructions, see "Exporting an encryption key" on page 471. h. *Only if charges for reading data from the target may apply.* Use the **Metered target** switch if you want HYCU to try to read the data from other locations first to avoid additional charges. With the Metered target switch enabled, HYCU will try to obtain the data from the snapshot if it is available, or from any other target that contains this data and for which no additional charges apply. If this is not possible, the data will be read from the target.

9. Click **Save**.

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

Setting up a Data Domain target

HYCU supports using a Data Domain target for storing data on a high performance and reliable storage system, providing you with the confidence that your data is safe and secure. By making a direct connection to the Data Domain system through the SDK, HYCU moves data securely without the need to expose mount points or use network shares, therefore allowing you efficient data transfer during backups and restores.

For details on Data Domain storage systems, see Dell documentation.

Prerequisites

- DD Boost must be enabled on the Data Domain system.
- A DD Boost storage unit must be created on the Data Domain system.
- The DD Boost license must be installed on the Data Domain system.
- Your Data Domain user must have the None role assigned.
- The supported Data Domain target version must be used. For a list of supported versions, see the *HYCU Compatibility Matrix*.
- *Only if retention lock is applied to data on your Data Domain system.* Automatic retention lock must be enabled. For details, see Dell documentation.

Limitation

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

Considerations

- DD Boost SDK libraries are embedded in HYCU.
- A Data Domain target that has retention lock (WORM) enabled is represented by the icon in the list of targets.
- *Only if backup data is not encrypted.* HYCU by default uses Data Domain Managed File Replication (MFR) for copying backup data from one Data Domain target to another. If for any reason you do not want HYCU to use MFR, contact HYCU Support.

Accessing the Targets panel

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

Procedure

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

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

   > **Important** The target that you use for archiving data cannot be used for backing up data or storing copies of backup data.
6. Click **Next**.
7. In the Size field, enter the maximum storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
8. Enter the host name or IP address of the Data Domain server that runs on the system.
9. Provide DD Boost user credentials that HYCU will use to access the storage unit on the Data Domain system.
10. In the Data Domain Storage Unit Name field, enter the name of the storage unit that will be used for storing data.
11. From the Data Domain Authentication Mode drop-down menu, select the required level of Data Domain authentication:
    - **None**
    - **Two-way password** *(the default mode)*
12. *Only if you selected the two-way password authentication mode.* From the Data Domain Encryption Strength drop-down menu, select the required level of Data Domain encryption that will be used when transferring data to the target:

- **None**
- **Medium**
- **High**

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

13. Use the **Target encryption** switch if you want HYCU to encrypt the data before it is stored on this target. Because enabling this option may affect the deduplication ratio on the Data Domain system, HYCU by default compresses the data before storing it on this target.

    > **Note** To be able to import the encrypted target for restoring virtual machines, applications, file shares, buckets, and volume groups, export the encryption key to a file and keep this file safe. For instructions, see "Exporting an encryption key" on page 471.

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

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

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

15. Click **Save**.

The target is added to the list of targets. For details on managing targets, see "Managing targets" on page 417.

Defining your backup strategy

HYCU enables you to schedule automatic backups to achieve the optimum level of data protection based on your recovery point and time objectives, and backup retention requirements. Backups can be scheduled to start each time the specific number of minutes, hours, days, weeks, or months has passed.

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

- Recovery Point Objective (RPO)

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

- Recovery Time Objective (RTO)

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

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

- Taking advantage of predefined policies

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

- Creating a custom policy

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

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

- If one of the predefined or custom policies meets all data protection goals of your environment, you can set such a policy as default. For details, see “Setting a default policy” on page 152.
- You can set up the automatic assignment of policies to virtual machines. For details, see “Setting up automatic policy assignment” on page 151.
- You can enable automatic detection of external threats and anomalies in your backup data. For details, see “Creating an R-Shield policy” on page 148.

Taking advantage of predefined policies

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

HYCU comes with the following predefined policies:

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

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

Creating a custom policy

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

Policy option	Description
Copy	Allows you to create a copy of backup data.
Archiving	Allows you to preserve your data for future reference.
Fast restore	Allows you to restore data to the original source in a fast way by keeping local snapshots for the specified retention time. With this option enabled, HYCU will keep more than one snapshot on the original source, depending on your retention settings. This will allow you to restore data in a fast way, reducing downtime.
Backup from replica	Allows you to back up your virtual machines and volume groups from their replicas in remote office/branch office (ROBO) environments. **Important** Make sure that the schedule interval you set for the Nutanix protection

Policy option	Description
	domains that include the virtual machines and volume groups you want to protect is less than or equal to the RPO set in the HYCU policy. Keep in mind that the replication retention for the respective snapshot on the Nutanix cluster is automatically adjusted to the RPO set in the HYCU policy. This allows HYCU to use the Changed Block Tracking (CBT) feature to get a list of changed data since the last snapshot and perform an incremental backup. For details on protecting virtual machines and volume groups through the Nutanix Prism web console, see Nutanix documentation.
R-Shield	Allows you to enable anomaly detection and malware scanning for your backup data.
Auto-assignment	Allows you to set up the automatic assignment of policies to virtual machines. You do this by first assigning tags to virtual machines in their native environment, and then specifying the corresponding keys and values in HYCU policies.

Before you start tailoring the policy to your needs, make sure that the policy option and the backup target type that you plan to select are supported for your environment:

Policy option	Supported for sources	Supported backup target type(s)
Copy	All sources	Target
Archiving	All sources	Target Snapshot[ce]
Fast restore	Nutanix clusters[b] vSphere[c] XenServer	Target

Policy option	Supported for sources	Supported backup target type(s)
	AWS GovCloud (US) Azure Azure Government File servers[d] Servers	
Backup from replica[a]	Nutanix clusters	Target Snapshot[f]
R-Shield	All sources	Target Snapshot[g]
Auto-assignment	Nutanix clusters vSphere XenServer AWS GovCloud (US) Azure Azure Government Servers	Target Snapshot[c]

[a] Not supported for the HYCU backup controller.
[b] The Fast restore option is not supported if the Backup from replica option is enabled and Nutanix DR is enabled in Prism Central.
[c] Not supported for vSphere virtual machines residing on VMFS or NFS datastores.
[d] Supported for Nutanix Files, Dell PowerScale OneFS, and NetApp ONTAP.
[e] Not supported for Azure Local environments and Hyper-V clusters.
[f] Not supported if Nutanix DR is enabled in Prism Central.
[g] Malware scanning is supported for Nutanix clusters and all the supported types of file servers, except for generic file servers.

Creating a policy

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

Prerequisites

- *Only if you plan to specify time windows for backup and backup copy jobs.* The time windows must be created. By specifying time windows, you define time frames when your backup and backup copy jobs are allowed to start. For details on time windows, see "Creating a time window" on page 140.
- *Only if you plan to enable the Archiving policy option.* A data archive must be created. For details on how to do this, see "Creating a data archive" on page 145.
- *Only if you plan to back up data from replicas in ROBO environments.*
  - A protection domain that includes the virtual machines and volume groups that you want to protect must be created and the specified schedule interval must be less than or equal to the RPO set in the HYCU policy. For details on protecting virtual machines and volume groups through the Nutanix Prism web console, see Nutanix documentation.
  - Both the central site Nutanix cluster and the branch office site cluster must be added to HYCU. For details, see "Adding a Nutanix cluster" on page 74.
  - *Only if Nutanix Disaster Recovery (Nutanix DR) is enabled in Prism Central.* The Nutanix cluster that hosts the virtual machines that you want to protect must be registered with Prism Central and your Prism Central user must have a role with sufficient permissions assigned. For details, see "Configuring Prism Central user permissions" on page 556.
- *Only if you plan to enable the Auto-assignment policy option.*
  - You must be familiar with the information described in "Setting up automatic policy assignment" on page 151.
  - *For Nutanix AHV clusters:* The Nutanix cluster that hosts the virtual machines that you want to protect must be registered with Prism Central and your Prism Central user must have a role with sufficient permissions assigned. For details, see "Configuring Prism Central user permissions" on page 556.
- *Only if you plan to enable the R-Shield option in the policy that is assigned to an entity.* At least one R-Shield policy must be created. For details on how to do this, see "Creating an R-Shield policy" on page 148.
- *Only if you plan to enable the R-Shield option and enable malware scanning in the related R-Shield policy.*
  - The time on the entities must be synchronized with the HYCU backup controller.

- *For Nutanix virtual machines and their attached volume groups:* If available, the snapshots will be scanned for malware, otherwise the backup data that is stored on the target will be scanned.

Limitations

Only if you plan to enable the Backup from replica option. Depending on whether Nutanix DR is enabled in Prism Central, the following limitations apply:

- *Nutanix DR is not enabled:* Backing up data from replicas is supported for virtual machines and volume groups.
- *Nutanix DR is enabled:* Backing up data from replicas is supported only for virtual machines.

Considerations

- *Only if you plan to select Snapshot as the backup target type.* When setting the RPO and the retention period, keep in mind that the number of snapshots that will be created by HYCU must not exceed source maximums and snapshot limitations.
- *Only if you plan to enable the Backup from replica option.* If Nutanix DR is enabled in Prism Central, make sure that you configure the protection policy in Prism Central in such a way that at least two snapshots are always available and snapshot retention is longer than the RPO set in the HYCU policy. For instructions, see Nutanix documentation.
- *Only if you plan to enable the R-Shield option.*
  - The initial malware scan is run on all files in the backup of an entity and may take some time to complete.

    The malware scans that follow are run on the files that change between the selected restore point and the last successful backup. For virtual machines and volume groups, another malware scan of all files is run after a change in the disk layout.
  - If you later unassign the policy from an entity, the R-Shield status for the entity and the available restore points is changed to Undefined.

Accessing the Policies panel

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

Procedure

1. In the Policies panel, click **New**. The New Policy dialog box opens.
2. Enter a name and, optionally, a description of your policy.
3. Add any of the following policy options to the list of the enabled options by clicking it:
   - **Backup** *(mandatory)*
   - **Copy**
   - **Archiving**
   - **Fast restore**
   - **Backup from replica**
   - **R-Shield**
   - **Auto-assignment**
4. In the Backup section, do the following:
   a. In the Backup every field, set the RPO (in months, weeks, days, hours, or minutes).
   b. *Only if Target is selected as the backup target type.* In the Recover within field, set the RTO (in months, weeks, days, hours, or minutes).
   c. In the Retention field, set a retention period (in months, weeks, days, or hours) for the data. The retention period defines when a restore point will be expired. For details on data retention, see "Managing data retention" on page 501.

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

   d. Under Backup target type, select the location for storing the protected data:
      - **Snapshot**
      - **Target**
   e. *Only if Target is selected as the backup target type.* Under start new backup chain, select when you want a new backup chain to be started:
      - **Backup threshold**

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

- **Backup chain length**

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

For details about the impact of the backup chain length on the backup time window, see "Creating a backup window" on page 141.

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

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

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

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

g. *Only if you want to specify a backup window.* Enable the **Use backup window** switch, and then, from the Backup Window drop-down menu, select a backup window for backup jobs. If no backup window is available and you want to create one, see "Creating a backup window" on page 141.

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

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

Enabled option	Procedure
Copy	To create a copy of backup data, in the Copy section, do the

Enabled option	Procedure
	following: a. Set a retention period (in months, weeks, or days) for the copy of backup data. b. From the Targets drop-down menu, select one or more targets that you want to use for storing the copy of backup data. If you want your target to be selected automatically, make sure the **Automatically selected** option is selected. The copy target will be different from the target for data safety reasons. > Note When there are several targets available for storing the copy of backup data and multiple copies of backup data are being created in parallel, HYCU distributes these copies accordingly among targets based on the estimated size of queued and running backups on them. c. *Only if you want to specify a copy window.* Enable the **Use copy window** switch, and then, from the Copy window drop-down menu, select a copy window for backup copy jobs. If no copy window is available and you want to create one, see "Creating a copy window" on page 143. By clicking **Manage**, you are automatically directed to the Time Windows dialog box from where you can manage copy windows.
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 "Creating a data archive" on page 145.
Fast restore	To keep more than one snapshot on the source, which allows a fast restore, in the Fast restore section, set a retention period (in months, weeks, days, hours, or minutes) for snapshots. For example, if you set the RPO to

Enabled option	Procedure
	two days and the snapshot retention period to four days, you will have two snapshots available on the source. **Note** The snapshot retention period cannot be shorter than the RPO or longer than the backup retention period.
Backup from replica	To back up data from replicas, in the Backup from replica section, from the Central site cluster drop-down menu, select the cluster on which the replicas of your entities reside.
R-Shield	To enable anomaly and malware detection for your backup data, in the R-Shield section, from the R-Shield Policy drop-down menu, select an R-Shield policy. If no R-Shield policy is available and you want to create one, see "Creating an R-Shield policy" on page 148.
Auto-assignment	To set up automatic policy assignment, in the Auto-assignment section, enter a key and a value, and then click **Add**. If required, repeat this step for all the keys and values that you want to add. For details on automatic policy assignment, see "Setting up automatic policy assignment" on page 151. **Important** If the category in Nutanix Prism includes more than one value and you want to add the same key with different values to HYCU, you must repeat this step for each value that you want to add.

6. Click **Save**.

The custom policy is created and added to the list of policies. For details on managing policies, see "Managing policies" on page 422.

Creating a time window

HYCU enables you to define time frames when your backup and backup copy jobs are allowed to start. If you use a time window, the backup or backup copy jobs are started only within the specified hours, therefore improving effectiveness and avoiding an overloaded environment. For example, you can

schedule your backup or backup copy jobs to run on non-production hours to reduce loads during peak hours.

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

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

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

- "Creating a backup window" below
- "Creating a copy window" on page 143

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

Creating a backup window

Considerations

- *Only if you selected Target as the backup target type in your policy.* During the Full/Incremental time frame, full or incremental backups are started depending on the backup chain settings, whereas during the Incremental Only time frame, only incremental backups are started. However, if for some reason (for example, due to the Copy policy option being enabled, a snapshot missing, a disk being added to the virtual machine, and so on) an incremental backup cannot be started, a full backup is started instead, also during the Incremental Only time frame.
- *Only if you selected Snapshot as the backup target type in your policy.* Because the backups performed by HYCU have a minimal effect on your data protection environment, they are started in both the Full/Incremental and the Incremental Only time frame.

Procedure

1. In the Time Windows dialog box, click **New**. The Select Window dialog box opens.
2. Select **Backup Window**, and then click **Next**.

3. Enter a name for your backup window and, optionally, its description.
4. From the Time Zone drop-down menu, specify the time zone for your backup window. You can click one of the displayed time zones (your local time zone or your HYCU backup controller time zone) or select one from the drop-down menu.
5. Click **Full/Incremental** or **Incremental Only** to schedule backups depending on the backup type.
6. Select the week days and hours during which you want backups of the selected backup type to start running. To specify time frames for backups of a different backup type, select another backup type, and then repeat this step.

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

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

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

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

- Specify a backup window when creating a new policy. For details, see "Creating a policy" on page 134.
- Assign a backup window to the existing policy. To do so, select the policy, click **Edit**, and then make the required modifications.

Example

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

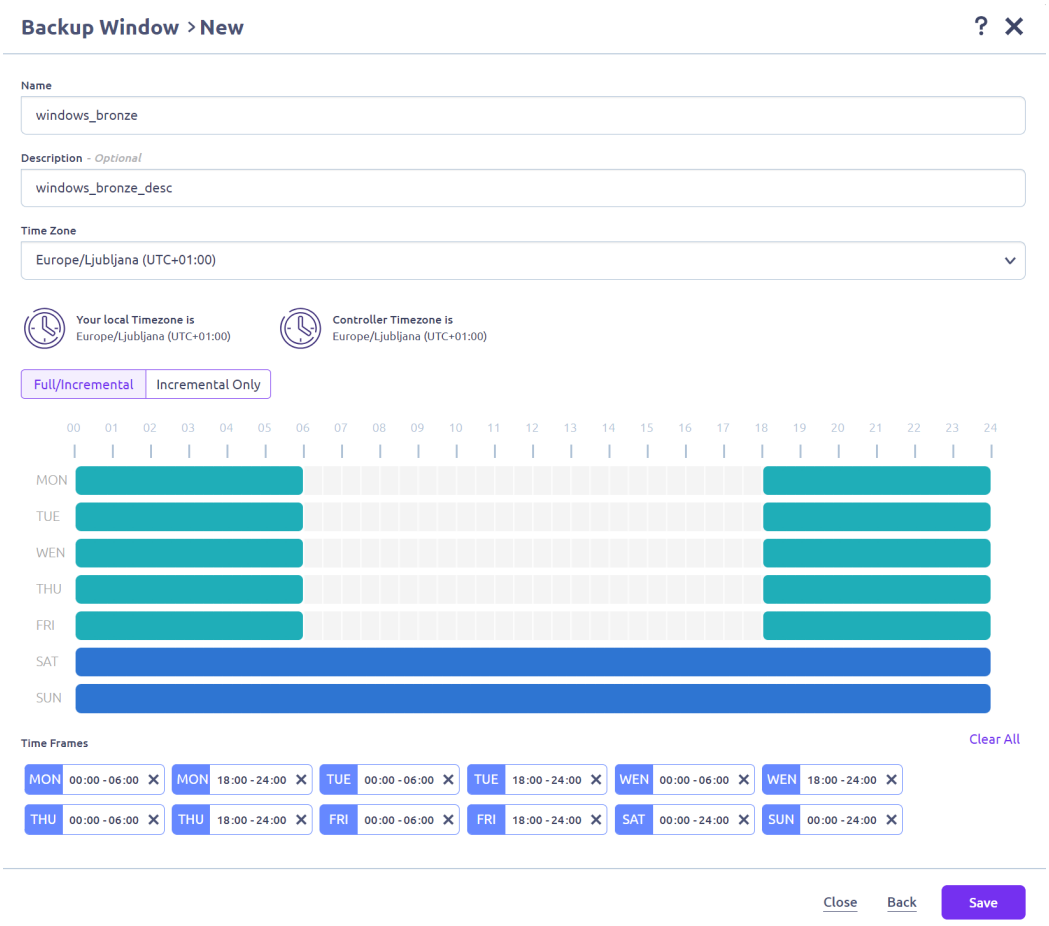

In this case, the backup jobs will be started every 24 hours (full backups will be started only during the weekends) at any point of time within the specified backup windows.

Creating a copy window

Procedure

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

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

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

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

   The selected time frames are displayed in the Time Frames field. If you want to delete any of the selected time frames, click ✕ for the relevant time frame.

6. Click **Save**.

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

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

- Specify a copy window when creating a new policy. For details, see "Creating a policy" on page 134.
- Assign a copy window to the existing policy. To do so, select the policy, click **Edit**, and then make the required modifications.

Example

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

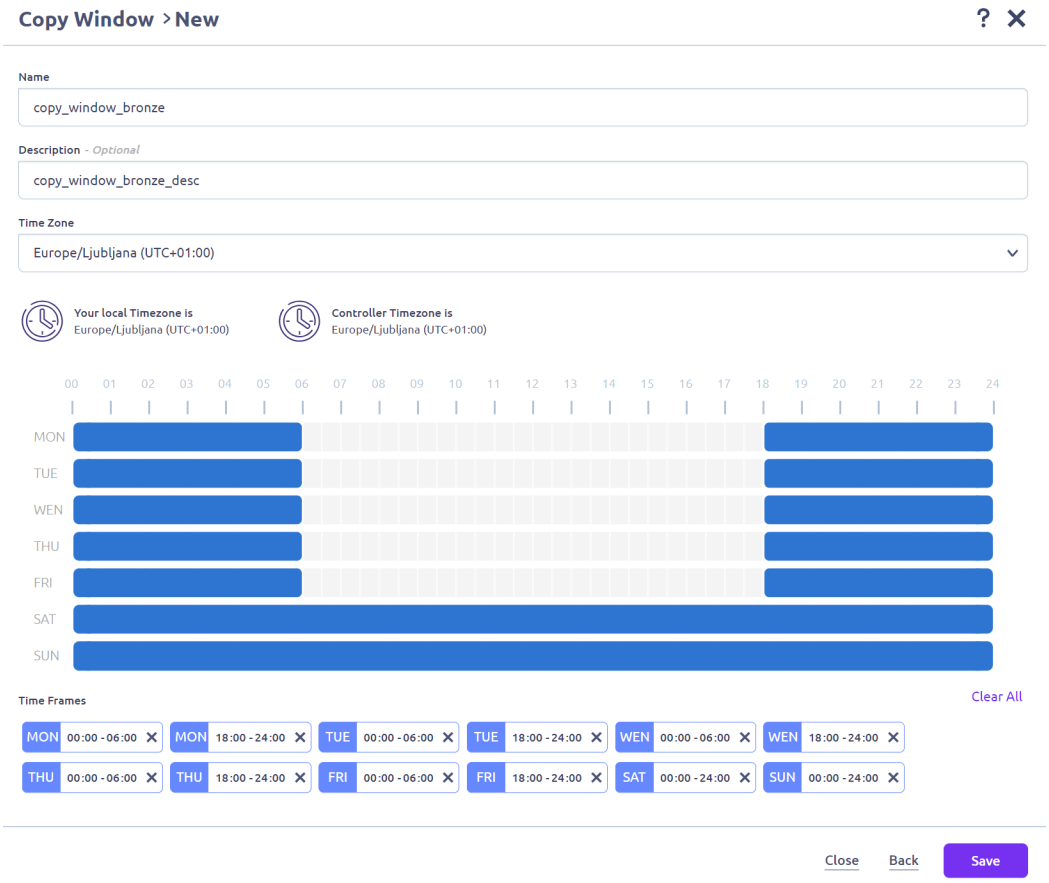

In this case, the backup copy jobs will be started every 24 hours at any point of time within the specified time frames.

Creating a data archive

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

Prerequisites

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

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

Limitations

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

Considerations

- If you edit a data archive, the changes that you make are applicable for archive data that will be created in the future and do not affect existing archive data.
- *Only if you selected Snapshot as the backup target type in your policy.* The configuration settings that HYCU uses for archiving are the ones that the virtual machine has at the time when archiving starts.

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

Procedure

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

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

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

Archiving option	Instructions
Daily	a. In the Recur every field, specify whether you want the data to be archived every day or every few days. b. Use the **Apply only on weekdays** switch if you want the data to be archived only on weekdays.
Weekly	a. In the Recur every field, specify whether you want the data to be archived every week or every few weeks. b. Select one or more days of the week on which you want the data to be archived. **Note** If you select more than one day, archive compliance is calculated by taking into account data archives of all the selected days, not only the latest data archive.
Monthly	a. In the Recur every field, specify whether you want the data to be archived every month or every few months. b. Select whether you want the data to be archived on the same day of the month (for example, on the fifth day of the month), or on a specific day of the month (for example, on the second Friday of the month).
Yearly	a. In the Recur every field, specify whether you want the data to be archived every year or every few years. b. Select whether you want the data to be archived on the same day of the preferred month (for example, on the fifth day of January), or on a specific day of the preferred month (for example, on the second Friday of April).

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

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

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

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

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

- Specify a data archive when creating a new policy. For details, see "Creating a policy" on page 134.
- Assign a data archive to the existing policy. To do so, select the policy, click **Edit**, and then make the required modifications.
- Archive data manually. For details, see "Archiving data manually" on page 435.

Creating an R-Shield policy

You can create an R-Shield policy that helps ensure integrity and security of your data. When configuring the R-Shield policy, you can enable the following:

- Anomaly detection to check the backup data size and detect changes in the standard deviation of the data size between backups. The result is the anomaly detection status.
- Malware scanning to scan backed-up data for threats that are recognizable by using the YARA rules. The result is the malware detection status.

Anomaly detection and malware scanning begin after the R-Shield option is enabled in a policy and that policy is assigned to an entity. For instructions, see "Creating a policy" on page 134.

After you create an R-Shield policy and enable the R-Shield option in a policy that is assigned to an entity, you can do the following:

- View the R-Shield status for the entities to which this policy is assigned. The R-Shield status is visible in the R-Shield column in the relevant panels and in the Detail view for all the available restore points. By hovering over the R-Shield icon, if enabled, you can also separately see the anomaly detection

and malware detection statuses. For details about the R-Shield statuses, see “Viewing entity details” on page 399.

- Run malware scans manually. For details, see “Running malware scans manually” on page 431.

Prerequisites

Only if you plan to enable malware scanning.

- *For virtual machines, applications, and volume groups:* You must own a premium tier Platform license. For details, see “Licensing” on page 488.
- *For file shares:* A HYCU instance is required. For details, see “HYCU instances” on page 25.

Limitations

Only if you plan to enable malware scanning. Malware scanning is not supported for the following entities:

- Volume groups that are not attached to virtual machines. Malware scanning can be run only for an entire virtual machine and the attached volume groups.
- Generic file shares.
- Buckets.

Considerations

- If you delete an R-Shield policy, the R-Shield option is automatically disabled in the policies that use the deleted R-Shield policy.
- If after enabling malware scanning you notice an increasing number of queued malware scan jobs in the Jobs panel, the available resources for malware scanning are insufficient. To make more resources available, you can do the following:
  - Follow the recommended system requirements for your HYCU backup controller. For details, see “Sizing resources for your HYCU backup infrastructure” on page 27.
  - Increase the available memory on your HYCU instance. To do so, sign in to the management console of the environment where the HYCU instance is deployed, and then modify the memory configuration as required.
  - Deploy additional HYCU instances in your data protection environment.

For instructions, see "Creating a HYCU instance by using the HYCU web user interface" on page 484

- *For malware scanning:* The YARA rules that are used for malware scanning are automatically updated from the public ReversingLabs repository.

Recommendation

Only if you plan to enable malware scanning. For optimal malware scanning performance, the following is recommended:

- A HYCU instance should be available in your data protection environment. If a HYCU instance is not available, malware scanning is performed by the HYCU backup controller. Running malware scans on the HYCU backup controller may have a significant impact on its resources and could cause delays in data protection tasks.
- To perform malware scanning with maximum speed, enable hardware-assisted virtualization on the HYCU instance (if available) or the HYCU backup controller.

Procedure

1. In the Policies panel, click **R-Shield**.
2. Click **New**.
3. Enter a name for your R-Shield policy and, optionally, its description.
4. *Only if you want to enable anomaly detection.* Enable the **Anomaly detection** switch.
   a. In the Detection Threshold field, set the threshold for the change in the standard deviation of data size between the backups. The threshold value of 100% means that the size of backup can deviate by up to 100% of the standard deviation from the average value. If the change is higher than the threshold that you set, HYCU will set the anomaly detection status to Suspicious.
   b. In the Detection Range field, set the range for calculating the change in the standard deviation of data size. To define the range, select one of the following:
      - To use the number of backups, select **Number unit: Backups**, and then enter the preferred number.
      - To use the amount of time, select **Time unit: Months**, **Weeks**, **Hours**, or **Minutes**, and then enter the preferred time interval length.

5. *Only if you want to enable malware scanning.* Enable the **Malware scanning** switch.
6. Click **Save**.

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

Setting up automatic policy assignment

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

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

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

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

Considerations

- If you want a predefined policy to be automatically assigned to a virtual machine, when specifying the tag and the matching key and value, you can use the name of the policy (Gold, Silver, Bronze, or Exclude). Keep in mind that if you use the Exclude value, the virtual machine will be excluded from the backup.
- Assigning policies automatically does not affect virtual machines that already have a policy assigned.
- If the default policy is set, it is never assigned to newly discovered virtual machines that have tags applied, but only to the ones for which no automatic assignment of policies is set up. For details on setting the default policy, see “Setting a default policy” on the next page.

- If the comparison of tags and keys and values returns multiple match results, the policy with the lowest RPO is assigned to the virtual machine.
- *For Nutanix ESXi clusters and vSphere environments:* After you restore a virtual machine for which you have set up automatic policy assignment, the tag value is kept on the restored virtual machine only if the original tag still exists in VMware vSphere.

Procedure

1. Sign in to the management console of your data protection environment.
2. Assign tags to virtual machines for which you want to set up automatic assignment. For instructions, see your platform documentation.
3. Sign in to the HYCU web user interface.
4. Specify the matching keys and values in HYCU policies as described in "Creating a policy" on page 134.

> ⓘ **Important** Depending on your data protection environment, the key and the value that you should enter represent the following:
>
> - *For Nutanix AHV clusters:* The name and the value of the category.
> - *For Nutanix ESXi clusters or vSphere environments:* The tag name and the category of the tag, or the attribute and the value of the custom attribute.
> - *For XenServer environments:* The name of the tag for both the key and the value.
> - *For AWS GovCloud (US), Azure, or Azure Government environments:* The name and the value of the tag.

Setting a default policy

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

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

> **Important** The default policy will be assigned to all existing entities that do not have an assigned policy only when you set it. After that, the default policy is automatically assigned only to newly discovered entities. Therefore, if you later unassign a policy from an existing entity, you must manually assign a new policy if you want the entity to be protected.

Accessing the Policies panel

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

Procedure

1. In the Policies panel, select the policy that you want to set as the default one, and then click **Set Default**. The Set Default Policy dialog box opens.
2. Select the entities to which you want the default policy to be assigned:
   - **Virtual Machines**
   - **Applications**

     > **Important** Setting the default policy for applications is possible only if the default policy is set also for virtual machines.

   - **Volume Groups**
   - **Shares**
   - **Buckets**
3. Depending on whether you want the default policy to be assigned to only newly discovered entities, or both newly discovered entities and already existing entities without an assigned policy, do one of the following:

I want the default policy to be assigned to...	Instructions
Only newly discovered entities.	Click **Save**.
Both newly discovered entities and all existing entities that do not have an assigned policy.	a. Enable the **Assign to entities without policy** switch. b. Click **Save**.

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

Chapter 4

Protecting virtual machines

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

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

Source	Item available for protection
Nutanix cluster	Volume groups (collections of logically related virtual disks) in storage containers **Important** If one or more volume groups are attached to a virtual machine at backup time, they are also backed up during the virtual machine backup. You can view such volume groups and their details in the Volume Groups panel together with all existing volume groups residing on Nutanix clusters that have been added to HYCU. For instructions on how to enable data protection for volume groups independent of virtual machine protection, see "Protecting volume groups" on page 343.
vSphere environment	Virtual machine templates (virtual machines that are used as templates to create other virtual machines)

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

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

- "Planning virtual machine protection" on the next page
- "Backing up virtual machines" on page 182
- "Restoring virtual machines" on page 184
- "Restoring individual files" on page 243

Planning virtual machine protection

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

- "Preparing your data protection environment" below
- "Preparing for disaster recovery of your data" on page 161
- "Preparing for the restore to a different source" on page 163
- "Server specifics" on page 168
- "SpinUp specifics" on page 170
- "Enabling access to data" on page 174
- "Setting up virtual machine backup configuration options" on page 178

Preparing your data protection environment

Prerequisites

- *For vSphere environments and Nutanix ESXi clusters:* VMware Tools of the latest version must be installed on the virtual machines. For detailed information about installing VMware Tools, see VMware documentation.
- *For vSphere environments:* If you want HYCU to use the HotAdd transport instead of the NBDSSL transport when performing data protection operations, the following applies:
  - You must create a HYCU instance on the same cluster on which the virtual machines that you plan to protect are running. For details on HYCU instances, see "HYCU instances" on page 25.
  - The host on which you create the HYCU instance must have access to the datastore on which the virtual machine disks are located.
  - *For VMware vSphere version 8.0 Update 2a or earlier:* If the virtual machine disks are located on an NFSv3 datastore, the HYCU instance must be located on the same host as the disks.
  - *Only if the virtual machine disks are located on a VSAN datastore.* You must create the HYCU instance on a datastore of the same type (VSAN).

  For details on VMware transport methods, see VMware documentation.

- *For Nutanix clusters:* If you plan to perform any of the data protection tasks described in "Configuring Prism Central user permissions" on page 556, the Nutanix cluster that hosts the virtual machines that you want to protect must must be registered with Prism Central and your Prism Central user must have a role with sufficient permissions assigned.
- *For ROBO environments:* If volume groups are attached to virtual machines that you plan to back up and you want these volume groups to be backed up during the virtual machine backup, make sure they are in the same Nutanix protection domain as the virtual machines.
- *For XenServer environments:* You must enable NBD connections on the XenServer. For instructions, see XenServer documentation.
- *For Linux virtual machines that you plan to clone to a XenServer environment:* The `xen-blkfront.ko` and `xen-netfront.ko` drivers must be installed on the virtual machines.
- *For archiving data to a QStar tape target:* 1 GiB of additional free memory must be available on the HYCU backup controller for each concurrent archive job.
- *Only if you plan to validate the virtual machine backup and specify a custom script.*
  - The script must be available on the virtual machine in the accessible folder and must have one of the following extensions:
    - Windows: `bat`, `ps1`, `cmd`
    - Linux: `sh`
  - *For Linux:* You must have permissions to run the script on the virtual machine with the assigned credentials.
- *For Linux virtual machines that you plan to clone to an AWS GovCloud (US) environment:* If you have SSH password authentication and the SSH root login enabled in the SSH service, and you want to use them on the restored virtual machine, the `cloud-init` service, if present, must be properly configured on the original virtual machine to not override the SSH service configuration on the restored virtual machine.
- *For AWS GovCloud (US), Azure, or Azure Government environments:* If you plan to restore individual files or applications, make sure that the HYCU backup controller and the virtual machine whose files or applications you plan to restore are located in the same subnet.
- *For Azure Local environments and Hyper-V clusters:*

  - If any snapshots are present for the virtual machines that you plan to protect, you must delete them before you start protecting the data with HYCU.
  - The administrative share must be present on the cluster storage. For details, see Azure or Hyper-V documentation.
  - *For Linux virtual machines:* Hyper-V Linux Integration Services must be installed. For details, see Azure or Hyper-V documentation.
- *Only if you plan to protect NDB-managed database server virtual machines.* You must own a premium tier Platform license. For details, see "Licensing" on page 488.

Limitations

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

- *For Nutanix ESXi clusters:* If you enabled the Backup from replica policy option, backing up virtual machines that have disks on different containers is not supported.
- *For vSphere environments:* Protecting virtual machines that have vSphere Fault Tolerance enabled is not supported.
- *For Azure Local environments and Hyper-V clusters:* Snapshots are not supported. Backup and restore operations that are performed from the snapshot are not available in HYCU.
- Restoring virtual machines without disks to a XenServer environment is not supported.

Considerations

- In large or medium size data protection environments with virtual machines of larger size (2–4 TiB), keep in mind, that the first backup of such virtual machines takes more time and resources. Consider protecting these virtual machines in such a way that they are not backed up simultaneously. You can assign a policy to a large virtual machine, wait until it gets protected, and then continue with protecting other virtual machines.
- *For Nutanix clusters, vSphere, AWS GovCloud (US), Azure, or Azure Government environments:* Archiving is performed from a snapshot if the snapshot is available. Otherwise, archiving is performed from the target (if Target is defined as the backup target type in your policy).
- *Only if archiving data to a cloud target.* If the snapshot of a previous data archive exists (if you enabled the Fast restore policy option or specified Snapshot as the backup target type in your policy), HYCU by default performs incremental archiving of data. However, in the case of an AWS GovCloud (US) environment, HYCU always performs full archiving of data. For limitations related to policy options and backup target types, see "Creating a policy" on page 134.
- *For Nutanix clusters:* You can back up vTPM-enabled virtual machines. However, because vTPM device data cannot be copied and saved for security reasons, a blank vTPM device will be attached to the virtual machine during the restore. For most security features, like Secure Boot and Credential Guard, this is not an issue because they rely on the existence of the vTPM device and not the data inside. For other features, like BitLocker, the data stored in the vTPM device is critical, and a blank vTPM device may cause issues. If you use BitLocker, consider keeping a copy of the BitLocker

recovery keys (stored externally or in the Active Directory) or using alternate methods like Nutanix drive encryption.

- *For vSphere environments:*
  - The number of snapshots that can be created per virtual machine may differ due to snapshot limitations. For details, see VMware documentation.
  - If something unexpected occurs during the backup of a virtual machine template (for example, a network problem), the virtual machine template that is converted to a virtual machine as part of the backup process will remain converted. In this case, make sure to convert the virtual machine back to the virtual machine template. For details on how to do this, see VMware documentation.
  - *Only if you use HotAdd and plan to restore individual files or applications.*
    - When restoring from a snapshot, using HotAdd is supported if all the prerequisites are met. For details about HotAdd prerequisites, see VMware documentation.
    - Using HotAdd is supported only for snapshots that were created automatically by HYCU, not for snapshots that were created by using the Recreate Snapshot option.
  - *Only if you want HYCU to use the HotAdd transport instead of the NBDSSL transport when performing data protection operations.* Each HYCU instance can perform up to 13 concurrent data protection operations by using the HotAdd transport. If additional operations are started after this number is reached, they are performed by using the NBDSSL transport. If you want to make sure that all data protection operations are performed by using the HotAdd transport, you can do the following:
    - *Recommended.* Create additional HYCU instances. For details, see "Deploying the HYCU virtual appliance" on page 23.
    - In vSphere, add up to three additional VMware Paravirtual SCSI controllers (PVSCSI controllers) to the HYCU instance. Each additional PVSCSI controller allows the HYCU instance to perform up to 15 additional concurrent data protection operations by using the HotAdd transport.
  - If a storage container is presented as an NFS datastore to the VMware infrastructure, a full backup of a virtual disk on such a storage container copies the entire allocated disk, not only the used blocks.

- *For Nutanix ESXi clusters, AWS GovCloud (US) environments, Azure environments, and Azure Government environments:* If the snapshot that HYCU used to perform a full backup is missing on the source, the next virtual machine backup will be a full backup.
- *For protection domains configured with NearSync:* Although snapshots in a protection domain are created in a 1 to 15-minute interval, HYCU uses only the snapshots that are created on an hourly basis for backing up and restoring from snapshots. This applies to the following environments:
  - Nutanix ESXi clusters
  - Nutanix clusters when using the Backup from replica option
- If you want the virtual machine details section in the Nutanix Prism web console and vSphere Client to contain the information on which HYCU policy is assigned to a virtual machine, in the HYCU `config.properties` file, set the `hycu.policy.description` configuration setting to `true`. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
- *For ROBO environments:* The number of snapshots in the protection domain may be higher than configured if HYCU uses these snapshots for backing up virtual machines and volume groups.
- *For XenServer environments:* If one of the servers is down, you cannot perform any of the following data protection operations:
  - Backing up virtual machines that are hosted on shared storage.
  - Restoring virtual machines that are hosted on shared storage.
  - Cloning virtual machines to a XenServer environment when shared storage is selected as the restore destination.

Recommendations

- *For ROBO environments:* If a volume group is attached to several virtual machines that you plan to back up and you want this volume group to be backed up as well, it is recommended that it is attached only to the virtual machines inside the same Nutanix protection domain.
- *For virtual machines in a ROBO environment:* To ensure that applications on virtual machines are up and running after restoring the virtual machines, it is recommended that application-consistent snapshots are created for them. For details on how to do this, see Nutanix documentation.
- *For Nutanix AHV clusters:* If your virtual machine is protected with the synchronous replication schedule, the virtual machine snapshots created by

HYCU are replicated also to the remote sites that HYCU cannot access. This results in HYCU not being able to perform the snapshot cleanup. Therefore, to avoid excessive storage consumption and having to manually delete snapshots, it is recommended to protect your virtual machines with the asynchronous or NearSync replication schedules.

- *For using the HotAdd transport when performing vSphere data protection operations:* To ensure optimal performance, it is recommended that the HYCU instance is located on the same type of datastore as the virtual machine disks.

Preparing for disaster recovery of your data

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

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

Target type	Required information for importing
NFS	• NFS server name or IP address • Shared folder
SMB	• Domain (if used) • User name (if used) • Password (if used) • SMB server host • Shared folder
Nutanix	• URL • User name • Password
Nutanix Objects	• Service endpoint • Bucket name • Access key ID • Secret access key • Path style access
iSCSI	• Target portal • Target name • User (if CHAP authentication is enabled) • Target secret (if CHAP authentication is enabled) • Perform mutual authentication (if CHAP authentication is enabled)
Amazon S3 / S3 Compatible	• Service endpoint • Bucket name • Access key ID • Secret access key • Path style access
Azure	• Storage account name • Secret access key • Storage container name
Google Cloud	• Bucket name • Google Cloud service account
QStar NFS	• User name • Password (if used) • Integral volume set name • Web service endpoint

	• Shared folder (if used)
QStar SMB	• Domain (if used) • User name • Password (if used) • Integral volume set name • Web service endpoint • Shared folder (if used)
Data Domain	• Data Domain server host name or IP address • User name • Password • Storage unit name • Authentication mode • Encryption strength

Considerations

- The RPO in the policy that is assigned to the HYCU backup controller should always be lower than any RPO already set for other protected entities in the data protection environment.
- *Only if you use more than one HYCU backup controller for data protection.* Each HYCU backup controller must be protected from within its own web user interface.

Recommendation

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

Preparing for the restore to a different source

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

Prerequisites

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

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

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

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

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

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

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

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

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

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

- *For Linux virtual machines that you plan to clone to an AWS GovCloud (US) environment:* If you have SSH password authentication and the SSH root

login enabled in the SSH service, and you want to use them on the restored virtual machine, the `cloud-init` service, if present, must be properly configured on the original virtual machine to not override the SSH service configuration on the restored virtual machine.

- *For Linux virtual machines that you plan to restore to an Azure Local environment or a Hyper-V cluster:* Hyper-V drivers (`hv_vmbus`, `hv_storvsc`, and `hv_netvs`) must be included in initramfs. To add the drivers, on the virtual machine, run the following command as the root user:

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

- *For restoring a virtual machine to an AWS GovCloud (US), Azure, or Azure Government environment:*
  - Access to the virtual machines through SSH or remote desktop connection must be enabled and a firewall must be configured to allow a remote desktop or SSH connection using a public network.
  - Appropriate credentials must be assigned to the virtual machines that you plan to restore or to the virtual machines on which the applications that you plan to restore are running. For instructions on how to assign credentials to a virtual machine, see “Enabling access to data” on page 174.
  - *For Linux virtual machines:*
    - DHCP must be enabled on the virtual machines that you want to migrate to cloud.
    - Privileged access to the Linux system as root or by using the sudo command without a password is required.
    - The use of persistent network device names based on MAC addresses must be disabled. For details on how to do this, see your Linux distribution documentation.
    - Hyper-V drivers (`hv_vmbus`, `hv_storvsc`, and `hv_netvs`) must be included in initramfs. To add the drivers, on the virtual machine, run the following command as the root user:

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

    - In the `/etc/fstab` system configuration file of the virtual machine, LABEL or UUID must be used instead of device names for file system device identification (for example, `UUID=8ff089c0-8e71-4320-a8e9-`

`dbab8f18a7e5`). If not, platform readiness check will issue a warning in the backup job report.

- *For disaster recovery:* The virtual machine must have the DR-ready status. This means that all backups in the current backup chain must be stored on the respective cloud target. You can check the DR-ready status of a virtual machine in the Virtual Machines panel.

Limitation

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

Considerations

- During the backup, HYCU performs the platform readiness check to ensure that the virtual machine can be successfully restored to a different source. You can view the platform readiness check status in the backup job report.
- If during a restore of the selected virtual machine you receive a warning message indicating that there is a guest operating system mismatch detected (between the guest operating system that is running on the virtual machine and the one specified during the configuration of the virtual machine) or a memory size mismatch detected while creating a new virtual machine, make sure to modify the virtual machine configuration after the restore by specifying the appropriate guest operating system or memory. By doing so, you make sure that the restored virtual machine has the same configuration as it had before the restore. For details on how to do this, see Nutanix or VMware documentation.
- *For virtual machines with attached volume groups:* You must reattach the volume groups to the virtual machine after the restore. For details on how to do this, see Nutanix and guest operating system documentation.
- Depending on your virtual machine original environment and target environment, some additional steps may be required after the restore. For details, see “After restoring a virtual machine to a different source” on page 636.

Recommendations

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

virtual machine will start after the restore (otherwise, you will need to perform additional manual steps as described in “After restoring a virtual machine to a Nutanix AHV cluster” on page 637):

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

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

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

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

The following output confirms that the VirtIO drivers are available:

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

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

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

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

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

```
97084 blocks
```

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

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

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

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

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

An output similar to the following one should appear:

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

For details, see Nutanix documentation.

Server specifics

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

Prerequisites

- Access to the file system data must be enabled. For instructions, see "Enabling access to data" on page 174.
- Sufficient disk space—estimated at up to 1.8 percent of the space of all volumes that you plan to back up—must be available for the index created by HYCU for data protection purposes at the following location:
  - Windows: `%programdata%\HYCU\hycuraw`
  - Linux: `/var/opt/hycu/hycuraw`
- *For Windows servers:*
  - The VSS service must be enabled and running, and the VSS writer status must be stable.
  - WinRM must be enabled and configured by using the `winrm quickconfig` command.
  - *For cloning a Windows server to a Nutanix AHV cluster:* Make sure the Nutanix VirtIO package is installed on the server before you back it up. For detailed information about installing Nutanix VirtIO, see Nutanix documentation.
- *For Linux servers:*
  - Access to the server through SSH must be enabled.
  - *Only if you plan to back up data by using LVM snapshots (the recommended approach).* Sufficient space in the volume group must be available for LVM snapshots. It is recommended that at least 10 percent of free space is available in each volume. However, the percent should be higher if a large number of writes to volumes is expected during the backup. For more information, see LVM documentation.
  - Privileged access to the Linux system as root or by using the sudo command without a password is required.
  - The dm-snapshot kernel module must be included in initramfs. To add the module, on the server, run the following command as the root user:

    ```
    dracut -f --add-drivers "dm-snapshot"
    ```

  - *For cloning a Linux server:* The following drivers must be added to the guest OS kernel:
    - *For cloning to a Nutanix AHV cluster:* Nutanix VirtIO drivers (`virtio_pci`, `virtio_blk`, `virtio_scsi`, `virtio_net`)

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

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

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

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

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

Limitations

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

Consideration

For Linux servers: By default, HYCU uses LVM snapshots for data protection. However, if you cannot provide the required space for LVM snapshot storage in each volume, you can configure HYCU to use device mapper (DM) snapshots as an alternative. For details, see "Enabling DM snapshots" on page 181.

SpinUp specifics

If you plan to use the SpinUp functionality to migrate your virtual machines across on-premises and cloud (AWS, Google Cloud, Azure, or Azure Government) environments, get familiar with all the prerequisites, limitations, considerations, and recommendations described in this section.

Prerequisites

- To be able to migrate your virtual machines across the on-premises and cloud environments, you must configure your environment to pass a platform readiness check during the virtual machine backup, or apply the configuration changes that are required for the migration to cloud as part of the migration procedure:

- To configure your environment to pass a platform readiness check during the virtual machine backup, make sure that the following prerequisites are met before backing up the virtual machine:
  - Access to the virtual machines through SSH or remote desktop connection must be enabled and a firewall must be configured to allow a remote desktop or SSH connection by using a public network.
  - Appropriate credentials must be assigned to the virtual machines that you plan to migrate or to the virtual machines on which the applications that you plan to migrate are running. For instructions on how to assign credentials to a virtual machine, see "Enabling access to data" on page 174.
  - *For migration of Windows virtual machines:*
    - *For migration to a Google Cloud environment:*
      - The Nutanix VirtIO package must be installed on the virtual machines that you plan to migrate.
      - DHCP must be enabled on the virtual machines that you want to migrate to cloud.
    - *For migration to an Azure environment:* DHCP must be enabled on the virtual machines that you want to migrate to cloud.
    - *For migration from an AWS or Azure environment to a Nutanix AHV cluster:* The Nutanix VirtIO package must be installed on the virtual machines that you plan to migrate.
  - *For migration of Linux virtual machines:*
    - DHCP must be enabled on the virtual machines that you want to migrate to cloud.
    - Privileged access to the Linux system as root or by using the sudo command without a password is required.
    - The use of persistent network device names based on MAC addresses must be disabled. For details on how to do this, see your Linux distribution documentation.
    - In the `/etc/fstab` system configuration file of the virtual machine, LABEL or UUID (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`) must be used instead of device names for file system device identification.
    - The following drivers must be included in initramfs:

- *For migration to an AWS environment:*`ixgbevf`, `ena`, `nvme`, `nvme-core`, `xen_netfront`, and `xen_blkfront`

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

  ```
  dracut -f --add-drivers "ixgbevf ena nvme nvme-core
  xen_netfront xen_blkfront"
  ```

- *For migration to a Google Cloud environment:*`virtio_pci`, `virtio_net`, and `virtio_scsi`

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

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

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

- *For migration to an Azure or Azure Government environment:* Hyper-V drivers (`hv_vmbus`, `hv_storvsc`, and `hv_netvsc`)

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

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

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

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

  ```
  dracut -f --add-drivers "virtio_net.ko virtio_
  scsi.ko virtio_pci.ko"
  ```

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

- If you plan to enable the **Adapt OS for migration** option to apply the configuration changes that are required for the migration to cloud as part of the migration procedure, make sure that the following prerequisites are met before backing up the virtual machine that you

plan to migrate:

- *For Windows 11 virtual machines:* The EMS and SAC Toolset for Windows (client edition) optional feature must be installed on the virtual machine.
- *For Linux virtual machines:* The HYCU backup controller must run on the same or higher CPU microarchitecture level (for example, x86-64-v3) as the virtual machine.

For details on the Adapt OS for migration option, see "Protecting data across on-premises and cloud environments" on page 586.

- *For disaster recovery to cloud:* The virtual machine that you plan to migrate must have the DR-ready status. This means that all backups in the current backup chain must be stored on the respective cloud target. You can check the DR-ready status of a virtual machine in the Virtual Machines panel.

Limitations

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

Considerations

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

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

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

command in PowerShell:

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

Recommendations

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

Enabling access to data

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

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

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

- You plan to use the SpinUp functionality to migrate your virtual machines and applications to cloud.

Prerequisites

- A firewall must be configured to allow inbound network traffic through the required TCP port.
- To ensure that HYCU can use Kerberos authentication, the user name must include the DNS domain name. The legacy NetBIOS names are not supported by Kerberos. For example, `hycu.local\administrator` is sufficient, whereas `HYCU\administrator` is not.
- *Only if the WinRM protocol over HTTPS will be used.* HYCU must be configured to use HTTPS for WinRM connections to virtual machines. For instructions, see "Enabling HTTPS for WinRM connections" on page 538.

Limitation

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

Considerations

- *For Windows virtual machines:* When specifying a user name, make sure to use one of the following formats:
  - If the virtual machine is added to an Active Directory domain: `<Domain>\<Username>` or `<Username>@<Domain>`
  - If the virtual machine is not added to an Active Directory domain: `<Username>`, `.\<Username>`, or `<Hostname>\<Username>` (in this case, `<Hostname>` is the value of the `COMPUTERNAME` variable).
- *For Linux virtual machines:* When specifying a user name, make sure to use the following format:
  - If the virtual machine is not joined to a realm or if the default realm is configured: `<Username>`.
  - If the realm domain is part of the sign-in: `<Username>@<Domain>` or `<Username>@<FQDN>`.
- *For virtual machines that you plan to back up from their replicas in ROBO environments:* Make sure that the most recent replica reflects the state of the virtual machine.

Accessing the Virtual Machines panel

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

Procedure

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

Protocol options	Instructions
Automatic	Select this option if you want HYCU to automatically select a protocol for accessing the virtual machine: the SSH protocol (TCP port 22) or the WinRM protocol (HTTPS transport and TCP port 5986, or HTTP transport and TCP port 5985), and then enter the user name and password of a user account that has required permissions to access the virtual machine. **Note** *For Linux virtual machines:* Password authentication is used by default. If you want to use public key authentication, select the **SSH** protocol option and make the required modifications.
SSH	Select this option if you want to use the SSH protocol, and then do the following: a. In the Port field, enter the SSH server port number. b. From the Authentication type drop-down menu, select the type of authentication you want to be used and provide the required information: • **Password authentication** Enter the user name and password of a user account that has required permissions to access the virtual machine.

Protocol options	Instructions

	• **Public key authentication**
	  • In the Username field, enter the user name of a user account that has required permissions to access the virtual machine.
	  • Choose a private key.
	  > Note *Only if you are signed in to HYCU as a self-service group administrator.* If you use Conjur for managing your HYCU secrets, you can enable the **Retrieve values from secrets manager** switch if you want to provide the secret instead of browsing for the file. For details on managing secrets, see "Managing secrets" on page 509.
	  • *Only if the private key is encrypted.* Enter the private key passphrase.

WinRM	Select this option if you want to use the WinRM protocol, and then do the following: a. From the Transport drop-down menu, select the type of transport you want to be used. b. In the Port field, enter the WinRM server port number. c. Enter the user name and password of a user account that has required permissions to access the virtual machine.

6. Click **Save**.
7. Click **Assign**.

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

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

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

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

Setting up virtual machine backup configuration options

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

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

I want to...	Instructions
Specify the pre/post-backup and pre/post-snapshot scripts.	“Specifying pre/post-backup and pre/post-snapshot scripts” on the next page
Specify any disks or volume groups to exclude or include when backing up a virtual machine.	“Excluding or including disks in the backup” on page 180
Applicable only for Linux servers. Configure HYCU to use DM snapshots instead of LVM snapshots for backing up data.	“Enabling DM snapshots” on page 181

Accessing the Virtual Machines panel

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

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

You can use the pre/post-backup and pre/post-snapshot scripts to perform necessary actions before the backup is performed or the snapshot is created (for example, to suspend application I/O), and after the backup is performed or the snapshot is created (for example, to resume application I/O). For details on how to specify the scripts, follow the procedure described in this section. For details on exit codes and exported environment variables, see "Using the pre and post scripts" on page 561.

Prerequisites

- Access to the virtual machine file system must be enabled. For instructions, see "Enabling access to data" on page 174.
- A script must be available in the accessible folder and must have one of the following extensions:
  - Windows: `bat`, `ps1`, `cmd`
  - Linux: `sh`
- *For Linux:* You must have permissions to run a script on the virtual machine with the assigned credentials.

Procedure

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

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

Excluding or including disks in the backup

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

- By excluding disks, you make sure that only the selected disks are not backed up. In this selection mode, if a disk is later added to the virtual machine, it is backed up.
- By including disks, you make sure that only the selected disks are backed up. In this selection mode, if a disk is later added to the virtual machine, it is not backed up.

Prerequisite

You must be an owner of the virtual machine whose disks you want to exclude from or include in the backup. For instructions on how to set ownership of a virtual machine, see "Setting ownership of virtual machines" on page 453.

Limitations

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

Considerations

- The next backup after changing the virtual machine backup scope will be a full backup.
- Excluding disks with protected applications may affect application protection.
- If any disks are excluded from the backup (manually or automatically), the virtual machine will be restored or migrated to cloud without such disks or with blank disks if you select the option to create excluded disks as blank. The corresponding restore point labels are marked with a red circle. For details, see "Viewing entity details" on page 399.

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

Procedure

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

I want to...	Instructions
Exclude disks and volume groups from the backup.	a. Click **Exclude selected vDisks**, and then select the disks or volume groups that you want to exclude from the backup. b. Click **Save**. **Important** *For vSphere environments:* If you plan to restore individual files, make sure not to exclude the operating system disk from the backup.
Include disks and volume groups in the backup.	a. Click **Include selected vDisks**, and then select the disks or volume groups that you want to include in the backup. b. Click **Save**.

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

Enabling DM snapshots

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

Considerations

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

Procedure

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

Backing up virtual machines

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

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

Prerequisites

- *For Nutanix ESXi clusters and vSphere environments:* You must have the required backup privileges assigned. For details, see "Assigning privileges to a vSphere user" on page 552.
- *For Linux virtual machines residing in cloud environments:* If you use LVM to span the logical volumes over multiple disks, you must freeze and unfreeze the relevant filesystems as a part of the backup procedure. To achieve this, use the `/usr/sbin/fsfreeze` utility in your pre/post-snapshot scripts. For details, see "Specifying pre/post-backup and pre/post-snapshot scripts" on page 179.

- *Only if you plan to protect servers or volume groups that are attached to a virtual machine by using iSCSI.* Credentials must be assigned to servers that you want to protect or to virtual machines whose volume groups you want to protect. For instructions, see "Enabling access to data" on page 174.

Limitation

For virtual machines that have Azure Disk Encryption enabled: The key vault is not protected by HYCU.

Considerations

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

Accessing the Virtual Machines panel

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

Procedure

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

   > **Tip** You can update the list of virtual machines by clicking **Refresh**. To narrow down the list of displayed virtual machines, you can use the filtering options described in "Filtering and sorting data" on page 406.

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

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

The backup is scheduled according to the values that you defined for your policy. If required, you can also perform a manual backup at any time. For details, see "Performing a manual backup" on page 424.

Restoring virtual machines

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

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

Prerequisites

- If you are restoring a virtual machine to the same source and you want the existing ISO image to be attached to the restored virtual machine, make sure the ISO image that was attached to the virtual machine at backup time still exists on the source at virtual machine restore time and its name and

location are the same.

- *For restoring data from tape:* If the tape target is being actively used for archiving data, its mode should be set to Read Only. For details on how to edit a target, see "Managing targets" on page 417.
- *For Nutanix ESXi clusters and vSphere environments:* You must have the required restore privileges assigned. For details, see "Assigning privileges to a vSphere user" on page 552.
- *For servers:* At least one hypervisor or cloud source must be added to HYCU to provide a storage container for storing the restore data. For details on how to add sources to HYCU, see "Adding sources" on page 73.
- If you are a self-service group user, you must do the following:
  - *For AWS GovCloud (US) environments:* Add an AWS GovCloud (US) account to HYCU. For details, see "Adding an AWS GovCloud (US) account" on page 465.
  - *For Azure environments:* Add an Azure service principal to HYCU. For details, see "Adding an Azure service principal" on page 468
  - *For Azure Government environments:* Add an Azure Government service principal to HYCU. For details, see "Adding an Azure Government service principal" on page 469.

Limitations

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

Considerations

- A restore is performed from the snapshot only if you are restoring to the same source (the source where the original virtual machine was running). If you are restoring to a different source, depending on the tier that you select

for the restore, the following will happen:

- If you select Snapshot, the restore will fail.
- If you select Automatic, the restore will be performed from the target if there is an available target. Otherwise, it will fail.

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data or validating the virtual machine backup.
- You cannot perform a restore of a virtual machine whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up` configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
- *For AWS GovCloud (US) virtual machines with encrypted volumes:* Depending on whether you are restoring such a virtual machine to the same or a different region, the following applies:
  - To the same region: The restored volumes will be encrypted with the same KMS key as the original ones.
  - To a different region: The restored volumes will be encrypted with the default KMS managed key for EBS encryption.
- *Only if you plan to restore data from a QStar tape target.* During the restore procedure, you can use the Tape Info option to view information about the media. If a note indicating that an empty response was received from QStar is displayed, make sure to check whether the data was archived to the media.

Restore options

Virtual machine restore options

You can select among the following virtual machine restore options:

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

VM restore option	Description
	original virtual machine with the restored one. For instructions, see "Restoring a virtual machine" on the next page. **Important** The Restore VM option cannot be used for restoring the following: • Azure Local virtual machines • Hyper-V virtual machines • Servers
Restore DB server VM	*Available only for NDB-managed database servers.* Enables you to restore an NDB-managed database server virtual machine by restoring the disks of the virtual machine on which the database server is running. For instructions, see "Restoring a database server virtual machine" on page 203.
Clone VM	Enables you to restore a virtual machine by creating a virtual machine clone on the same or a different source. Select this option if you want to keep the original virtual machine. For instructions, see "Cloning a virtual machine" on page 205. **Important** If you plan to restore a virtual machine to a different source, keep in mind that depending on your virtual machine original environment and target environment, you might have to perform some additional steps after the restore. For details, see "After restoring a virtual machine to a different source" on page 636.
Validate VM backup	Enables you to validate the virtual machine backup by creating a virtual machine clone. Select this option if you want to verify that the virtual machine has no corrupted backups. For instructions, see "Validating the virtual machine backup" on page 228. **Note** Performing the backup validation is not supported for the HYCU backup controller and for virtual machines running in AWS GovCloud (US),

VM restore option	Description
	Azure, or Azure Government environments.

Disk restore options

You can select among the following disk restore options:

Disk restore option	Description
Restore vDisks	Enables you to restore virtual disks. Select this option if you want to replace the original virtual disks with the restored ones. For instructions, see “Restoring virtual disks” on page 233. **Important** You cannot restore server disks by using the Restore vDisks option.
Clone vDisks	Enables you to restore virtual disks by creating their clones. Select this option if you want to keep the original virtual disks. For instructions, see “Cloning virtual disks” on page 235. **Important** You cannot restore server disks by using the Clone vDisks option.
Export vDisks	Enables you to restore virtual disks to an NFS or SMB share. Select this option if you want to make the virtual disks available to users with specific access permissions, or if you want to use the virtual disks later to restore data to a server or to a source not supported by HYCU or not added to HYCU. For instructions, see “Exporting virtual disks” on page 238.

Restoring a virtual machine

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

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

- “Restoring a virtual machine to a Nutanix cluster or a vSphere environment” below
- “Restoring a virtual machine to a XenServer environment” on page 193
- “Restoring a virtual machine to an AWS GovCloud (US) environment” on page 196
- “Restoring a virtual machine to an Azure or Azure Government environment” on page 200

For details on how to restore an NDB-managed database server virtual machine, see the following topic:

- “Restoring a database server virtual machine” on page 203

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

Considerations

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

- *Only if you plan to restore a vSphere virtual machine from a datacenter that was removed from HYCU.* After you remove the datacenter, the protection status of the virtual machine changes from Protected to Protected deleted. When restoring such a virtual machine, consider the following:
  - If restoring to a datacenter that is added to HYCU, the protection status of the virtual machine changes back to Protected.
  - If restoring to any datacenter that is not added to HYCU, the protection status of the virtual machine stays Protected deleted.
- *Only if you plan to restore a virtual machine running on a Nutanix ESXi cluster.* If Snapshot is selected as the backup target type in your policy, the `NVRAM` file will not be restored.

Accessing the Virtual Machines panel

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

Procedure

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

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Virtual Machine Options**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. In the General section, do the following:
   a. From the Storage container drop-down menu, select where you want to restore the virtual machine. By default, the original storage container is selected.

      > **Note** If you decide to restore the virtual machine to another storage container, keep in mind the following:
      > - Restore from the Snapshot tier cannot be performed to another storage container.

> - If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.

b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**

c. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

   If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

   - In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
   - In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

     > Note The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

   - In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.

d. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore. The original virtual machine will be deleted automatically.

   > Important *Only if you are restoring a vSphere virtual machine to a vSphere environment and you have disabled the Power virtual machine on switch.* When you try to power on the virtual machine and you are prompted to answer whether the virtual machine has been moved or copied, make sure to answer **I Moved It**.

e. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones

to be created and attached to the restored virtual machine.

f. *For volume groups attached to the virtual machine:* Use the **Restore volume groups** switch if you also want to restore the volume groups that are attached to the virtual machine.

7. In the Network section, review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

   Depending on whether the original networks are available, proceed as follows:

   - If the original networks are available, you can leave the default values and restore the virtual machine with the original network settings, or you can modify the network settings.
   - If the original networks are not available, you must modify the network settings.

   Modifying network settings

Original networks are…	Instructions
Available	You can do the following: • Add a new network adapter by clicking **New**, and then selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking **Delete**.
Unavailable	You can do the following: • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the affected network adapter by selecting it,

Original networks are...	Instructions
	Delete. • Add a new network adapter by clicking **New**, and then selecting the preferred network.

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

8. Click **Restore**.

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

Restoring a virtual machine to a XenServer environment

Accessing the Virtual Machines panel

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

Procedure

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

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Virtual Machine Options**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. In the General section, do the following:
   a. From the Storage Container drop-down menu, select where you want to restore the virtual machine. By default, the original storage container is selected.

> **Note** If you decide to restore the virtual machine to another storage container, keep in mind the following:
>
> - The restore from the Snapshot tier cannot be performed to another storage container.
> - If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.

b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**

c. Enable the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

   If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

   - In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
   - In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

     > **Note** The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

   - In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.

d. Enable the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore. The original virtual machine will be deleted automatically.

e. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.

7. In the Network section, review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

   Depending on whether the original networks are available, proceed as follows:

   - If the original networks are available, you can leave the default values and restore the virtual machine with the original network settings, or you can modify the network settings.

     > **Note** The network name may differ from the one that is configured in XenCenter.

   - If the original networks are not available, you must modify the network settings.

   Modifying network settings

Original networks are...	Instructions
Available	You can do the following: • Add a new network adapter by clicking **New**, and then selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking **Delete**.
Unavailable	You can do the following: • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking **Delete**. • Add a new network adapter by clicking **New**,

Original networks are...	Instructions
	and then selecting the preferred network.

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

8. Click **Restore**.

Restoring a virtual machine to an AWS GovCloud (US) environment

Considerations

- Make sure that the virtual machine you are restoring is not deleted from AWS GovCloud (US). If you delete a virtual machine from AWS GovCloud (US), you cannot restore it even if it still has a valid restore point available in HYCU (that is, even if its status is Protected deleted).
- When restoring a virtual machine, the original virtual machine disks are deleted and replaced with the restored ones.

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

Procedure

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

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Virtual Machine Options**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. The following information is displayed and preselected:

- The AWS GovCloud (US) account to which the virtual machine will be restored.
- The account ID of the AWS GovCloud (US) account to which the virtual machine will be restored.
- The region to which the virtual machine will be restored.
- *Only if found by HYCU.* The key pair name for connection to the restored virtual machine.

  > **Important** If HYCU does not find the name of your key pair and you want to use it, you can select it from the Key pair name drop-down menu.

- The availability zone to which the virtual machine will be restored.

7. Click **Next**.
8. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**
9. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

   If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:
   - In the vCPU threads field, enter the number of CPUs for the restored virtual machine multiplied by the number of cores per CPU and the number of threads per core.
   - In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
   - From the Virtual machine type drop-down menu, select the virtual machine type.

     > **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

10. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
11. Under Network interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the Virtual Private Cloud (VPC) to which the original virtual machine belongs. If required, you can also modify network settings.

    ### Modifying network settings

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

    > **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same Virtual Private Cloud (VPC). The maximum number of network interfaces that you can add depends on the selected virtual machine type.

    Depending on how you want to modify network settings, do one of the following:

    - Click **Add network interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow these steps:
      a. The Virtual Private Cloud (VPC) to which the network interface will be added is displayed and preselected.
      b. From the Subnets drop-down menu, select the subnet to which the network interface should be assigned.
      c. From the Security groups drop-down menu, select one or more security groups that will be associated with the network interface. If you want to select all the available security groups, select **Select all**.
      d. In the Public address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Auto-assign	An automatically allocated public IP address will be

Option	Description
	assigned to the network interface on the restored virtual machine.
Elastic IP (Reserved)	An elastic public IP address that you reserved in AWS GovCloud (US) will be assigned to the network interface on the restored virtual machine.
Elastic IP (New)	An elastic public IP address will be assigned to the network interface on the restored virtual machine.

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

Option	Description
Auto-assign	An automatically allocated private IP address will be assigned to the network interface on the restored virtual machine.
Custom	A private IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

12. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
13. Under Operating system license, select one of the following options:

OS license option	Select this option if you want to...
Keep existing license	Keep the existing OS license on the restored virtual machine.

OS license option	Select this option if you want to...
	ⓘ Important Make sure that the existing license is applicable also in AWS.
Available only for the Windows Server OS. **Replace existing license with AWS license**	Replace the existing OS license with an AWS license on the restored virtual machine.

14. Click **Restore**.

Restoring a virtual machine to an Azure or Azure Government environment

Consideration

If you want the restored virtual machine to have the same static IP address as the original virtual machine, do one of the following:

- Before the restore, in Azure or Azure Government, disassociate the IP address from the original virtual machine, and then select this IP address for the network interface during the restore in HYCU.
- During the restore, select a different IP address for the network interface. After the restore, in Azure or Azure Government, assign the preferred IP address to the restored virtual machine.

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

Procedure

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

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Virtual Machine Options**, and then click **Next**.

5. Select **Restore VM**, and then click **Next**.
6. From the Availability Zone drop-down menu, select the zone for the restored virtual machine. If you do not want to restore data to any zone, select None.
7. Click **Next**.
8. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**
9. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

   If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:
   - In the vCPU cores field, enter the number of virtual CPUs for the restored virtual machine.
   - In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
   - From the Virtual machine type drop-down menu, select the virtual machine type.

     > Note The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.
10. Under Network Interfaces, you can view all the network interfaces that are attached to the virtual machine. Keep in mind that the IP address that is assigned to the primary IP configuration of the network interface will be assigned to the network interface on the restored virtual machine. If required, you can also modify network settings.

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

> **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add Network Interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow these steps:
  a. *Only if you are adding a network interface.* From the Virtual Network drop-down menu, select the network for the network interface.

     > **Note** The list of available networks includes only the ones within the region you selected for the restored virtual machine.

  b. From the Subnet drop-down menu, select the subnet to which the network interface should be assigned.
  c. Under NIC Network Security Group, select the network security group for the network interface. You can select among the following options:

Option	Description
None	The network interface will not be assigned to a network security group.
Basic	The network interface will be assigned to Azure's basic network security group.
Advanced	The network interface will be assigned to the network security group that you select from the drop-down menu. By default, the network security group of the original virtual machine is selected.

  d. Under Public IP Address Type, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.

Option	Description
Static	A static IP address will be assigned to the network interface on the restored virtual machine.
Existing	A preferred public IP address resource that you have created in Azure or Azure Government will be assigned to the network interface on the restored virtual machine.

e. Under Private IP Address Type, select the private IP address for the network interface. You can select between the following options:

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

11. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
12. Click **Restore**.

Restoring a database server virtual machine

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

Considerations

- A virtual machine that contains an NDB-managed database is represented by the icon in the list of virtual machines.

- If volume groups are attached, the original volume group and its virtual disks will be deleted.
- The original virtual disks are deleted.
- The original virtual machine is restarted automatically.

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

Procedure

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

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Virtual Machine Options**, and then click **Next**.
5. Select **Restore Database Server VM**, and then click **Next**.
6. In the General section, do the following:
   a. From the Storage container drop-down menu, select where you want to restore the virtual machine. By default, the original storage container is selected.

      > **Note** If you decide to restore the virtual machine to another storage container, the fast restore or restore from Snapshot tier cannot be performed because the restore will be performed from the target and not from the snapshot.

   b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
      - **Automatic**: Ensures the fastest restore to the latest state.
      - **Backup**
      - **Copy**
      - **Archive**
      - **Snapshot**

7. *Only applicable for Time Machines with a Continuous SLA.* Enable the **Run a Tail Log Backup on NDB before restore** switch to ensure that an NDB Tail Log Backup is started before a database restore. A Tail Log Backup is started automatically after a restore.
8. Click **Restore**.

Cloning a virtual machine

You can create a clone of the original virtual machine by restoring the virtual machine to its original or a new location on the same or a different source. In this case, the original virtual machine will not be overwritten.

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

- "Cloning a virtual machine to a Nutanix cluster, a vSphere environment, an Azure Local environment, or a Hyper-V cluster" below
- "Cloning a virtual machine to a XenServer environment" on page 211
- "Cloning a virtual machine to an AWS GovCloud (US) environment" on page 215
- "Cloning a virtual machine to an Azure or Azure Government environment" on page 220

> ⓘ **Important** After you clone the virtual machine, make sure that everything works as expected by going through the considerations and the recommendations listed in "After cloning a virtual machine" on page 224.

Cloning a virtual machine to a Nutanix cluster, a vSphere environment, an Azure Local environment, or a Hyper-V cluster

Prerequisites

- *For virtual machines that you plan to clone to a new location:* The source to which you plan to clone the virtual machine must be added to HYCU. For details on how to do this, see "Adding sources" on page 73.
- *For Linux servers:* In the `/etc/fstab` system configuration file of the server, UUIDs (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`) must be used instead of device names for file system device identification.
- *For virtual machines that you plan to clone to an Azure Local environment:* The virtual hard disk path and the virtual machine path must point to the cluster

shared volume location. For example:

- Virtual Machines Path: `C:\ClusterStorage\UserStorage\VirtualMachines`
- Virtual Hard Disks Path: `C:\ClusterStorage\UserStorage\VirtualHardDisks`

Limitation

For vSphere environments: Attaching the ISO image to the restored virtual machine is not supported.

Considerations

- *Only if volume groups are attached to the virtual machine that you are cloning.* You can choose to restore the volume groups together with the virtual machine if they were attached to it at backup time. In this case, the original volume groups are kept alongside of the restored ones. If the volume groups are also attached to other virtual machines, the following applies (depending on how they are attached to the virtual machines):
  - Directly: Volume groups are automatically attached only to the cloned virtual machine.
  - By using iSCSI: Volume groups are automatically attached to all virtual machines to which they were attached at backup time.
- *For restoring a virtual machine running on a Nutanix AHV cluster to a Nutanix ESXi cluster:* If virtual machine disks are attached to the PCI bus, the bus type will be automatically changed to SCSI after the restore. Because of this configuration change, the restore finishes with a warning.
- *For Linux virtual machines running on a Nutanix ESXi cluster:* If after restoring a virtual machine that was created through the vSphere Client, the virtual machine does not boot, follow the steps described in "After restoring a virtual machine to a Nutanix ESXi cluster" on page 638.
- After you restore a virtual machine, it might happen that the order of virtual disks differs from the one on the original virtual machine if you performed the restore:
  - From a Nutanix AHV cluster to a Nutanix ESXi cluster or a vSphere environment
  - From a Nutanix ESXi to another Nutanix ESXi cluster
  - From a vSphere environment to a Nutanix ESXi cluster

In this case, make the necessary adjustments, including the selection of the correct boot disk.

- *Only if you plan to restore vSphere virtual machine data to the original storage container.* If the storage container is mounted to several hosts and the original host is powered off or in maintenance mode at restore time, data will be restored to the same storage container on a different host.
- *Only if ownership is set for the virtual machine.* The same owner is automatically assigned to the restored virtual machine.
- *Only if you plan to restore a virtual machine running on a Nutanix ESXi cluster.* If Snapshot is selected as the backup target type in your policy, the `NVRAM` file will not be restored.
- *Only if the original virtual machine resides on a source other than a vSphere environment.* Make sure to modify the virtual machine configuration by specifying the appropriate guest operating system.

Recommendation

For Linux virtual machines: It is recommended that the use of persistent network device names based on MAC addresses is disabled. Otherwise, you will have to configure the network manually. For details on how to disable the use of persistent network device names, see your Linux distribution documentation.

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

Procedure

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

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Virtual Machine Options**, and then click **Next**.

5. Select **Clone VM**, and then click **Next**.
6. From the Destination source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. In the General section, do the following:
   a. From the Storage container drop-down menu, select the storage container where you want to restore the virtual machine.

      > **Note** By default, the original storage container is selected. If you decide to restore the virtual machine to another storage container, keep in mind the following:
      > - Restore from the Snapshot tier cannot be performed to another storage container.
      > - If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.
      > - If the selected storage container is on a different source, additional prerequisites apply. For details, see "Preparing for the restore to a different source" on page 163.

   b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
      - **Automatic**: Ensures the fastest restore to the latest state.
      - **Backup**
      - **Copy**
      - **Archive**
      - **Snapshot**

      > **Note** When selecting the restore point tier, consider the following:
      > - *Only if you select the Archive tier and the data is stored on a QStar tape target.*
      >
      >   To view tape target information, click **Tape Info**. The following information is displayed:
      >   - Target name and UUID
      >   - Barcode label of the media
      >   - Information on whether the media is online or offline
      >   - Location of the media

You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

- The Snapshot tier is not available for Azure Local environments and Hyper-V clusters.

c. In the New VM name field, specify a new name for the virtual machine.

> **Important** *For Azure Local environments and Hyper-V clusters:* The name that you specify must not contain spaces.

d. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

   If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

   - In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
   - In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

     > **Note** The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

   - In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.

e. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore.

> **Important** Make sure to consider the following:
>
> - This option is disabled for virtual machines that have volume groups attached by using iSCSI. For details on what needs to be done before turning on the restored virtual machine, see "After cloning a virtual machine" on page 224.
> - *Only if you are cloning a virtual machine from a Nutanix cluster, a vSphere environment, an Azure Local environment, or a Hyper-V cluster.* If you turn the restored virtual machine on, the original virtual machine will be turned off automatically.
> - *Only if you are cloning a vSphere virtual machine to a vSphere environment and you have disabled the Power virtual machine on*

> *switch.* When you try to power on the virtual machine and you are prompted to answer whether the virtual machine has been moved or copied, make sure to answer **I Copied It**.

f. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.

g. *For volume groups attached to the virtual machine:* Use the **Clone volume groups** switch if you also want to restore the volume groups that are attached to the virtual machine.

8. In the Network section, do the following:

   a. Review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

      Depending on whether the original networks are available, proceed as follows:

      - If the original networks are available, you can leave the default values and clone the virtual machine with the original network settings, or you can modify the network settings.
      - If the original networks are not available, you must modify the network settings.

      Modifying network settings

Original networks are…	Instructions
Available	You can do the following: • Add a new network adapter by clicking **New** and selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting the virtual adapter, and then clicking **Edit** and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking

Original networks are...	Instructions
	Delete.
Unavailable	You can do the following: • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking **Delete**. • Add a new network adapter by clicking **New**, and then selecting the preferred network.

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

b. *Only if you are restoring the virtual machine to a different source.* Use the **Keep original MAC address** switch if you want the restored virtual machine to keep the original MAC address. Keep in mind that this is applicable only if at least one network adapter has a MAC address assigned.

9. Click **Restore**.

Cloning a virtual machine to a XenServer environment

Prerequisites

- *For virtual machines that you plan to clone to a new location:* The source to which you plan to clone the virtual machine must be added to HYCU. For details on how to do this, see "Adding sources" on page 73.
- *For Linux servers:* In the `/etc/fstab` system configuration file of the server, UUIDs (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`) must be used instead of device names for file system device identification.

Consideration

Only if ownership is set for the virtual machine. The same owner is automatically assigned to the restored virtual machine.

Recommendation

For Linux virtual machines: It is recommended that the use of persistent network device names based on MAC addresses is disabled. Otherwise, you will have to configure the network manually. For details on how to disable the use of persistent network device names, see your Linux distribution documentation.

Accessing the Virtual Machines panel

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

Procedure

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

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Virtual Machine Options**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination Source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. In the General section, do the following:
   a. From the Storage Container drop-down menu, select the storage container where you want to restore the virtual machine.

      > **Note** By default, the original storage container is selected. If you decide to restore the virtual machine to another storage container, keep in mind the following:
      >
      > - The restore from the Snapshot tier cannot be performed to another storage container.
      > - If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.

- If the selected storage container is on a different source, additional prerequisites apply. For details, see "Preparing for the restore to a different source" on page 163.

b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**

   Note *Only if you select the Archive tier and the data is stored on a QStar tape target.* Consider the following:
   To view tape target information, click **Tape Info**. The following information is displayed:
   - Target name and UUID
   - Barcode label of the media
   - Information on whether the media is online or offline
   - Location of the media

   You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

c. In the New VM Name field, specify a name for the restored virtual machine.

d. Enable the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

   If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:
   - In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
   - In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

     Note The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.

e. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore.

> **Important** If you turn the restored virtual machine on, the original virtual machine will be turned off automatically.

f. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.

8. In the Network section, do the following:

a. Review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

Depending on whether the original networks are available, proceed as follows:

- If the original networks are available, you can leave the default values and clone the virtual machine with the original network settings, or you can modify the network settings.

> **Note** The network name may differ from the one that is configured in XenCenter.

- If the original networks are not available, you must modify the network settings.

Modifying network settings

Original networks are...	Instructions
Available	You can do the following: • Add a new network adapter by clicking ✚ **New** and selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting the virtual adapter, and then clicking

Original networks are...	Instructions
	Edit and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking **Delete**.
Unavailable	You can do the following: • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking **Delete**. • Add a new network adapter by clicking **New**, and then selecting the preferred network.

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

b. *Only if you are restoring the virtual machine to a different source.* Use the **Keep original MAC address** switch if you want the restored virtual machine to keep the original MAC address. Keep in mind that this is applicable only if at least one network adapter has a MAC address assigned.

9. Click **Restore**.

Cloning a virtual machine to an AWS GovCloud (US) environment

Prerequisites

- *For virtual machines that you plan to restore to a new location:* The AWS GovCloud (US) region to which you plan to restore the virtual machine must be added to HYCU. For instructions, see "Adding an AWS GovCloud (US) region" on page 82.
- *Only if the virtual machine that you plan to restore resides on a source other than*

AWS GovCloud (US). A premium tier Platform license is required. For details, see "Licensing" on page 488.

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for restoring data to a new location.
- *For virtual machines that have BitLocker volumes encrypted with TPM-based keys:* Restoring such volumes is not supported.

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

Procedure

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

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Virtual Machine Options**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination Source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. From the AWS GovCloud (US) Account drop-down menu, select the account to which the virtual machine will be restored.

   The following information is displayed and preselected:

   - The account ID of the AWS GovCloud (US) account to which the virtual machine will be restored.
   - The region to which the virtual machine will be restored.
8. *Optional.* From the Key Pair Name drop-down menu, select the key pair name that you want to use for connection to the restored virtual machine.

   > **Important** *For Windows virtual machines:* The key pair name that you select can be used only if the EC2Config or EC2Launch service was

> configured on the original virtual machine or if you configure it later on the restored virtual machine.

9. From the Availability Zone drop-down menu, select the availability zone to which the virtual machine will be restored.
10. Click **Next**.
11. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
    - **Automatic**: Ensures the fastest restore to the latest state.
    - **Backup**
    - **Copy**
    - **Archive**
    - **Snapshot**
12. In the New VM Name field, specify a name for the restored virtual machine.
13. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

    If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

    - In the vCPU threads field, enter the number of CPUs for the restored virtual machine multiplied by the number of cores per CPU and the number of threads per core.
    - In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
    - From the Virtual machine type drop-down menu, select the virtual machine type.

      > Note The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

14. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
15. Under Network Interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network

interface from the Virtual Private Cloud (VPC) to which the original virtual machine belongs. If required, you can also modify network settings.

Modifying network settings

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

> **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same Virtual Private Cloud (VPC). The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add Network Interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow these steps:
  a. The Virtual Private Cloud (VPC) to which the network interface will be added is displayed and preselected.
  b. From the Subnets drop-down menu, select the subnet to which the network interface should be assigned.
  c. From the Security Groups drop-down menu, select one or more security groups that will be associated with the network interface. If you want to select all the available security groups, select **Select all**.
  d. Under Public Address Type, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Auto-assign	An automatically allocated public IP address will be assigned to the network interface on the restored virtual machine.
Elastic IP (Reserved)	An elastic public IP address that you reserved in AWS GovCloud (US) will be assigned to the network interface on the restored virtual machine.
Elastic IP	An elastic public IP address will be assigned to the

Option	Description
(New)	network interface on the restored virtual machine.

e. Under Private Address Type, select the private IP address for the network interface. You can select between the following options:

Option	Description
Auto-assign	An automatically allocated private IP address will be assigned to the network interface on the restored virtual machine.
Custom	A private IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

16. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
17. Under Operating System License, select one of the following options:

OS license option	Select this option if you want to...
Keep existing license	Keep the existing OS license on the restored virtual machine. **Important** Make sure that the existing license is applicable also in AWS.
Available only for the Windows Server OS. **Replace existing license with AWS license**	Replace the existing OS license with an AWS license on the restored virtual machine.

18. Click **Restore**.

Cloning a virtual machine to an Azure or Azure Government environment

Prerequisites

- *For virtual machines that you plan to restore to a new location:* The Azure or Azure Government subscription to which you plan to restore the virtual machine must be added to HYCU. For details on how to do this, see "Adding an Azure subscription" on page 83 or "Adding an Azure Government subscription" on page 84.
- *For virtual machines that have Azure Disk Encryption enabled:* The key vault must be available on the location to which you are restoring the virtual machine.
- *Only if the virtual machine that you plan to restore resides on a source other than Azure or Azure Government.* A premium tier Platform license is required. For details, see "Licensing" on page 488.

Limitation

If a restore point contains only a Snapshot tier, you cannot use it for restoring data to a new location.

Considerations

- If the resource validation fails during a restore (for example, if the resource is no longer available on the Azure marketplace or the offer was renamed), the restore completes with warnings. In this case, HYCU restores only disks to the specified resource group. You can use the restored virtual disks to manually create the virtual machine.
- *Only if you plan to keep the original network settings on the restored virtual machine.* If the original static IP address is still associated with the original or another virtual machine, a new random private IP address will be assigned to the restored virtual machine.

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

Procedure

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

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Virtual Machine Options**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination Source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. From the Service Principal drop-down menu, select the service principal that has access to the required resources (the source from which and to which you are restoring the virtual machine).
8. From the Subscription drop-down menu, select the subscription for the restored virtual machine.
9. From the Resource Group drop-down menu, select the resource group for the restored virtual machine.
10. From the Location drop-down menu, select the geographic region for the restored virtual machine.
11. From the Availability Zone drop-down menu, select the zone for the restored virtual machine.

    > **Note** The selected geographic region and the size of the virtual machine determine to which zones you can restore data. If you do not want to restore data to any zone, select **None**.

12. Click **Next**.
13. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
    - **Automatic**: Ensures the fastest restore to the latest state.
    - **Backup**
    - **Copy**

- **Archive**
- **Snapshot**

14. In the New VM Name field, specify a name for the restored virtual machine.
15. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

    If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

    - In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
    - In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
    - From the Virtual machine type drop-down menu, select the virtual machine type.

      > **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

16. Under Network Interfaces, you can view all the network interfaces that are attached to the virtual machine. Keep in mind that the IP address that is assigned to the primary IP configuration of the network interface will be assigned to the network interface on the restored virtual machine. If required, you can also modify network settings.

    ### Modifying network settings

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

    > **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on the selected virtual machine type.

    Depending on how you want to modify network settings, do one of the following:

    - Click **Add Network Interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow

these steps:

a. *Only if you are adding a network interface.* From the Virtual Network drop-down menu, select the network for the network interface.

> **Note** The list of available networks includes only the ones within the region you selected for the restored virtual machine.

b. From the Subnet drop-down menu, select the subnet to which the network interface should be assigned.

c. Under NIC Network Security Group, select the network security group for the network interface. You can select among the following options:

Option	Description
None	The network interface will not be assigned to a network security group.
Basic	The network interface will be assigned to Azure's basic network security group.
Advanced	The network interface will be assigned to the network security group that you select from the drop-down menu. By default, the network security group of the original virtual machine is selected.

d. Under Public IP Address Type, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Static	A static IP address will be assigned to the network interface on the restored virtual machine.
Existing	A preferred public IP address resource that you have created in Azure or Azure Government will be assigned to the network interface on the restored virtual machine.

e. Under Private IP Address Type, select the private IP address for the

network interface. You can select between the following options:

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

17. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
18. Click **Restore**.

After cloning a virtual machine

After cloning a virtual machine, go through considerations and recommendations listed in this section to make sure that everything works as expected.

Considerations

- If a new MAC address is assigned to a network adapter on the cloned virtual machine, make sure that the guest operating system is configured appropriately to connect the cloned virtual machine to the selected network.
- If after cloning a virtual machine from a Nutanix AHV cluster to a Nutanix ESXi cluster or a vSphere environment, the virtual machine does not turn on due to an IDE device not being configured properly, you must edit the IDE device configuration manually. For details on how to do this, see VMware documentation.
- *For vSphere environments:* Some operating systems (for example, Red Hat Enterprise Linux 7) might require network configuration. For details, see VMware documentation.

- *For virtual machines to which volume groups are attached by using iSCSI:* Because the original virtual machine and the restored one have the same network and iSCSI configuration settings after the restore, make sure both the virtual machines are not turned on at the same time to avoid any potential issues. As one way of preventing issues, you can disconnect the restored virtual machine from the network before turning it on and make the required changes such as replacing the network adapter and updating the iSCSI settings on it.
- *For servers:*
  - *Only if you cloned a Windows server to a Nutanix ESXi cluster.* Make sure to modify the machine configuration after the restore by specifying the appropriate guest OS and to install the latest version of VMware Tools on the machine. For detailed information, see VMware documentation.
  - *Only if you cloned a Linux server that uses UEFI firmware to a Nutanix AHV cluster.* If the virtual machine does not boot after the restore, reboot the machine.

Recommendations

- *For Azure Local environments and Hyper-V clusters:* In Windows Admin Center, check the network settings of the cloned virtual machine and make sure that the isolation mode and the VLAN identifier match the network configuration of the cluster to which you cloned the virtual machine.
- *For Linux servers:* Because the original boot loader of the server is replaced with a temporary one during the backup, it is recommended to update the boot configuration after the restore. Depending on what firmware the server uses, see one of the following sections for instructions on how to do this:
  - “Updating the boot configuration of Linux servers that use BIOS firmware” below
  - “Updating the boot configuration of Linux servers that use UEFI firmware” on the next page

Updating the boot configuration of Linux servers that use BIOS firmware

Procedure

1. In the `/etc/default/grub` system configuration file, do the following:
   a. Edit the `GRUB_CMDLINE_LINUX` option and remove the following kernel parameters (if present):

- `rd.lvm.` (except `rd.lvm=0`)
- `rd.md.` (except `rd.md=0`)
- `rd.dm.` (except `rd.dm=0`)
- `rd.luks.`

b. Set the resume device on the virtual machine to match the resume device UUID on the original server. For example, if the resume device on the original server is `resume=/dev/mapper/cl-swap`, the resume device on the virtual machine should be `resume=UUID=4044243b-612b-42bc-ba22-4736c4eadde6`.

2. *Optional.* If you want to speed up the boot process and skip mounting non-existent volumes, in the `/etc/fstab` system configuration file, comment all the lines for volumes for which a warning was triggered at backup time.

   Example
   The following warning message was triggered:

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

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

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

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

4. Install the boot loader on the boot disk by running the following command:

   ```
   grub2-install /dev/sdc
   ```

   > **Tip** The boot disk is the one that contains the boot partition. To identify the boot partition, run the following command:
   >
   > ```
   > findmnt -nT /boot -o SOURCE
   > ```

5. Reboot the virtual machine.

Updating the boot configuration of Linux servers that use UEFI firmware

Procedure

1. *Only if you cloned the server to a Nutanix ESXi cluster or a vSphere environment.* When the virtual machine enters the firmware setup mode, select the **Boot from file** option, and then specify the

`<EFIPartition>/EFI/hycu/shimx64.efi` file. For details, see Nutanix or VMware documentation.

2. In the `/etc/default/grub` system configuration file, do the following:
   a. Edit the `GRUB_CMDLINE_LINUX` option and remove the following kernel parameters (if present):
      - `rd.lvm.` (except `rd.lvm=0`)
      - `rd.md.` (except `rd.md=0`)
      - `rd.dm.` (except `rd.dm=0`)
      - `rd.luks.`
   b. Set the resume device on the virtual machine to match the resume device UUID on the original server. For example, if the resume device on the original server is `resume=/dev/mapper/cl-swap`, the resume device on the virtual machine should be `resume=UUID=4044243b-612b-42bc-ba22-4736c4eadde6`.
3. *Optional.* If you want to speed up the boot process and skip mounting non-existent volumes, in the `/etc/fstab` system configuration file, comment all the lines for volumes for which a warning message was triggered at backup time.

   Example
   The following warning message was triggered:

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

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

4. Update the GRUB configuration by running the following command:
   - For Red Hat Enterprise Linux and Oracle Linux:

     ```
     grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
     ```

   - For CentOS:

     ```
     grub2-mkconfig -o /boot/efi/EFI/centos/grub.cfg
     ```

5. Reboot the virtual machine.
6. *Only if secure boot was enabled on the original server and you use third-party kernel modules.* Enroll the Machine Owner Key (MOK) used to sign third-party kernel modules. For details on how to do this, see the respective operating system documentation.

7. Create the default boot entry in the UEFI firmware setup. The boot entry should point to the following system file:
   - For Red Hat Enterprise Linux and Oracle Linux:

     ```
     <EFIPartition>/EFI/redhat/shimx64.efi
     ```

   - For CentOS:

     ```
     <EFIPartition>/EFI/centos/shimx64.efi
     ```

Validating the virtual machine backup

You can validate the virtual machine backup by creating a virtual machine clone. In this case, the original virtual machine will not be overwritten and turned off. You can also specify whether you want to keep the virtual machine clone after the backup validation is performed.

> **Note** You can also set up a validation policy and schedule the backup validation according to the values that you define in your validation policy. For details on how to do this, see "Setting up a validation policy" on page 425.

Prerequisites

- If you are cloning the virtual machine to a vSphere environment, the latest version of VMware Tools must be installed on the virtual machine.
- *Only if you plan to specify the Advanced validation type.*
  - Credentials must be assigned to the virtual machine. For prerequisites, limitations, considerations, and instructions, see "Enabling access to application data" on page 252.
  - A network card must be added to the virtual machine.

Limitation

Performing the backup validation is not supported for the following:

- The HYCU backup controller and virtual machines running in AWS GovCloud (US), Azure, or Azure Government environments.
- Virtual machines that have volume groups attached by using iSCSI.

Considerations

- Network conflicts may occur during the backup validation if the virtual machine is configured with a static IP address, resulting in unreliable backup validation data.
- *Only if you plan to specify the Advanced validation type when performing the backup validation for a Windows virtual machine.* Checking for disk errors may fail in some cases, which does not mean that your virtual machine is corrupted. However, it is highly recommended that you check the status of such a virtual machine manually.
- After you perform the backup validation, consider the following:
  - You can view the backup validation status of a virtual machine in the Validation column in the Virtual Machines panel (represented by an icon). By pausing on the icon, you can also see the assigned validation policy, if you have set it up, and detailed information about the last performed backup validation.
  - The Exclude policy is automatically assigned to the cloned virtual machine.

Procedure

1. In the Virtual Machines panel, click the virtual machine for which you want to perform the backup validation. The Detail view appears at the bottom of the screen.

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Virtual machine options**, and then click **Next**.
5. Select **Validate VM backup**, and then click **Next**.
6. From the Storage container drop-down menu, select where you want to clone the virtual machine for which you are performing the backup validation.
7. From the Restore from drop-down menu, select which tier you want to use for the backup validation. Your restore point can contain one or more tiers among which you can select:

- **Automatic**
- **Backup**
- **Copy**
- **Archive**
- **Snapshot**

> **Note** When selecting the restore point tier, consider the following:
>
> - If you select Automatic, the tier for the backup validation is by default selected in the following priority order: Backup > Copy > Archive > Snapshot. This means that HYCU will always use the first available tier in the specified order for the backup validation. However, you can at any time change this default behavior by customizing the `backup.validation.restore.source.priority.order` configuration setting in the HYCU `config.properties` file and adjusting the tier order to your data protection needs. For details on how to customize HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
> - *Only if you select the Archive tier and the data is stored on a QStar tape target.*
>
> To view tape target information, click **Tape Info**. The following information is displayed:
> - Target name and UUID
> - Barcode label of the media
> - Information on whether the media is online or offline
> - Location of the media
>
> You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

8. In the New VM name field, specify a name for the cloned virtual machine.
9. Use the **Use original VM configuration** switch if you want the cloned virtual machine to have the same configuration settings as the original virtual machine.

   If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

   - In the vCPU(s) field, enter the number of virtual CPUs for the cloned virtual machine.
   - In the Cores per vCPU field, enter the number of cores per virtual CPU for the cloned virtual machine.

> **Note** The total number of cores of the cloned virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

- In the Memory field, set the amount of memory (in GiB or MiB) for the cloned virtual machine.

10. From the Keep VM after validation drop-down menu, depending on whether you want to keep the virtual machine after the backup validation is performed, select one of the following options:

Option	Description
Always	The virtual machine will be kept after the backup validation is performed.
On validation error	The virtual machine will be kept after the backup validation is performed only if a validation error occurs during the validation.
Never	The virtual machine will be automatically deleted after the backup validation is performed.

11. From the Validation type drop-down menu, select one of the following types:

Validation type	Description
Basic	During the backup validation, the following tasks will be performed: • The virtual machine will be cloned and turned on. • The guest OS will be shut down.
Advanced	During the backup validation, the following tasks will be performed: • The virtual machine will be cloned and turned on. • Any applications running on the virtual machine will be discovered. • Virtual disks will be validated, which includes checking the virtual machine file system and existing disks on the virtual machine. For Windows virtual machines,

Validation type	Description
	checking for disk errors is also performed. • The custom scripts will be run, if specified. • The guest OS will be shut down.

12. *Only if you selected the Advanced validation type.* Do the following:
    a. Enable the **Run custom script** switch if you want the custom script to be run on the virtual machine as part of the backup validation process, and then make sure that the proper path to the script is specified.

    > **Note** The script returns an exit code of 0 for success and any other value for failure.

    b. In the Network section, review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

    Depending on whether the original networks are available, proceed as follows:

    - If the original networks are available, you can leave the default values and clone the virtual machine with the original network settings, or you can modify the network settings.
    - If the original networks are not available, you must modify the network settings.

    Modifying network settings

Original networks are...	Instructions
Available	You can do the following: • Add a new network adapter by clicking **New**, and then selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting it, and then clicking **Edit** and selecting the preferred network.

Original networks are...	Instructions
	• Delete the network adapter you do not need anymore by selecting it, and then clicking **Delete**.
Unavailable	You can do the following: • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking **Delete**. • Add a new network adapter by clicking **New**, and then selecting the preferred network.

13. Click **Validate**.

Restoring virtual disks

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

Prerequisite

For Azure Local environments and Hyper-V clusters: Make sure no snapshots are present for the virtual machine whose disks you plan to restore.

Limitations

- Restoring server disks by using the Restore vDisks option is not supported.
- For virtual machines running in AWS GovCloud (US), Azure, or Azure Government environments, restoring virtual disks is not supported.
- *Only if you plan to restore vSphere virtual disks.* The virtual disks that you plan to restore must be available on the original virtual machine. If the virtual disks are not available, the restore will fail.
- You cannot perform a granular restore of NDB-managed database server virtual disks if the disk layout changes after the selected restore point is created.

Considerations

- If any virtual disks were excluded from the backup, you cannot select them for the restore. The corresponding restore point labels are marked with a red circle. For details, see "Viewing entity details" on page 399.
- The original virtual disks are deleted and the restored ones are automatically attached to all virtual machines to which they were attached at backup time.
- *Only if restoring volume groups attached to the virtual machine.* The virtual machines to which the volume groups are attached must be turned off.
- If you plan to restore vSphere virtual disks, keep in mind that the virtual machine will be powered off during restore.

Accessing the Virtual Machines panel

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

Procedure

1. In the Virtual Machines panel, click the virtual machine whose virtual disks you want to restore.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

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

3. Click **Restore**.
4. Select **Disk options**, and then click **Next**.
5. Select **Restore vDisks**, and then click **Next**.
6. From the list of virtual disks that are available for the restore, select the ones that you want to restore, and then click **Next**.

   > **Important** *Only if restoring volume groups attached to the virtual machine.* You cannot select individual disks, but only the whole volume group.

7. From the Storage container drop-down menu, select where you want to restore the virtual disks.

   > **Note** By default, the original storage container is selected. If you decide to restore the virtual disks to another storage container, they will

> not be restored from the snapshot, but from the target. Therefore, no fast restore will be performed.

8. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: This type of restore ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot** *(Nutanix only)*

   > Note *Only if you select the Archive tier and the data is stored on a QStar tape target.*
   > To view tape target information, click **Tape Info**. The following information is displayed:
   >
   > - Target name and UUID
   > - Barcode label of the media
   > - Information on whether the media is online or offline
   > - Location of the media
   >
   > You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

9. *Only if you are restoring disks of an NDB-managed database server, for Time Machines with a Continuous SLA.*
   Enable the **Run a Tail Log Backup on NDB before restore** switch to ensure that an NDB Tail Log Backup is started before a database restore.

   > Note A Tail Log Backup is started automatically after a restore.

10. Click **Restore**.

Cloning virtual disks

You can create clones of virtual disks by restoring them to their original or a new location. In this case, the original virtual disks will not be overwritten.

Prerequisites

- *For vSphere environments:* Make sure the number of disks on the virtual machine to which you are restoring data does not exceed 15 on SCSI controller 0.
- *For Azure Local environments and Hyper-V clusters:* Make sure no snapshots are present for the virtual machine whose disks you plan to restore.

Limitations

- For virtual machines running in AWS GovCloud (US), Azure, or Azure Government environments, restoring virtual disks is not supported.
- Restoring server disks by using the Clone vDisks option is not supported.
- Restoring virtual disks to the original virtual machine is not possible if the original virtual machine is powered off.
- Restoring virtual disks to a virtual machine running on a different source is not supported.

Considerations

- If any virtual disks are excluded from backup, you cannot select them for restore. The corresponding restore point labels are marked with a red circle. For details, see "Viewing entity details" on page 399.
- *Only if restoring volume groups attached to the virtual machine.*
  - The original volume groups are kept alongside of the restored ones and the following applies regarding their attachment:
    - If the virtual machine to which you are restoring the volume groups is the original one, the volume groups will be attached to all the virtual machines to which they were attached at backup time.
    - If the virtual machine to which you are restoring the volume groups is not the original one and is running on a Nutanix AHV cluster, the volume groups will be automatically attached to the selected virtual machine.
    - If the virtual machine to which you are restoring the volume groups is not the original one and is running on a Nutanix ESXi cluster, you must manually attach the volume groups to the selected virtual machine after the restore.
  - The name format of the cloned volume groups is as follows:

    `<OriginalVGName>-<Timestamp>`
- *For virtual machine disks:*

- The original virtual machine disks are kept alongside the restored ones that are automatically attached to the virtual machine as the first available interface index (per interface type). For example, if you have the `scsi.0`, `scsi.1`, and `scsi.4` virtual disks already attached to your virtual machine, the restored one will be `scsi.2`.
- If the bus type of the original virtual disks is IDE, it is automatically changed to SCSI during the restore.

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

Procedure

1. In the Virtual Machines panel, click the virtual machine whose virtual disks you want to restore.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

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

3. Click **Restore**.
4. Select **Disk options**, and then click **Next**.
5. Select **Clone vDisks**, and then click **Next**.
6. From the list of virtual disks that are available for the restore, select the ones that you want to restore, and then click **Next**.

   > **Important** *Only if restoring volume groups attached to the virtual machine.* You cannot select individual disks, but only the whole volume group.

7. From the Select VM drop-down menu, select the virtual machine to which you want to attach the restored virtual disks. The restored virtual disks can be attached to the original virtual machine (the default selection) or any other virtual machine. Consider the following:
   - If you are attaching the virtual disks to the original virtual machine, make sure it is turned on.
   - You cannot attach the restored disks to a server.

8. From the Storage container drop-down menu, select where you want to restore the virtual disks.

   > **Note** *For virtual machines:* You can select only among the storage containers that are created on the Nutanix cluster on which the selected virtual machine resides.

9. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: This type of restore ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot** *(Nutanix clusters only)*

   > **Note** *Only if you select the Archive tier and the data is stored on a QStar tape target.*
   > To view tape target information, click **Tape Info**. The following information is displayed:
   >
   > - Target name and UUID
   > - Barcode label of the media
   > - Information on whether the media is online or offline
   > - Location of the media
   >
   > You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

10. Click **Restore**.

Exporting virtual disks

You can restore virtual disks to an NFS or SMB share. You can use exported virtual disks to restore data to a server. For details, see "Restoring data to a server" on page 241.

Prerequisites

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

- *Only if you are restoring data that is stored in the archive access tier on an Azure target.* You must recreate a snapshot and use this snapshot for restoring data, or manually rehydrate data. For details on how to recreate a snapshot, see "Recreating snapshots" on page 436. For details on how to manually rehydrate data, see Azure documentation.

Consideration

If any virtual disks were excluded from the backup, you cannot select them for the restore. The corresponding restore point labels are marked with a red circle. For details, see "Viewing entity details" on page 399.

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

Procedure

1. In the Virtual Machines panel, click the virtual machine whose virtual disks you want to restore.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

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

3. Click **Restore**.
4. Select **Disk options**, and then click **Next**.
5. Select **Export vDisks**, and then click **Next**.

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

6. From the list of virtual disks that are available for the restore, select the ones that you want to restore, and then click **Next**.
7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: This type of restore ensures the fastest restore to the latest state.
   - **Backup**

- **Copy**
- **Archive**
- **Snapshot** *(Nutanix clusters only)*

> **Note** *Only if you select the Archive tier and the data is stored on a QStar tape target.*
> To view tape target information, click **Tape Info**. The following information is displayed:
>
> - Target name and UUID
> - Barcode label of the media
> - Information on whether the media is online or offline
> - Location of the media
>
> You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

8. From the Type drop-down menu, select where you want to restore the virtual disks, and then provide the required information:
   - **SMB**
     a. *Optional.* Enter the domain and user credentials.
     b. Enter the SMB server host and the path to the SMB shared folder from the root of the server (for example, `/backups/HYCU`).

     > **Important** To ensure that HYCU can use Kerberos authentication, you must enter the fully qualified domain name (FQDN) for the SMB server host.

   - **NFS**

     Enter the NFS server name or IP address and the path to the NFS shared folder from the root of the server (for example, `/backups/HYCU`).
9. Click **Restore**.

After exporting virtual disks

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

Data is restored to the following location:

`/<SharedPath>/<VMName>/<Timestamp>/<Filename>`

In this instance, <*SharedPath*> is the path to the shared folder, <*VMName*> is the virtual machine name, <*Timestamp*> is the time of the restore, and <*Filename*>is the virtual machine disk UUID.

What kind of files are created by the restore depends on the environment in which the virtual machine whose virtual disks you restored was backed up. Depending on the type of source in your environment, the following files are created for each selected disk:

Source	Files
Nutanix AHV	<*DiskName*> (without extensions) A raw image of the disk, including unallocated space as zeroes.
Nutanix ESXi	
vSphere	• <*DiskName*>-flat.vmdk A raw image of the disk • <*DiskName*>.vmdk A VMDK descriptor file, referencing <*DiskName*>-flat.vmdk
XenServer	<*DiskID*> (without extensions) A raw image of the disk.
Azure Local	
Hyper-V	

Restoring data to a server

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

Prerequisites

- The server to which you want to restore data must have the same number of disks as the original machine and the disk size must be equal to or greater than the original size.
- Make sure you downloaded a Linux live CD (for example, Ubuntu) and booted it on the server where you want to restore your data.

Considerations

- Make sure you run all the commands as root.
- You can safely ignore the following error message:

```
The backup GPT table is corrupt, but the primary appears OK, so
that will be used.
```

Procedure

1. Identify your destination disk.

   Because HYCU performs the backup at the disk level, you must identify the path of each disk to which you will restore data. To list all the disks on your system, run the following command:

   ```
   fdisk -l
   ```

   The following is an example of the output:

   ```
   Disk /dev/sda: 32 GiB, 34359738368 bytes, 67108864 sectors
   Disk /dev/sdb: 5 GiB, 5368709120 bytes, 10485760 sectors
   ```

2. Mount the share to which you exported the disks.
3. Identify the path to the exported disks on the mounted share by running the following commands:

   ```
   cd /<SharedPath>/<VMName>/<Timestamp>
   ```

   ```
   ls
   ```

   The following is an example of the output:

   ```
   ServerDisk0 ServerDisk1
   ```

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

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

   For example:

   ```
   fdisk -l ServerDisk0
   ```

   The information about the exported disk (for example, disk size and a list of partitions) is displayed. Use this information to identify a suitable destination disk for restoring the data. For example, the size of exported disk `ServerDisk0` matches the size of disk `/dev/sda`. Therefore, disk `ServerDisk0` can be restored to disk `/dev/sda`.

   The following is an example of the output:

   ```
   Disk ServerDisk0: 32 GiB, 34359738368 bytes, 67108864 sectors
   Units: sectors of 1 * 512 = 512 bytes
   Sector size (logical/physical): 512 bytes / 512 bytes
   ```

```
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x36bab260
Device Boot Start End Sectors Size Id Type
ServerDisk0p1 * 2048 718847 716800 350M 7
HPFS/NTFS/exFAT
ServerDisk0p2 718848 67106815 66387968 31.7G 7
HPFS/NTFS/exFAT
```

5. Restore data by running the following command for each disk:

   ```
   dd if=<ExportedDiskName> of=<DestinationDiskPath> bs=1024
   status=progress
   ```

   For example:

   ```
   dd if=ServerDisk0 of=/dev/sda bs=1024k status=progress
   ```

   The following is an example of the output:

   ```
   33540483072 bytes (34 GB, 31 GiB) copied, 229 s, 146 MB/s
   33554432+0 records in
   33554432+0 records out
   34359738368 bytes (34 GB, 32 GiB) copied, 229.78 s, 150 MB/s
   ```

6. Eject the Linux live CD and reboot the server.

Restoring individual files

You can restore individual files to the same or a different virtual machine, to an SMB or NFS share, or to the local machine. This alternative to restoring an entire virtual machine allows you to restore only one or more files that have become corrupted or have been deleted for some reason and are now missing on the virtual machine.

Individual files can be restored from a target or a snapshot. A restore is always performed from the snapshot if the snapshot is available for the selected restore point (this speeds up the restore process). Otherwise, the restore is performed from the target if Target is selected as the backup target type in your policy (this saves space in your environment). If you want to restore individual files from a snapshot and no snapshot is available for the selected virtual

machine restore point, HYCU enables you to manually recreate it. For details on how to do this, see “Recreating snapshots” on page 436.

You can use the pre-restore and post-restore scripts to perform necessary actions before and after the restore of individual files is performed. For details on how to specify the scripts, follow the procedure described in this section. For details on exit codes and exported environment variables, see “Using the pre and post scripts” on page 561.

Prerequisites

Windows virtual machines	• The NTFS, FAT, or FAT32 file system must be used. • For improved restore performance, the startup type of the Microsoft iSCSI Initiator Service may not be set to `Disabled`. • *For restoring files to a virtual machine:* ◦ *For Windows 8 and 10 virtual machines:* WinRM must be enabled and configured by using the `winrm quickconfig` command. ◦ A Windows operating system user account must exist. This account must have WinRM permissions granted and must be a member of the virtual machine's local Administrators group. ◦ Access to the virtual machine file system must be enabled. For instructions, see “Enabling access to data” on page 174. ◦ *For pre/post-restore scripts:* A script must be available in the accessible folder and must have one of the following extensions: `bat`, `ps1`, `cmd`. ◦ *For AWS GovCloud (US), Azure, and Azure Government environments:* The virtual machine to which you plan to restore files must be in the same virtual network as the HYCU backup controller.
Linux virtual machines	• The FAT32, xfs, ext4/ext3/ext2, reiserfs, or btrfs file system must be used. • *For restoring individual system files with the non-root sudo user and better overall restore performance:* ◦ The sudo user must have the NOPASSWD option set.

	For example, to set the NOPASSWD option for the user on a Red Hat Enterprise Linux 8.x system, add the following line to the `/etc/sudoers` file: `restoreuser ALL=(ALL) NOPASSWD: ALL` ◦ The `cifs-utils` package must be installed on virtual machines whose files you plan to restore. ◦ *For AWS GovCloud (US), Azure, and Azure Government environments:* The virtual machine to which you plan to restore files must be in the same virtual network as the HYCU backup controller. • References in the `/etc/fstab` system configuration file entries must use universally unique identifiers (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`) rather than device names (for example, `/dev/sda1`) unless the entries refer to logical volumes (for example, `/dev/mapper/ol-root`). • *For restoring files to a virtual machine:* ◦ Access to the virtual machines through ssh must be enabled. ◦ Access to the virtual machine file system must be enabled. For instructions, see "Enabling access to data" on page 174. ◦ *For pre/post-restore scripts:* A script must be available in the accessible folder and must have the `sh` extension. You must have permissions to run the script on the virtual machine with the assigned credentials.
Nutanix ESXi clusters	• *For restoring files to a virtual machine:* The latest versions of VMware Tools and NGT must be 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. • You must have the required restore privileges assigned. For details, see "Assigning privileges to a vSphere user" on page 552.

vSphere environments	• You must have the required restore privileges assigned. For details, see "Assigning privileges to a vSphere user" on page 552.
XenServer environments	• XenServer VM Tools must be installed on the client virtual machine. For detailed information about installing XenServer VM Tools, see XenServer documentation.
All environments	• *Only if you are restoring data that is stored in the archive access tier on an Azure target.* You must recreate a snapshot and use this snapshot for restoring data, or manually rehydrate data. For details on how to recreate a snapshot, see "Recreating snapshots" on page 436. For details on how to manually rehydrate data, see Azure documentation.

Limitations

- Restoring individual files on multi-boot systems is not supported.
- Restoring individual files is not supported for virtual machines with encrypted disks, folders, or files.
- Restoring individual file data from tape is not supported.
- On Linux, you can restore symbolic links and soft links only to the original location.
- Restoring files from the same snapshot simultaneously by two different users is not possible.
- You cannot restore individual files if you excluded all virtual machine disks from the backup and left only the attached volume groups.
- File browsing is supported if the file names contain only characters in the Unicode Basic Multilingual Plane (BMP).
- *For Linux virtual machines:* Restoring individual files is supported if the file names contain only characters in the Unicode Basic Multilingual Plane (BMP).
- *For restoring files to a different virtual machine:* You can restore files only to a virtual machine that belongs to the same operating system family as the original one.
- *For restoring files to a local machine:* You can download only a data archive whose size is less than or equal to 2 GiB.

- *For Windows virtual machines running on a Nutanix cluster that have Storage Replica enabled:* Restoring individual files to a virtual machine is supported only if the restore is performed from the target.
- *For Azure and Azure Government environments:* You cannot restore individual files if Azure Disk Encryption is enabled on the virtual machine.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- You cannot perform a restore of a virtual machine whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up` configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
- *For restoring files to a virtual machine:* To be able to restore some types of files (for example, system files), the account you specify to access a virtual machine must be a member of the virtual machine's local Administrators group on Windows or have root permissions on Linux.
- If any virtual disks are excluded from backup, you cannot select them for restore. The corresponding restore point labels are marked with a red circle. For details, see "Viewing entity details" on page 399.
- *For using the Backup from replica option:* When restoring to the central or remote site (the original location), the restore is always performed from the snapshot on the central site.
- *For pre/post-restore scripts:* You can specify pre/post-restore scripts only when restoring files to a virtual machine.
- *Only if restoring files to an external distributed SMB share.* Make sure that the folder for the restore is created on the share and the shared path leads to this folder.

Consideration when restoring symbolic links

Depending on the file or folder and the location that you select for the restore, symbolic links and their targets are restored in the following ways:

Selected item	Restore location	Restored resource
Symbolic link (file or folder)	Original location on the original or a different virtual machine	Symbolic link and its target
	• Alternate location on the original virtual machine • Alternate location on a different virtual machine • External share or download	Symbolic links' target
Folder that contains symbolic links	Any location on the original or a different virtual machine	Folder that contains symbolic links
	External share and download	Folder that contains symbolic links' targets

Recommendation

Only if restoring a large number of files. Instead of restoring individual files, it is highly recommended to restore disks hosting these files by using the Clone vDisks option. For instructions, see "Cloning virtual disks" on page 235.

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

Procedure

1. In the Virtual Machines panel, click the virtual machine that contains the files that you want to restore to open the Detail view.

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

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

3. Click **Restore Files**.
4. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**

   > Note The Snapshot tier is not available for Azure Local environments and Hyper-V clusters.

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

   > Tip If there are too many files to be displayed on one page, you can move between the pages by clicking > and <.
   > You can also search for a file or a folder by entering its name and then pressing **Enter** in the Search field.

7. Depending on where you want to restore the selected files (to the same or a different virtual machine, an external SMB or NFS share, or the local machine), select the preferred restore option, click **Next**, and then follow the instructions:

Restore option	Instructions
Restore to virtual machine	a. On the General tab, do the following: i. From the Virtual machine drop-down menu, select the virtual machine to which you want to restore the files. You can restore the files to the same or a different virtual machine. ii. Select whether you want to restore the files to the original location or an alternate location. If you select an alternate location, specify the path in the format that is supported by your operating system. iii. Specify which action should be performed during the restore operation if a file with the

Restore option	Instructions
	same name already exists in the selected location (rename the original file, rename the restored file, skip the file, or overwrite the file). iv. Use the **Restore ACL** switch if you want to restore the original access control list. > ⓘ Important If the virtual machine is not accessible due to various reasons (for example, credentials are not assigned to it, discovery was not successful, or it is turned off or deleted from the source), you cannot select it for restoring the individual files. b. *Optional.* Click the **Pre/Post Scripts** tab, and then do the following: i. In the Run pre-restore script field, enter the path to the script that HYCU will run before the restore is performed. ii. In the Run post-restore script, enter the path to the script that HYCU will run after the restore is performed. c. Click **Restore**.
Restore to external share	a. Select **NFS** or **SMB** for the share type, and then specify the path to a shared folder in the following format: `\\server\<Path>` b. *For SMB:* Optionally, provide user credentials to access the SMB share. c. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file). d. Click **Restore**.

Restore option	Instructions
Download	Click **Download** to restore the selected files to the local machine. > ⓘ Important Do not refresh the page or navigate away from the page until the download process job finishes.

Chapter 5

Protecting applications

HYCU enables you to protect your application data with fast and reliable backup and restore operations. After you enable HYCU to access an application running on a virtual machine, complete the required preparatory steps, and back up the application, you can choose to restore either the whole application or only specific application items.

> **Note** The instructions for protecting applications residing on virtual machines apply also to applications residing on servers except where specifically stated otherwise.

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

- "Enabling access to application data" below
- "Planning application protection" on page 256
- "Backing up applications" on page 270
- "Restoring whole applications" on page 272
- "Restoring SQL Server databases" on page 313
- "Restoring Exchange Server databases, mailboxes, and public folders" on page 317
- "Restoring Oracle database instances and tablespaces" on page 321
- "Restoring PostgreSQL database clusters" on page 324

Enabling access to application data

After you assign credentials to virtual machines as described in "Enabling access to data" on page 174, the process of application discovery starts automatically.

When the application discovery job completes, the discovered applications are listed in the Applications panel. HYCU supports different types of applications

on virtual machines and servers. For a list of supported applications, see the *HYCU Compatibility Matrix*.

Depending on the Discovery status of the applications that you want to protect, do one of the following:

✅	HYCU can access the discovered applications that you want to protect with the virtual machine credentials and you can start protecting such applications. For instructions, see "Backing up applications" on page 270. **Note** Access to Active Directory and SAP HANA is always granted with the virtual machine credentials.
❌	The virtual machine credentials do not have proper permissions and HYCU cannot access applications. To enable HYCU to access applications, do one of the following: • If you want to use virtual machine credentials, reassign credentials to virtual machines so that they have proper permissions. For instructions on how to assign credentials to a virtual machine, see "Enabling access to data" on page 174. • If you want to use application-specific credentials, follow the procedure described in this section.

Prerequisites

Windows virtual machines	• *For Windows 8 and 10:* WinRM must be enabled and configured by using the `winrm quickconfig` command. • A Windows user account with WinRM permissions must exist. This account should have access to the application and be a member of the virtual machine's local Administrators group. • Access to the virtual machine file system must be enabled. For instructions, see "Enabling access to data" on page 174 • To ensure that HYCU can use Kerberos authentication when accessing data or running the application discovery process on virtual machines by using WinRM connection, your DNS server must be configured to perform reverse DNS lookups.

Linux virtual machines	• Access to the virtual machines through SSH must be enabled. • Access to the virtual machine file system must be enabled. For instructions, see "Enabling access to data" on page 174
Nutanix ESXi clusters	VMware Tools and NGT must be 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.

Application-specific prerequisites

SQL Server	• Access should be enabled on all virtual machines where the SQL Server failover cluster and SQL Server Always On Availability Group instance resides. • *For SQL Server Always On Availability Group:* An availability group must be created by using automatic seeding.
Oracle	• The OS user must have sudo privileges and the `NOPASSWD` option set.
PostgreSQL	• The OS user must have sudo privileges and the `NOPASSWD` option set. • The database user must have permissions to execute the following PostgreSQL functions: `pg_tablespace_location`, `pg_backup_start`, `pg_backup_stop`, `pg_switch_wal`, and `pg_promote`. • The `pg_hba.conf` file must be configured to allow the database user to connect to the database cluster from the local host. • *For standby database clusters:* ◦ The database user must be able to connect to the primary database cluster from the standby database cluster system. ◦ If you want HYCU to automatically retrieve the primary database cluster host and port, the database user must have the `pg_read_all_settings` role granted on the database cluster.

Considerations

- *For an Oracle application*: When an operating system is used to authenticate Oracle database users, the Oracle database can be accessed with the OS user credentials, which allows you to skip the procedure of providing access to application data. To enable such authentication mode, contact the Oracle database administrator.
- *For PostgreSQL:*
  - The application discovery process discovers only the PostgreSQL database clusters that were started with the `-D` parameter and are running.
  - HYCU supports peer and password-based authentication for the connection to the database cluster. The type of authentication for the database user is defined in the `pg_hba.conf` file:
    - Peer authentication if the virtual machine credentials are used for the connection.
    - Password-based authentication if the application-specific credentials are used for the connection.
- *For Windows virtual machines*: To ensure that HYCU can use Kerberos authentication, the virtual machine must use the credentials of an Active Directory domain account. If not, the legacy NTLM authentication will be used.

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

Procedure

1. In the Applications panel, select the applications that you want to protect.
2. Click **Configuration**. The Configuration dialog box opens.
3. On the Credentials tab, depending on the credentials that you want to use, do one of the following:
   - If you want to use virtual machine credentials, click **Save**.
   - If you want to use the application-specific credentials, do the following:
     a. Disable the **Use VM credentials with access to the application** switch.

b. Enter credentials for a user account with required permissions and access to the applications. Make sure the following requirements are met:
   - *For applications running on Windows virtual machines:* The specified account must be a member of the virtual machine's local Administrators group.
   - *For SQL Server:* The specified account must have the sysadmin role on the SQL Server application instance. The SQL Server account that connects by using SQL Server Authentication is not supported.
   - *For Exchange Server:* The specified account must be a member of the View-Only Organization Management role group, and it must have the Databases and Disaster Recovery roles assigned.

c. Click **Save**.

A new process of application discovery is started with the modified credentials for all virtual machines that have these credentials assigned. After this is done, the status of your applications should be ✅ and you can continue with protecting application data as described in "Backing up applications" on page 270.

You can later unassign the credentials from a virtual machine by clicking **Unassign** or delete the virtual machine credentials that you do not need anymore by clicking 🗑 **Delete**.

> ⓘ Important You can unassign or delete credentials from a virtual machine only if the discovered applications running on it do not have assigned policies or available restore points. Therefore, before unassigning or deleting credentials, make sure to unassign policies or expire restore point tiers.

Planning application protection

Before performing an application backup, you must plan your application protection strategy to address all the needs of your data protection environment. Planning the application protection consists of the following tasks:

Task	Instructions
1. Get familiar with general requirements to determine if your environment is ready for application data protection.	Make sure that you are familiar with all the prerequisites, limitations, considerations, and/or recommendations related to the virtual machines on which the applications that you plan to protect are running. For details, see "Planning virtual machine protection" on page 155.
2. Prepare for your application data protection.	Depending on the type of application that you want to protect, see one of the following sections: • "Preparing for SQL Server application protection" below • "Preparing for Active Directory application protection" on page 262 • "Preparing for Exchange Server application protection" on page 262 • "Preparing for Oracle application protection" on page 265 • "Preparing for SAP HANA application protection" on page 267 • "Preparing for PostgreSQL application protection" on page 268

Preparing for SQL Server application protection

Preparing for SQL Server application data protection includes the following tasks:

Task	Instructions
1. Get familiar with SQL Server application specifics.	"Getting familiar with SQL Server application specifics" on the next page
2. Configure SQL Server application	"Configuring SQL Server application

Task	Instructions
backup options.	backup options" on page 259
3. Exclude databases from the SQL Server application backup.	"Excluding databases from the SQL Server application backup" on page 261

Getting familiar with SQL Server application specifics

When setting up your environment for data protection, you must get familiar with all prerequisites, limitations, considerations, and/or recommendations that are specific to protecting SQL Server applications.

Prerequisites

- *For Nutanix clusters:* Databases must reside on the local disks on the Nutanix cluster.
- *For restoring an SQL Server database to a point in time:* The database must be set to the full or bulk-logged recovery model during the backup.
- *For SQL Server failover cluster:*
  - All virtual machines where an SQL Server failover cluster resides must be discovered by HYCU.
  - Policies must be assigned to all virtual machines on which the application instance is running.

Limitations

- The `tempdb` system database is excluded from all backups.
- The `master`, `model`, and `msdb` system databases cannot be restored individually. You can restore them only as part of the whole instance restore.
- Backing up a database that is set to single-user mode is not possible if it is already in use.
- *For Always On Basic Availability Groups:* No backups on a secondary replica are possible.

Considerations

- *Only if you have upgraded your SQL Server to a newer version.* HYCU recognizes the upgraded application as a new application and at the same time changes

the status of the old one to Protected deleted. Therefore, to ensure data protection for the upgraded application, do the following:

1. Assign credentials to the upgraded application to enable HYCU to access it. For details, see "Enabling access to application data" on page 252.
2. Assign a policy to the upgraded application to protect it. For details, see "Backing up applications" on page 270.

- Backing up transaction logs of an SQL Server database with the `AUTO CLOSE` option set to `TRUE` may fail if the database has the `RECOVERING` status.

Recommendation

It is recommended to use a dedicated disk of a sufficient size for storing temporary files generated during a backup. Otherwise, this data will be stored on the volume with the most free space (for SQL Server) or an operating system disk volume (for NDB-managed SQL Server), which may affect the restore performance.

Configuring SQL Server application backup options

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

Limitation

The Back up and truncate SQL transaction logs option is not available for the databases that are registered in NDB. Consequently, you cannot perform the point-in-time restore of such databases.

Accessing the Options tab

To access the Options tab, follow these steps:

1. In the navigation pane, click **Applications**.
2. From the list of discovered applications, select the one for which you want to specify the application backup options, and then click **Configuration**.
3. Click the **Options** tab.

Backup option	Instructions
Back up and truncate SQL transaction logs	Use the switch if you want your SQL Server transaction logs to be backed up and truncated in the

Backup option	Instructions
(enabled by default)	SQL Server database automatically as part of the HYCU application backup. In this case, you can use HYCU to recover the SQL Server database. If disabled, HYCU does not back up and truncate the SQL Server transaction logs. In this case, to recover the SQL Server database, you should apply the transaction logs manually after restoring data.
Enter path for temporary transaction log backup and metadata files *(optional)*	If specified, the backup copies of the SQL Server temporary files (transaction logs and metadata files) are stored to this location. Otherwise, these backup copies are stored to the `.hycu` folder on the root of the disk with the largest amount of free space. **Note** For better restore performance, it is recommended to use a dedicated disk for storing backup copies of temporary files.
Optimized SQL Server HADR protection	*This option is available for Windows virtual machines hosting SQL Server databases that are part of an Always On Availability Group, and is not available for such virtual machines residing in an AWS GovCloud (US), Azure, or Azure Government environment.* Enable this option if you want to run backups only on the secondary replica with the highest backup priority. If only the primary replica is available, the backups are run on the primary replica. **Important** If you plan to enable the Optimized SQL Server HADR protection option, take into account the following: • When this option is enabled, HYCU backs up only the operating system disk and the disks where the secondary replicas with the highest backup priorities are located. Make sure that each database replica is located on a separate disk for this optimization to save backup storage space.

Backup option	Instructions
	• *For servers hosting SQL Server databases that are part of an Always On Availability Group:* Storage exclusion is performed at the volume level, and not at the disk level.

Excluding databases from the SQL Server application backup

By default, all databases belonging to SQL Server applications are backed up during the application backup. However, if you want specific databases to be excluded from the backup, HYCU enables you to select these databases before the application backup is performed. You can exclude databases on standalone SQL Server instances, SQL Server Always On Availability Group instances, and SQL Server failover cluster instances.

Considerations

- The disks where the excluded databases are located will be backed up.
- *Only if you exclude a database that is part of an Always On Availability Group from the backup.* The database is excluded from all SQL Server instances participating in the group.
- *Only if you exclude a database on a failover cluster from the backup.* The database is excluded from the active node and all passive nodes.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, select the SQL Server application whose databases you want to exclude from the backup, and then click **Configuration**.
2. Click the **Exclude Databases** tab.
3. From the list of all databases, select the ones that you want to exclude from the backup.

   > **Important** The `master`, `model`, and `msdb`SQL Server system databases cannot be excluded from the backup and are grayed out.

4. Click **Save**.

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

Preparing for Active Directory application protection

Before you start protecting Active Directory applications, you must get familiar with all prerequisites, limitations, considerations, and/or recommendations that are specific to protecting Active Directory applications.

Prerequisites

- No volume groups may be attached to the client virtual machine.
- *For Nutanix clusters:* NGT must be installed and enabled on the client virtual machine. For instructions, see Nutanix documentation.
- *For vSphere environments:* VMware Tools of the latest version must be installed on virtual machines on which the applications you want to protect are running.
- *For Windows servers:* The VSS service must be enabled and running, and the VSS writer status must be stable.

Limitations

- Protecting multiple instances of the same application type running on a virtual machine is not supported.
- Protecting applications running on volume groups or on virtual machines with attached volume groups is not supported.
- Protecting applications that are running on NDB-managed database server virtual machines is not supported.
- *For Nutanix clusters:* Protecting applications that are running on virtual machines with IDE disks is not possible.
- *For XenServer, AWS GovCloud (US), Azure, and Azure Government environments:* Protecting Active Directory applications is not supported.

Preparing for Exchange Server application protection

Preparing for Exchange Server application data protection includes the following tasks:

Task	Instructions
1. Get familiar with Exchange Server application specifics.	"Getting familiar with Exchange Server application specifics" below
2. Configure Exchange Server application backup options.	"Configuring Exchange Server application backup options" on the next page

Getting familiar with Exchange Server application specifics

When setting up your environment for data protection, you must get familiar with all prerequisites, limitations, considerations, and/or recommendations that are specific to protecting Exchange Server applications.

Prerequisites

- No volume groups may be attached to the client virtual machine.
- All databases must be mounted.
- The Active Directory application must be protected.

  Because Exchange Server stores all configuration information in Active Directory, make sure that you also back up your Active Directory application so that you can retrieve the information about the configuration if required. For example, if an entire database is deleted by accident and you want to restore it, you must first restore the Active Directory application, and then you can restore this database by performing the Exchange Server restore. However, if only the contents of the database are deleted, you must restore only the Exchange Server application.
- *For Nutanix clusters:* NGT must be installed and enabled on the client virtual machine. For details on how to do this, see Nutanix documentation.
- *For vSphere environments:* VMware Tools of the latest version must be installed on virtual machines on which the applications you want to protect are running.
- *For Windows servers:* The VSS service must be enabled and running, and the VSS writer status must be stable.
- You must run a custom PowerShell session configuration. Run the following PowerShell commands on the virtual machine:

```
New-PSSessionConfigurationFile -Path .\HycuJEA.pssc -
RunAsVirtualAccount
```

```
Register-PSSessionConfiguration -Path .\HycuJEA.pssc -Name
hycu.powershell -Force
```

Limitations

- Backing up multiple instances of the same application type running on a virtual machine is not supported.
- Backing up the applications running on the volume groups or on the virtual machines with the attached volume groups is not supported.
- *For Nutanix clusters:* Protecting applications that are running on virtual machines with IDE disks is not possible.
- *For XenServer, AWS GovCloud (US), Azure, and Azure Government environments:* Protecting Exchange Server applications is not supported.

Configuring Exchange Server application backup options

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

Accessing the Options tab

To access the Options tab, follow these steps:

1. In the navigation pane, click **Applications**.
2. From the list of discovered applications, select the one for which you want to specify the application backup options, and then click **Configuration**.
3. Click the **Options** tab.

Backup option	Instructions
Priority for Exchange Server restore requests	Specify the priority in which the restore requests for a mailbox restore are processed on the Exchange Server: Lowest, Lower, Low, Normal (the default value), High, Higher, Highest, Emergency.
Optimized Exchange Server DAG protection	*Available for Windows servers running Exchange mailbox servers that are members of a database availability group*

Backup option	Instructions
	(DAG). Enable this option if you want to back up only the volumes hosting the passive database copies with the highest activation preference number (including the system volume). Only the DAG members with the assigned policies are taken into account when searching for the database copies with the highest activation preference number. If no passive database copies are available, active database copies will be backed up. **Important** Optimized Exchange Server DAG protection is effective only if separate databases are stored on separate volumes.

Preparing for Oracle application protection

Preparing for Oracle application data protection includes the following tasks:

Task	Instructions
1. Get familiar with Oracle application specifics.	"Getting familiar with Oracle application specifics" below
2. Configure Oracle application backup options.	"Configuring Oracle application backup options" on the next page

Getting familiar with Oracle application specifics

When setting up your environment for data protection, you must get familiar with all prerequisites, limitations, considerations, and/or recommendations that are specific to protecting Oracle applications.

Prerequisites

- The SSH service must be enabled on the Oracle server and must be listening on port 22 for incoming connections.
- The Oracle database user must have the SYSDBA privilege.
- The database must be running in ARCHIVELOG mode.
- Tablespaces must be online.

- Additional disk space must be provided for temporary files created between two database backups. For optimal restore performance, separate disks must be specified for the temporary and database files.

Limitation

Backing up Oracle Real Application Clusters (RAC) databases is not supported. Consequently, assigning policies to such databases is not possible.

Recommendation

It is recommended to use a dedicated disk of a sufficient size for storing temporary files generated during a backup. Otherwise, this data will be stored on an operating system disk volume, which may affect the restore performance.

Configuring Oracle application backup options

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

Limitation

The Back up and truncate Oracle archive logs option is not available for the databases that are registered in NDB. Consequently, you cannot perform the point-in-time restore of such databases.

Accessing the Options tab

To access the Options tab, follow these steps:

1. In the navigation pane, click **Applications**.
2. From the list of discovered applications, select the one for which you want to specify the application backup options, and then click **Configuration**.
3. Click the **Options** tab.

Backup option	Instructions
Back up and truncate Oracle archive logs *(enabled by default)*	Use the switch if you want your Oracle archive logs to be backed up and truncated in the Oracle database automatically as part of the HYCU application backup. In this case, you can use HYCU to recover the Oracle

Backup option	Instructions
	database. If disabled, HYCU does not back up and truncate the Oracle archive logs. In this case, to recover the Oracle database, you should apply the transaction logs manually after restoring data.
Enter path for temporary Oracle files *(optional)*	If specified, the temporary Oracle files will be stored to this location. **Note** For better restore performance, it is recommended to use a dedicated disk for storing backup copies of temporary files.

Preparing for SAP HANA application protection

Before you start protecting SAP HANA applications, you must get familiar with all prerequisites, limitations, considerations, and/or recommendations that are specific to protecting SAP HANA applications.

Prerequisites

- SAP HANA savepoints must be enabled.
- *For multiple volume groups:* All data volumes and log volumes must belong to the same volume group.
- *For distributed (multi-host) environments:*
  - All virtual machines where SAP HANA resides must be discovered by HYCU.
  - Policies must be assigned to all virtual machines on which the application instance is running.

Limitation

Backing up multiple instances of the same application type running on a virtual machine is not supported.

Preparing for PostgreSQL application protection

Preparing for PostgreSQL application data protection includes the following tasks:

Task	Instructions
1. Get familiar with PostgreSQL application specifics.	"Getting familiar with PostgreSQL application specifics" below
2. Configure PostgreSQL application backup options.	"Configuring PostgreSQL application backup options" below

Getting familiar with PostgreSQL application specifics

When setting up your environment for data protection, you must get familiar with all prerequisites, limitations, considerations, and/or recommendations that are specific to protecting PostgreSQL database clusters.

Prerequisite

Only if you plan to protect a standby database cluster. The database cluster must be configured to run in hot standby mode.

Configuring PostgreSQL application backup options

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

Prerequisite

Only if you plan to protect a standby database cluster. The application-specific credentials must be assigned to the database cluster. For details on assigning credentials to entities, see "Enabling access to application data" on page 252.

Limitation

The Continuous WAL archiving for point-in-time restore option is not available for the databases that are registered in NDB. Consequently, you cannot perform the point-in-time restore of such databases.

Considerations

Only if you plan to enable the Continuous WAL archiving for point-in-time restore option. Consider the following:

- Depending on whether the archive mode is already enabled on the database cluster whose data you want to protect, the following happens:
  - If archive mode is enabled: HYCU sets up an additional archive location that will be used exclusively for point-in-time restore purposes.
  - If archive mode is not enabled: HYCU automatically enables the archive mode on the database cluster and creates an archive location that will be used exclusively for point-in-time restore purposes. In this case, HYCU prompts you to confirm the database cluster restart.
- There must be enough disk space in the HYCU archive location to store WAL segment files that are generated between two backups.
- If you do any modifications related to archiving on the database cluster after enabling the Continuous WAL archiving for point-in-time restore option, you must disable this option and enable it again in HYCU for the changes to take effect.

Accessing the Options tab

To access the Options tab, follow these steps:

1. In the navigation pane, click **Applications**.
2. From the list of discovered applications, select the one for which you want to specify the application backup options, and then click **Configuration**.
3. Click the **Options** tab.

Backup option	Instructions
Default database	From the Default Database drop-down menu, select the database that you want HYCU to use as the default database for executing queries and functions on the database cluster. The default is postgres.
Immediate checkpoint	Enable the **Immediate checkpoint** switch if you want to request an immediate checkpoint during the backup. In this case, you do not wait for a regular checkpoint that is scheduled by the system, but enforce an immediate one that uses the maximum I/O

Backup option	Instructions
	throughput on the database cluster and finishes as fast as possible.
Continuous WAL archiving for point-in-time restore	1. Enable the **Continuous WAL archiving for point-in-time restore** switch if you want to configure continuous WAL archiving for the database cluster, which enables you to perform the point-in-time restore. 2. *Optional.* In the Path to temporary WAL segments field, enter the path to the temporary WAL segment files. If you leave this field empty, the temporary WAL segment files will be stored in `/opt/hycu`. 3. *For a standby database cluster:* If the primary database cluster host and port are not automatically retrieved and displayed, provide the required information. **Note** HYCU automatically retrieves the primary database cluster host and port only if the database user has the `pg_read_all_settings` role granted on the database cluster.

Backing up applications

An application-aware backup allows a consistent backup of discovered applications.

Accessing the Applications panel

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

Limitations

- A tape target cannot be used for storing application data.
- Backing up multiple application types running on a virtual machine is not supported.

- Backing up applications running on virtual machines in ROBO environments is not supported.

Considerations

- If during virtual machine synchronization, a virtual machine cannot be found in a source environment, the status of this virtual machine and any discovered applications running on it is set to PENDING_REMOVAL. The policy is still assigned to the virtual machine and the applications, but you cannot perform any data protection actions (they are grayed out in HYCU). Depending on whether this virtual machine is found in the source environment during the time interval of two automatic virtual machine synchronization processes, the following happens:
  - *The virtual machine is found in the source environment:* Its status and the status of the applications running on it is changed to Protected.
  - *The virtual machine is not found in the source environment:* If the virtual machine still has at least one valid restore point available, its status and the status of the applications running on it is changed to Protected deleted. This means that the virtual machine that is deleted from the source is still considered protected and is not removed from HYCU.
- To ensure that HYCU can use Kerberos authentication, the virtual machine must use the credentials of an Active Directory domain account. If not, the legacy NTLM authentication will be used.

Procedure

1. In the Applications panel, select applications that you want to back up.

   > **Tip** To narrow down the list of all displayed applications, you can use the filtering options described in "Filtering and sorting data" on page 406.

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

   > **Note** When you assign the policy to the selected applications, the same policy is also assigned to the virtual machines on which they are running. If these virtual machines already have an assigned policy, the policy assigned to the applications takes precedence over the policy assigned to the virtual machines and is automatically assigned to the virtual machines.

The backup is scheduled according to the values that you defined for your policy. If required, you can also perform a manual backup of any application at any time. For details, see “Performing a manual backup” on page 424.

Restoring whole applications

With HYCU, you can restore a whole application to its original or a new location by restoring the virtual machine and attached volume groups on which the application is running.

Prerequisites

- *For Nutanix ESXi clusters and vSphere environments:* You must have the required restore privileges assigned. For details, see “Assigning privileges to a vSphere user” on page 552.
- *For applications with the Protected deleted status whose backups are stored on the imported targets:* Such applications must be discovered. For details, see “Enabling access to application data” on page 252.
- *For servers:* At least one hypervisor or cloud source must be added to HYCU to provide a storage container for storing the restore data. For details on how to add sources to HYCU, see “Adding sources” on page 73.

Considerations

- A restore is performed from the snapshot only if you are restoring to the same source (the source where the original virtual machine was running). If you are restoring to a different source, depending on the tier that you select for the restore, the following will happen:
  - If you select Snapshot, the restore will fail.
  - If you select Automatic, the restore will be performed from the target if there is an available target. Otherwise, it will fail.
- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- You cannot perform a restore of an application whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up`

configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.

- Restoring applications is not supported for NDB-managed applications if multiple application instances run on the same virtual machine.
- To ensure that HYCU can use Kerberos authentication, the virtual machine must use the credentials of an Active Directory domain account. If not, the legacy NTLM authentication will be used.
- *For Oracle:* If you disabled the Back up and truncate Oracle archive logs option, the database must be taken out of backup mode.
- *For AWS GovCloud (US) virtual machines with encrypted volumes:* Depending on whether you are restoring such a virtual machine to the same or a different region, the following applies:
  - To the same region: The restored volumes will be encrypted with the same KMS key as the original ones.
  - To a different region: The restored volumes will be encrypted with the default KMS managed key for EBS encryption.
- *Only if you are cloning a server running an application for which the Optimized Exchange Server DAG protection option was enabled during the backup.* After cloning the server, the volumes that were excluded during the backup will show up as RAW partitions in the Disk Management console. To recover data from the RAW partitions, do the following:
  1. Reformat the RAW partitions to NTFS or REFS.
  2. Reseed the database copies.

Restore options

You can select between the following restore options:

Restore option	Description
Restore VM	Enables you to restore an application by restoring the virtual machine on which the application is running to the same source. Select this option if you want to replace the original virtual machine with the restored one. For instructions, see "Restoring a virtual machine" on the next page.

Restore option	Description
	ⓘ **Important** The Restore VM option cannot be used for restoring the following: • Applications running on Azure Local and Hyper-V virtual machines • Applications running on servers
Restore Database Server VM	*Applicable only for NDB managed database servers.* Enables you to restore a database server virtual machine by restoring the virtual disks or volume groups to the same source. For instructions, see "Restoring a database server virtual machine" on page 292.
Clone VM	Enables you to restore an application by creating a clone of the virtual machine on which the application is running to the same or a different source. Select this option if you want to keep the original virtual machine. For instructions, see "Cloning a virtual machine" on page 293.

Restoring a virtual machine

You can restore an application by restoring the virtual machine on which the application is running to its original or a new location on the same source. In this case, the original virtual machine will be overwritten.

> ⚠ **Caution** When you are restoring the application to the original location, the restored data overrides the data in the original location. To avoid data loss, make sure that you back up the potentially unprotected data—the data that appeared between the last successful backup and the restore. To start a manual backup, see "Performing a manual backup" on page 424.

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

- "Restoring a virtual machine to a Nutanix cluster or a vSphere environment" on the next page
- "Restoring a virtual machine to a XenServer environment" on page 281

- “Restoring a virtual machine to an AWS GovCloud (US) environment” on page 284
- “Restoring a virtual machine to an Azure or Azure Government environment” on page 288

For details on how to restore an NDB-managed database server virtual machine, see “Restoring a database server virtual machine” on page 292.

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

Considerations

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

    - If restoring to a datacenter that is added to HYCU, the protection status of the virtual machine changes back to Protected.
    - If restoring to any datacenter that is not added to HYCU, the protection status of the virtual machine stays Protected deleted.
- *Only if you plan to restore a virtual machine running on a Nutanix ESXi cluster.* If Snapshot is selected as the backup target type in your policy, the `NVRAM` file will not be restored.
- *For Active Directory applications:* HYCU by default performs a non-authoritative restore of the Active Directory application, which means that the application data that has changed since the last backup is not replicated to other domain controllers. However, you can also perform an authoritative restore of the Active Directory application (a full restore or a granular restore of individual objects) and replicate all the restored application data to all the remaining domain controllers.

  Depending on what kind of restore you plan to perform, consider the following:
    - For the non-authoritative restore: Use the procedure described in this section.
    - For the authoritative-restore:
        1. Use the procedure described in this section, having in mind that you must prevent the virtual machine from turning on after the restore (you do this by disabling the **Power virtual machine on** switch).
        2. After you restore the Active Directory application, you must perform additional steps to make sure all the restored application data is replicated to all the remaining domain controllers. For instructions, see “After restoring an Active Directory application” on page 280.

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

   > **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click **Restore**.
4. Select **Restore Whole Server**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. In the General section, do the following:
   a. From the Storage container drop-down menu, select where you want to restore the virtual machine. By default, the original storage container is selected.

      > **Note** If you decide to restore the virtual machine to another storage container, keep in mind the following:
      > - Restore from the Snapshot tier cannot be performed to another storage container.
      > - If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.

   b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
      - **Automatic**: Ensures the fastest restore to the latest state.
      - **Backup**
      - **Copy**
      - **Archive**
      - **Snapshot**

   c. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

      If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:
      - In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.

- In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

  > **Note** The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.

d. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore. The original virtual machine will be deleted automatically.

   > **Important** *Only if you are restoring a vSphere virtual machine to a vSphere environment and you have disabled the Power virtual machine on switch.* When you try to power on the virtual machine and you are prompted to answer whether the virtual machine has been moved or copied, make sure to answer **I Moved It**.

e. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.

f. *For volume groups attached to the virtual machine:* Use the **Restore volume groups** switch if you also want to restore the volume groups that are attached to the virtual machine.

7. In the Network section, review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

   Depending on whether the original networks are available, proceed as follows:

   - If the original networks are available, you can leave the default values and restore the virtual machine with the original network settings, or you can modify the network settings.
   - If the original networks are not available, you must modify the network settings.

Modifying network settings

Original networks are...	Instructions
Available	You can do the following: • Add a new network adapter by clicking **New**, and then selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking **Delete**.
Unavailable	You can do the following: • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking **Delete**. • Add a new network adapter by clicking **New**, and then selecting the preferred network.

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

8. Click **Restore**.

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

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

After restoring an Active Directory application

If you want your Active Directory application to be restored in the authoritative mode and replicated to all domain controllers, after you restore the application, you must complete the additional steps described in this section.

> ⓘ **Important** Follow this procedure only if you are performing an authoritative restore of the Active Directory application.

Procedure

1. Sign in to the Nutanix Prism web console or the vSphere Client.
2. Disconnect the virtual machine from the network. For instructions, see Nutanix or VMware documentation.
3. Turn on the virtual machine, and then launch the web console.
4. Connect to the domain controller by using the local administrator user ID and password.
5. From the PowerShell console, run the following command:

   ```
   bcdedit /set safeboot dsrepair
   ```

6. Restart the domain controller. You can do this from the PowerShell console by running the following command:

   ```
   shutdown -t 0 -r
   ```

7. Connect the virtual machine to the network, and then turn on the virtual machine.

   After the virtual machine turns on, the domain controller starts up in the safe boot mode.
8. Log on to Directory Services Restore Mode (DSRM) by using the DSRM user name and password (a local administrator account).
9. Run `NTDSUtil` from the PowerShell console, and then do the following:
   a. Activate the Active Directory database instance by running the `activate instance ntds` command.
   b. Initiate the authoritative restore by running the `authoritative restore` command.
   c. Select a subtree or an object of the Active Directory application for which you want to perform the authoritative restore:

- For the subtree, run the following command:

  ```
  restore subtree "<DistinguishedName>"
  ```

  For example:

  ```
  restore subtree "DC=example,DC=com"
  ```

- For the object, run the following command:

  ```
  restore object "<DistinguishedName>"
  ```

  For example:

  ```
  restore object "CN=John
  Doe,OU=Sales,OU=Employees,DC=example,DC=com"
  ```

d. In the Authoritative Restore Confirmation dialog box, click **Yes**.

e. After the authoritative restore completes successfully, type `quit`.

10. Open the PowerShell console, and then run the following command:

    ```
    bcdedit /deletevalue safeboot
    ```

11. Restart the domain controller. You can do this from the PowerShell console by running the following command:

    ```
    shutdown -t 0 -r
    ```

Restoring a virtual machine to a XenServer environment

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

> **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click **Restore**.
4. Select **Restore Whole Server**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. In the General section, do the following:
   a. From the Storage Container drop-down menu, select where you want to restore the virtual machine. By default, the original storage container is selected.

      > **Note** If you decide to restore the virtual machine to another storage container, keep in mind the following:
      > - The restore from the Snapshot tier cannot be performed to another storage container.
      > - If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.

   b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
      - **Automatic**: Ensures the fastest restore to the latest state.
      - **Backup**
      - **Copy**
      - **Archive**
      - **Snapshot**

   c. Enable the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

      If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:
      - In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
      - In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

> **Note** The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.

d. Enable the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore. The original virtual machine will be deleted automatically.

e. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.

7. In the Network section, review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

   Depending on whether the original networks are available, proceed as follows:

   - If the original networks are available, you can leave the default values and restore the virtual machine with the original network settings, or you can modify the network settings.

     > **Note** The network name may differ from the one that is configured in XenCenter.

   - If the original networks are not available, you must modify the network settings.

Modifying network settings

Original networks are...	Instructions
Available	You can do the following: • Add a new network adapter by clicking **New**, and then selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting it, and then clicking **Edit** and selecting the

Original networks are...	Instructions
	preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking **Delete**.
Unavailable	You can do the following: • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking **Delete**. • Add a new network adapter by clicking **New**, and then selecting the preferred network.

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

8. Click **Restore**.

Restoring a virtual machine to an AWS GovCloud (US) environment

Considerations

- Make sure that the virtual machine you are restoring is not deleted from AWS GovCloud (US). If you delete a virtual machine from AWS GovCloud (US), you cannot restore it even if it still has a valid restore point available in HYCU (that is, even if its status is Protected deleted).
- When restoring a virtual machine, the original virtual machine disks are deleted and replaced with the restored ones.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

   > **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click **Restore**.
4. Select **Restore Whole Server**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. The following information is displayed and preselected:
   - The AWS GovCloud (US) account to which the virtual machine will be restored.
   - The account ID of the AWS GovCloud (US) account to which the virtual machine will be restored.
   - The region to which the virtual machine will be restored.
   - *Only if found by HYCU.* The key pair name for connection to the restored virtual machine.

     > **Important** If HYCU does not find the name of your key pair and you want to use it, you can select it from the Key pair name drop-down menu.

   - The availability zone to which the virtual machine will be restored.
7. Click **Next**.
8. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**

- **Archive**
- **Snapshot**

9. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

   If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

   - In the vCPU threads field, enter the number of CPUs for the restored virtual machine multiplied by the number of cores per CPU and the number of threads per core.
   - In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
   - From the Virtual machine type drop-down menu, select the virtual machine type.

     > **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

10. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
11. Under Network interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the Virtual Private Cloud (VPC) to which the original virtual machine belongs. If required, you can also modify network settings.

    ### Modifying network settings

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

    > **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same Virtual Private Cloud (VPC). The maximum number of network interfaces that you can add depends on the selected virtual machine type.

    Depending on how you want to modify network settings, do one of the following:

- Click **Add network interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow these steps:
  a. The Virtual Private Cloud (VPC) to which the network interface will be added is displayed and preselected.
  b. From the Subnets drop-down menu, select the subnet to which the network interface should be assigned.
  c. From the Security groups drop-down menu, select one or more security groups that will be associated with the network interface. If you want to select all the available security groups, select **Select all**.
  d. In the Public address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Auto-assign	An automatically allocated public IP address will be assigned to the network interface on the restored virtual machine.
Elastic IP (Reserved)	An elastic public IP address that you reserved in AWS GovCloud (US) will be assigned to the network interface on the restored virtual machine.
Elastic IP (New)	An elastic public IP address will be assigned to the network interface on the restored virtual machine.

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

Option	Description
Auto-assign	An automatically allocated private IP address will be assigned to the network interface on the restored virtual machine.
Custom	A private IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

12. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
13. Under Operating system license, select one of the following options:

OS license option	Select this option if you want to...
Keep existing license	Keep the existing OS license on the restored virtual machine. **Important** Make sure that the existing license is applicable also in AWS.
Available only for the Windows Server OS. **Replace existing license with AWS license**	Replace the existing OS license with an AWS license on the restored virtual machine.

14. Click **Restore**.

Restoring a virtual machine to an Azure or Azure Government environment

Consideration

If you want the restored virtual machine to have the same static IP address as the original virtual machine, do one of the following:

- Before the restore, in Azure or Azure Government, disassociate the IP address from the original virtual machine, and then select this IP address for the network interface during the restore in HYCU.
- During the restore, select a different IP address for the network interface. After the restore, in Azure or Azure Government, assign the preferred IP address to the restored virtual machine.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

   > **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click **Restore**.
4. Select **Restore Whole Server**, and then click **Next**.
5. Select **Restore VM**, and then click **Next**.
6. From the Availability Zone drop-down menu, select the zone for the restored virtual machine. If you do not want to restore data to any zone, select None.
7. Click **Next**.
8. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**
9. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

   If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU cores field, enter the number of virtual CPUs for the restored virtual machine.
- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- From the Virtual machine type drop-down menu, select the virtual machine type.

  > **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

10. Under Network Interfaces, you can view all the network interfaces that are attached to the virtual machine. Keep in mind that the IP address that is assigned to the primary IP configuration of the network interface will be assigned to the network interface on the restored virtual machine. If required, you can also modify network settings.

    ### Modifying network settings

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

    > **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on the selected virtual machine type.

    Depending on how you want to modify network settings, do one of the following:

    - Click **Add Network Interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow these steps:
      a. *Only if you are adding a network interface.* From the Virtual Network drop-down menu, select the network for the network interface.

         > **Note** The list of available networks includes only the ones within the region you selected for the restored virtual machine.

      b. From the Subnet drop-down menu, select the subnet to which the network interface should be assigned.

c. Under NIC Network Security Group, select the network security group for the network interface. You can select among the following options:

Option	Description
None	The network interface will not be assigned to a network security group.
Basic	The network interface will be assigned to Azure's basic network security group.
Advanced	The network interface will be assigned to the network security group that you select from the drop-down menu. By default, the network security group of the original virtual machine is selected.

d. Under Public IP Address Type, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Static	A static IP address will be assigned to the network interface on the restored virtual machine.
Existing	A preferred public IP address resource that you have created in Azure or Azure Government will be assigned to the network interface on the restored virtual machine.

e. Under Private IP Address Type, select the private IP address for the network interface. You can select between the following options:

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

11. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
12. Click **Restore**.

Restoring a database server virtual machine

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

Considerations

- An NDB-managed database server virtual machine is represented by the icon in the list of virtual machines.
- *Only if volume groups are attached.* The original volume group and its virtual disks are deleted during restore.
- The original database server virtual machine is restarted automatically.

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

   > **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click **Restore**.
4. Select **Restore Whole Server**, and then click **Next**.
5. Select **Restore Database Server VM**, and then click **Next**.
6. In the General section, do the following:
   a. From the Storage container drop-down menu, select where you want to restore the virtual machine. By default, the original storage container is selected.

      > **Note** If you decide to restore the virtual machine to another storage container, the fast restore or restore from Snapshot tier cannot be performed because the restore will be performed from the target and not from the snapshot.

   b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
      - **Automatic**: Ensures the fastest restore to the latest state.
      - **Backup**
      - **Copy**
      - **Archive**
      - **Snapshot**
7. *Only applicable for Time Machines with a Continuous SLA.* Enable the **Run a Tail Log Backup on NDB before restore** switch to ensure that an NDB Tail Log Backup is started before a database restore. A Tail Log Backup is started automatically after a restore.
8. Click **Restore**.

Cloning a virtual machine

You can restore an application by creating a clone of the virtual machine on which the application is running to its original or a new location on the same or a different source. In this case, the original virtual machine will not be overwritten.

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

- “Cloning a virtual machine to a Nutanix cluster, a vSphere environment, an Azure Local environment, or a Hyper-V cluster” on the next page
- “Cloning a virtual machine to a XenServer environment” on page 300

- “Cloning a virtual machine to an AWS GovCloud (US) environment” on page 304
- “Cloning a virtual machine to an Azure or Azure Government environment” on page 308

> **Important** After you clone the virtual machine, make sure that everything works as expected by going through the considerations and the recommendations listed in “After cloning a virtual machine” on page 224.

Cloning a virtual machine to a Nutanix cluster, a vSphere environment, an Azure Local environment, or a Hyper-V cluster

Prerequisites

- *For virtual machines that you plan to clone to a new location:* The source to which you plan to clone the virtual machine must be added to HYCU. For details on how to do this, see “Adding sources” on page 73.
- *For Linux servers:* In the `/etc/fstab` system configuration file of the server, UUIDs (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`) must be used instead of device names for file system device identification.
- *For virtual machines that you plan to clone to an Azure Local environment:* The virtual hard disk path and the virtual machine path must point to the cluster shared volume location. For example:
  - Virtual Machines Path: `C:\ClusterStorage\UserStorage\VirtualMachines`
  - Virtual Hard Disks Path: `C:\ClusterStorage\UserStorage\VirtualHardDisks`

Limitation

For vSphere environments: Attaching the ISO image to the restored virtual machine is not supported.

Considerations

- *Only if volume groups are attached to the virtual machine that you are cloning.* You can choose to restore the volume groups together with the virtual machine if they were attached to it at backup time. In this case, the original volume groups are kept alongside of the restored ones. If the volume groups are also attached to other virtual machines, the following applies (depending

on how they are attached to the virtual machines):

  - Directly: Volume groups are automatically attached only to the cloned virtual machine.
  - By using iSCSI: Volume groups are automatically attached to all virtual machines to which they were attached at backup time.

- *For restoring a virtual machine running on a Nutanix AHV cluster to a Nutanix ESXi cluster:* If virtual machine disks are attached to the PCI bus, the bus type will be automatically changed to SCSI after the restore. Because of this configuration change, the restore finishes with a warning.
- *For Linux virtual machines running on a Nutanix ESXi cluster:* If after restoring a virtual machine that was created through the vSphere Client, the virtual machine does not boot, follow the steps described in "After restoring a virtual machine to a Nutanix ESXi cluster" on page 638.
- After you restore a virtual machine, it might happen that the order of virtual disks differs from the one on the original virtual machine if you performed the restore:
  - From a Nutanix AHV cluster to a Nutanix ESXi cluster or a vSphere environment
  - From a Nutanix ESXi to another Nutanix ESXi cluster
  - From a vSphere environment to a Nutanix ESXi cluster

  In this case, make the necessary adjustments, including the selection of the correct boot disk.
- *Only if you plan to restore vSphere virtual machine data to the original storage container.* If the storage container is mounted to several hosts and the original host is powered off or in maintenance mode at restore time, data will be restored to the same storage container on a different host.
- *Only if ownership is set for the virtual machine.* The same owner is automatically assigned to the restored virtual machine.
- *Only if you plan to restore a virtual machine running on a Nutanix ESXi cluster.* If Snapshot is selected as the backup target type in your policy, the `NVRAM` file will not be restored.
- *Only if the original virtual machine resides on a source other than a vSphere environment.* Make sure to modify the virtual machine configuration by specifying the appropriate guest operating system.

Recommendation

For Linux virtual machines: It is recommended that the use of persistent network device names based on MAC addresses is disabled. Otherwise, you will have to configure the network manually. For details on how to disable the use of persistent network device names, see your Linux distribution documentation.

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

   > **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click **Restore**.
4. Select **Restore Whole Server**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination Source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. Under General, do the following:
   a. From the Storage Container drop-down menu, select the storage container where you want to restore the virtual machine.

      > **Note** By default, the original storage container is selected. If you decide to restore the virtual machine to another storage container, keep in mind the following:
      >
      > - Restore from the Snapshot tier cannot be performed to another storage container.

- If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.
- If the selected storage container is on a different source, additional prerequisites apply. For details, see "Preparing for the restore to a different source" on page 163.

b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**

   Note The Snapshot tier is not available for Azure Local environments and Hyper-V clusters.

c. In the New VM Name field, specify a new name for the virtual machine.

   Important *For Azure Local environments and Hyper-V clusters:* The name that you specify may not contain spaces.

d. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

   If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

   - In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
   - In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

     Note The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

   - In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.

e. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore.

> **Important** Make sure to consider the following:
> - This option is disabled for virtual machines that have volume groups attached by using iSCSI. For details on what needs to be done before turning on the restored virtual machine, see "After cloning a virtual machine" on page 224.
> - *Only if you are cloning a virtual machine from a Nutanix cluster, a vSphere environment, an Azure Local environment, or a Hyper-V cluster.* If you turn the restored virtual machine on, the original virtual machine will be turned off automatically.
> - *Only if you are cloning a vSphere virtual machine to a vSphere environment and you have disabled the Power virtual machine on switch.* When you try to power on the virtual machine and you are prompted to answer whether the virtual machine has been moved or copied, make sure to answer **I Copied It**.

f. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.

g. *For volume groups attached to the virtual machine:* Use the **Clone volume groups** switch if you want to also restore the volume groups that are attached to the virtual machine.

8. Under Network, do the following:

   a. Review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

      Depending on whether the original networks are available, proceed as follows:

      - If the original networks are available, you can leave the default values and clone the virtual machine with the original network settings, or you can modify the network settings.
      - If the original networks are not available, you must modify the network settings.

Modifying network settings

Original networks are...	Instructions
Available	You can do the following: • Add a new network adapter by clicking **New** and selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting the virtual adapter, and then clicking **Edit** and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking **Delete**.
Unavailable	You can do the following: • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking **Delete**. • Add a new network adapter by clicking **New**, and then selecting the preferred network.

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

b. *Only if you are restoring the virtual machine to a different source.* Use the **Keep original MAC address** switch if you want the restored virtual machine to keep the original MAC address. Keep in mind that this is applicable only if at least one network adapter has a MAC address assigned.

9. Click **Restore**.

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

There are some considerations that you should be aware of after cloning a virtual machine. For details, see “After cloning a virtual machine” on page 224.

Cloning a virtual machine to a XenServer environment

Prerequisites

- *For virtual machines that you plan to clone to a new location:* The source to which you plan to clone the virtual machine must be added to HYCU. For details on how to do this, see “Adding sources” on page 73.
- *For Linux servers:* In the `/etc/fstab` system configuration file of the server, UUIDs (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`) must be used instead of device names for file system device identification.

Consideration

Only if ownership is set for the virtual machine. The same owner is automatically assigned to the restored virtual machine.

Recommendation

For Linux virtual machines: It is recommended that the use of persistent network device names based on MAC addresses is disabled. Otherwise, you will have to configure the network manually. For details on how to disable the use of persistent network device names, see your Linux distribution documentation.

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

   > **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for

restoring the application.

3. Click **Restore**.
4. Select **Restore Whole Server**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination Source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. In the General section, do the following:
   a. From the Storage Container drop-down menu, select the storage container where you want to restore the virtual machine.

      Note By default, the original storage container is selected. If you decide to restore the virtual machine to another storage container, keep in mind the following:

      - The restore from the Snapshot tier cannot be performed to another storage container.
      - If you select the Automatic tier, the fast restore cannot be performed because the restore will be performed from the target and not from the snapshot.
      - If the selected storage container is on a different source, additional prerequisites apply. For details, see "Preparing for the restore to a different source" on page 163.

   b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
      - **Automatic**: Ensures the fastest restore to the latest state.
      - **Backup**
      - **Copy**
      - **Archive**
      - **Snapshot**

      Note *Only if you select the Archive tier and the data is stored on a QStar tape target.* Consider the following:
      To view tape target information, click **Tape Info**. The following information is displayed:

      - Target name and UUID
      - Barcode label of the media

- Information on whether the media is online or offline
- Location of the media

You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

c. In the New VM Name field, specify a name for the restored virtual machine.

d. Enable the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
- In the Cores per vCPU field, enter the number of cores per virtual CPU for the restored virtual machine.

> **Note** The total number of cores of the restored virtual machine will be the number of virtual CPUs multiplied by the number of cores per virtual CPU.

- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.

e. Use the **Power virtual machine on** switch if you want to turn the restored virtual machine on after the restore.

> **Important** If you turn the restored virtual machine on, the original virtual machine will be turned off automatically.

f. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.

8. In the Network section, do the following:

a. Review the list of network adapters that were added to the virtual machine at backup time (including the networks to which the virtual machine was connected). If any of the original networks is no longer available, N/A is shown.

Depending on whether the original networks are available, proceed as follows:

- If the original networks are available, you can leave the default values and clone the virtual machine with the original network settings, or you can modify the network settings.

  > **Note** The network name may differ from the one that is configured in XenCenter.

- If the original networks are not available, you must modify the network settings.

Modifying network settings

Original networks are...	Instructions
Available	You can do the following: • Add a new network adapter by clicking **New** and selecting the preferred network. • Edit the existing network adapter to connect the virtual machine to a different network by selecting the virtual adapter, and then clicking **Edit** and selecting the preferred network. • Delete the network adapter you do not need anymore by selecting it, and then clicking **Delete**.
Unavailable	You can do the following: • Edit the affected network adapter to connect the virtual machine to a new network by selecting it, and then clicking **Edit** and selecting the preferred network. • Delete the affected network adapter by selecting it, and then clicking **Delete**. • Add a new network adapter by clicking **New**, and then selecting the preferred network.

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

b. *Only if you are restoring the virtual machine to a different source.* Use the **Keep original MAC address** switch if you want the restored virtual machine to keep the original MAC address. Keep in mind that this is applicable only if at least one network adapter has a MAC address assigned.

9. Click **Restore**.

Cloning a virtual machine to an AWS GovCloud (US) environment

Prerequisite

- *For virtual machines that you plan to restore to a new location:* The AWS GovCloud (US) region to which you plan to restore the virtual machine must be added to HYCU. For instructions, see "Adding an AWS GovCloud (US) region" on page 82.
- *Only if the virtual machine that you plan to restore resides on a source other than AWS GovCloud (US).* A premium tier Platform license is required. For details, see "Licensing" on page 488.

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for restoring data to a new location.
- *For virtual machines that have BitLocker volumes encrypted with TPM-based keys:* Restoring such volumes is not supported.

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

> ⓘ Important If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click **Restore**.
4. Select **Restore Whole Server**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination Source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. From the AWS GovCloud (US) Account drop-down menu, select the account to which the virtual machine will be restored.

   The following information is displayed and preselected:

   - The account ID of the AWS GovCloud (US) account to which the virtual machine will be restored.
   - The region to which the virtual machine will be restored.
8. *Optional.* From the Key Pair Name drop-down menu, select the key pair name that you want to use for connection to the restored virtual machine.

   > ⓘ Important *For Windows virtual machines:* The key pair name that you select can be used only if the EC2Config or EC2Launch service was configured on the original virtual machine or if you configure it later on the restored virtual machine.
9. From the Availability Zone drop-down menu, select the availability zone to which the virtual machine will be restored.
10. Click **Next**.
11. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
    - **Automatic**: Ensures the fastest restore to the latest state.
    - **Backup**
    - **Copy**
    - **Archive**
    - **Snapshot**
12. In the New VM Name field, specify a name for the restored virtual machine.
13. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

- In the vCPU threads field, enter the number of CPUs for the restored virtual machine multiplied by the number of cores per CPU and the number of threads per core.
- In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
- From the Virtual machine type drop-down menu, select the virtual machine type.

  > **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

14. *Only if virtual disks were excluded from the backup (manually or automatically).* Enable the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the restored virtual machine.
15. Under Network Interfaces, you can view the network interface that will be added to the restored virtual machine. By default, this is the first network interface from the Virtual Private Cloud (VPC) to which the original virtual machine belongs. If required, you can also modify network settings.

    Modifying network settings

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

    > **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same Virtual Private Cloud (VPC). The maximum number of network interfaces that you can add depends on the selected virtual machine type.

    Depending on how you want to modify network settings, do one of the following:

    - Click **Add Network Interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow these steps:
      a. The Virtual Private Cloud (VPC) to which the network interface will be added is displayed and preselected.

b. From the Subnets drop-down menu, select the subnet to which the network interface should be assigned.

c. From the Security Groups drop-down menu, select one or more security groups that will be associated with the network interface. If you want to select all the available security groups, select **Select all**.

d. Under Public Address Type, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Auto-assign	An automatically allocated public IP address will be assigned to the network interface on the restored virtual machine.
Elastic IP (Reserved)	An elastic public IP address that you reserved in AWS GovCloud (US) will be assigned to the network interface on the restored virtual machine.
Elastic IP (New)	An elastic public IP address will be assigned to the network interface on the restored virtual machine.

e. Under Private Address Type, select the private IP address for the network interface. You can select between the following options:

Option	Description
Auto-assign	An automatically allocated private IP address will be assigned to the network interface on the restored virtual machine.
Custom	A private IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

16. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
17. Under Operating System License, select one of the following options:

OS license option	Select this option if you want to...
Keep existing license	Keep the existing OS license on the restored virtual machine. **Important** Make sure that the existing license is applicable also in AWS.
Available only for the Windows Server OS. **Replace existing license with AWS license**	Replace the existing OS license with an AWS license on the restored virtual machine.

18. Click **Restore**.

Cloning a virtual machine to an Azure or Azure Government environment

Prerequisites

- *For virtual machines that you plan to restore to a new location:* The Azure or Azure Government subscription to which you plan to restore the virtual machine must be added to HYCU. For details on how to do this, see "Adding an Azure subscription" on page 83 or "Adding an Azure Government subscription" on page 84.
- *For virtual machines that have Azure Disk Encryption enabled:* The key vault must be available on the location to which you are restoring the virtual machine.
- *Only if the virtual machine that you plan to restore resides on a source other than Azure or Azure Government.* A premium tier Platform license is required. For details, see "Licensing" on page 488.

Limitation

If a restore point contains only a Snapshot tier, you cannot use it for restoring data to a new location.

Consideration

Only if you plan to keep the original network settings on the restored virtual machine. If the original static IP address is still associated with the original or another virtual machine, a new random private IP address will be assigned to the restored virtual machine.

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

Procedure

1. In the Applications panel, click the application that you want to restore to open the Detail view.

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

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

   > **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the application.

3. Click **Restore**.
4. Select **Restore Whole Server**, and then click **Next**.
5. Select **Clone VM**, and then click **Next**.
6. From the Destination Source drop-down menu, select where you want to restore the virtual machine, and then click **Next**.
7. From the Service Principal drop-down menu, select the service principal that has access to the required resources (the source from which and to which you are restoring the virtual machine).
8. From the Subscription drop-down menu, select the subscription for the restored virtual machine.
9. From the Resource Group drop-down menu, select the resource group for the restored virtual machine.
10. From the Location drop-down menu, select the geographic region for the restored virtual machine.

11. From the Availability Zone drop-down menu, select the zone for the restored virtual machine.

    > **Note** The selected geographic region and the size of the virtual machine determine to which zones you can restore data. If you do not want to restore data to any zone, select **None**.

12. Click **Next**.
13. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
    - **Automatic**: Ensures the fastest restore to the latest state.
    - **Backup**
    - **Copy**
    - **Archive**
    - **Snapshot**
14. In the New VM Name field, specify a name for the restored virtual machine.
15. Use the **Use original VM configuration** switch if you want the restored virtual machine to have the same configuration settings as the original virtual machine.

    If you want to change any of the configuration settings, disable the **Use original VM configuration** switch, and then do the following:

    - In the vCPU(s) field, enter the number of virtual CPUs for the restored virtual machine.
    - In the Memory field, set the amount of memory (in GiB or MiB) for the restored virtual machine.
    - From the Virtual machine type drop-down menu, select the virtual machine type.

      > **Note** The list of virtual machine types is based on the number of virtual CPUs and the amount of memory that you specified. If no virtual machine type matches the specified values, the list is empty, and you must adjust the specified values.

16. Under Network Interfaces, you can view all the network interfaces that are attached to the virtual machine. Keep in mind that the IP address that is assigned to the primary IP configuration of the network interface will be assigned to the network interface on the restored virtual machine. If required, you can also modify network settings.

Modifying network settings

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

> **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add Network Interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow these steps:
  a. *Only if you are adding a network interface.* From the Virtual Network drop-down menu, select the network for the network interface.

     > **Note** The list of available networks includes only the ones within the region you selected for the restored virtual machine.

  b. From the Subnet drop-down menu, select the subnet to which the network interface should be assigned.
  c. Under NIC Network Security Group, select the network security group for the network interface. You can select among the following options:

Option	Description
None	The network interface will not be assigned to a network security group.
Basic	The network interface will be assigned to Azure's basic network security group.
Advanced	The network interface will be assigned to the network security group that you select from the drop-down menu. By default, the network security group of the original virtual machine is selected.

  d. Under Public IP Address Type, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the restored virtual machine.
Static	A static IP address will be assigned to the network interface on the restored virtual machine.
Existing	A preferred public IP address resource that you have created in Azure or Azure Government will be assigned to the network interface on the restored virtual machine.

e. Under Private IP Address Type, select the private IP address for the network interface. You can select between the following options:

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the restored virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the restored virtual machine.

f. Click **Add** or **Save**.

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

17. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
18. Click **Restore**.

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

Restoring SQL Server databases

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

Prerequisites

- *For point-in-time restore:* The database recovery model must be set to full or bulk-logged.
- *For restoring the whole SQL Server failover cluster instance:* The SQL Server service must be stopped by using the Failover Cluster Manager. For details on how to do this, see SQL Server documentation.
- For improved restore performance, the startup type of the Microsoft iSCSI Initiator Service may not be set to `Disabled`.
- *Only if you are restoring data that is stored in the archive access tier on an Azure target.* You must recreate a snapshot and use this snapshot for restoring data, or manually rehydrate data. For details on how to recreate a snapshot, see "Recreating snapshots" on page 436. For details on how to manually rehydrate data, see Azure documentation.
- *For AWS GovCloud (US), Azure, and Azure Government environments:* The virtual machine that hosts the SQL Server instance to which you plan to restore the SQL Server databases must be in the same virtual network as the HYCU backup controller.

Limitations

- Restoring SQL Server databases to another SQL Server application instance is supported only if you are restoring to the same or later version of the application.
- Databases that are part of an Always On Availability Group can be restored only to a primary node (from a secondary or primary node). However, keep in mind that in the case of an Always On Basic Availability Group, the databases can be restored only from a primary node.
- *Only if you plan to use the Archive restore point.* Performing the point-in-time restore is not supported.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives

missing or being stored on a deactivated target), you cannot use this tier for restoring data.

- If you are restoring the databases to a different SQL Server instance, they will be renamed and copied to the default SQL Server location of the selected target.
- If a virtual machine is deleted from the source, but it still has at least one valid restore point available, it is considered protected. In this case, the status of the virtual machine or any discovered applications running on it is Protected deleted. When restoring application items of such an application, keep in mind that you cannot restore them to the original application instance.
- You cannot perform a restore of an application whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up` configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
- If any databases were excluded from the backup, you cannot select them for the restore.
- *For SQL Server failover clusters:*
  - The restore needs to be redirected to the active SQL Server failover cluster instance.
  - The Overwrite existing databases option can be enabled for a redirected restore only if the database location also exists on the target virtual machine.
- *For restoring an SQL Server database to a different SQL Server instance:* The Overwrite existing databases option should be enabled only when restoring to an SQL Server instance which is on a different server and has identical database paths.
- *For NDB-managed SQL Servers:*
  - If additional databases are located on the same disk as the database that you select for restore, all databases on this disk will be restored.
  - An SQL Server instance managed by NDB is represented by the icon in the list of applications.
  - You cannot restore a whole instance.

- You cannot perform a granular restore of an NDB-managed database if the disk layout changes after the selected restore point is created.

Accessing the Applications panel

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

Procedure

1. In the Applications panel, click the application whose databases you want to restore to open the Detail view. The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open the Detail view.

   > **Note** With the SQL Server Always On Availability Group, you can expand the application item to view the discovered Availability Groups.

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

   > **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the databases.

3. Click **Restore**. The Restore MS SQL Server dialog box opens.

   > **Note** If the Back up and truncate SQL transaction logs option was disabled during the backup, you are prompted that database recovery must be performed after the restore.

4. Select **Restore databases**, and then click **Next**.
5. From the Target instance drop-down menu, select where you want to restore the databases.
6. *For SQL Server Always On Availability Group:* From the Destination availability group drop-down menu, select one of the available Availability Groups to restore the databases to this group or leave the field empty to restore the databases to the SQL Server.
7. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**

- **Archive**
- **Snapshot**

8. *Only if you want to restore the whole instance.* Select the **Whole instance** check box.
9. *Only if you want to restore databases.* Select the **Database name** check box to restore all databases or select the databases that you want to restore.

   > Note *Only if you are restoring an NDB-managed database.* If you select a database that is part of a group of databases that can be restored only as a whole group, all databases that are part of this group are automatically selected for restore.

10. *Optional.* Specify a point in time to which you want to restore data. The databases will be restored to the state they were in at the specified time.

    > Note To perform a point-in-time restore, select a backup that was performed before the specified point in time so that the database instance can be brought to the appropriate state by applying the transaction log files from the next backup.

11. *Only if you are restoring an NDB-managed database, for Time Machines with a Continuous SLA.* Enable the **Run a Tail Log Backup on NDB before restore** switch to ensure that an NDB Tail Log Backup is started before a database restore. A Tail Log Backup is started automatically after a restore.
12. Click **Next**.
13. Use the **Leave databases in restoring state** switch if you want to leave the databases in the restoring state. By doing so, you can apply transaction logs to the databases after the restore and perform a manual point-in-time restore.
14. Use the **Overwrite existing databases** switch if you want to overwrite existing databases when performing a restore. In this case, the backups will be restored to their original location and all data will be overwritten. Keep in mind that if you are restoring the databases to another SQL Server instance, all the databases that have the same names (and not necessarily the contents) will be overwritten.

    Otherwise, to restore data to a different location on the same or another SQL Server instance, specify a database prefix that will be given to the databases, a new database file location, and a new database log location.

    > Important If you are restoring the whole instance, you can only overwrite existing databases. In this case, the Overwrite existing

> databases option is enabled by default and you cannot disable it.

15. Click **Restore**.
16. *Only if the Back up and truncate SQL transaction logs option was disabled during the backup.* Recover the SQL Server databases by applying the transaction logs manually.
17. *Only if using SQL Server 2014 Always On Availability Groups.* Join the restored databases to an Always On Availability Group by using SQL Server Management Studio. For details on how to do this, see Microsoft documentation.

    > **Note** After you join the restored databases to the Always On Availability Group, it is recommended to perform a new backup of your Always On Availability Group.

18. *Only if restoring the whole SQL Server failover cluster instance.* Start the SQL Server service and all other related services by using the Failover Cluster Manager. For details on how to do this, see SQL Server documentation.

Restoring Exchange Server databases, mailboxes, and public folders

With HYCU, you can restore Exchange Server databases, mailboxes, and public folders. When restoring Exchange Server databases, you can choose between restoring to the original mailbox server and, if the mailbox server is a member of a Database Availability Group (DAG), to another mailbox server inside the DAG. When restoring mailboxes and public folders, the recovery database can be restored to the original mailbox server or any other mailbox server that is part of your Exchange Server organization. From there, the actual restore is performed to any mailbox or public folder within the organization.

Prerequisites

- *For restoring mailboxes:*
  - The mailbox to which you are restoring data must exist on the server and be initialized.

    - *Only if the original mailbox to which you plan to restore data was deleted from the server.* You must create a new mailbox with the same or a different name, and make sure it is initialized (to do so, sign in to it with your Exchange client).
- *For restoring public folders:* The public folder must exist in the public folder mailbox. If it does not exist, recreate it manually with the same name it had at backup time.
- For improved restore performance, the startup type of the Microsoft iSCSI Initiator Service may not be set to `Disabled`.
- *Only if you are restoring data that is stored in the archive access tier on an Azure target.* You must recreate a snapshot and use this snapshot for restoring data, or manually rehydrate data. For details on how to recreate a snapshot, see "Recreating snapshots" on page 436. For details on how to manually rehydrate data, see Azure documentation.

Limitation

Restoring data to the `hycu` subfolder (the Restore to subfolder option) is currently not supported for public folders.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- You cannot perform a restore of an application whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up` configuration setting in the HYCU `config.properties` file to `true`. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.

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

Procedure

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

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

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

   > **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring application items.

3. Click **Restore**. The Restore Exchange dialog box opens.
4. Select which application items you want to restore:
   - **Restore databases**
     a. From the Recovery database server drop-down menu, select the server for restoring the data. When specifying a recovery database server, keep in mind that you can select it only if your mailbox server is a member of a DAG and you want to restore data to another mailbox server inside the DAG. Otherwise, you can restore only to the original mailbox server.

        > **Important** *For restoring a mailbox server that is a member of a DAG:* Make sure to select the recovery database server on which the databases are currently active.

     b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
        - **Automatic**: Ensures the fastest restore to the latest state.
        - **Backup**
        - **Copy**
        - **Archive**
        - **Snapshot**

     c. Select the **Database name** check box to restore all databases or select the databases that you want to restore.

d. Use the **Enable restore to recovery database** switch if you want to enable restoring data to a recovery database. If enabled, provide a recovery database path. The default one is `C:\ProgramData\Hycu`.

- **Restore mailboxes and/or public folders**

a. From the Recovery database server drop-down menu, select the mailbox server for restoring the data. You can select among the mailbox servers that are part of your Exchange Server organization.

b. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**

c. From the list of mailboxes and/or public folders that are available for a restore, select the ones that you want to restore.

> Tip If there are too many mailboxes and/or public folders to be displayed on one page, you can move between the pages by clicking › and ‹. You can also use ⌄ to set the number of mailboxes and/or public folders to be displayed per page.
> You can search for a mailbox and/or public folder by entering its name and then pressing **Enter** in the Search field.

d. Enable the **Use non-default domain controller** switch if you want to use a domain controller other than the default one, and then, in the Domain controller field, enter the FQDN or IP address of the preferred domain controller.

e. Click **Next**.

f. Select where you want to restore data:
   - **Original mailbox**
   - **Alternate mailbox**, and then enter an alternate mailbox name.

g. Select the mode for restoring data:
   - **Restore in place**

     Enables you to restore data to the original location.
   - **Restore to subfolder** *(not supported for public folders)*

Enables you to restore data to the `hycu` subfolder that is created automatically.

h. *For restoring data to the original location:* Use the **Conflict resolution** switch if you want to resolve any potential data conflict by keeping the most recent version of the items in conflict. Otherwise, HYCU will overwrite the existing items with the ones from the backup.

i. Enter a temporary recovery database path. The default one is `C:\ProgramData\Hycu`.

5. Click **Restore**.

Restoring Oracle database instances and tablespaces

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

Prerequisites

- On the original virtual machine, references in the `/etc/fstab` system configuration file entries must use universally unique identifiers (for example, `UUID=8ff089c0-8e71-4320-a8e9-dbab8f18a7e5`) rather than device names (for example, `/dev/sda1`) unless they refer to logical volumes (for example, `/dev/mapper/ol-root`).
- The `bashrc` and `.bash_profile` scripts may not write to standard output (`STDOUT`) or standard error (`STDERR`) for the user whose credentials are used for application discovery.
- *Only if you are restoring data that is stored in the archive access tier on an Azure target.* You must recreate a snapshot and use this snapshot for restoring data, or manually rehydrate data. For details on how to recreate a snapshot, see “Recreating snapshots” on page 436. For details on how to manually rehydrate data, see Azure documentation.

Limitations

- Tablespaces can be restored only from the latest restore point in the backup chain and cannot be restored to a point in time.
- *Only if you plan to use the Archive restore point.* Performing the point-in-time restore is not supported.

Considerations

- When performing a database instance restore, keep in mind the following:
  - You can perform a complete or point-in-time restore:
    - Complete restore

      If you select the latest restore point and do not specify a point in time, HYCU performs a complete restore of the whole database instance to the latest possible time.
    - Point-in-time restore after the last backup

      If you select the latest restore point and specify a point in time after the last backup, the control file and archive log files are not restored. Instead, only the database files are restored and the existing archive log files are applied.
    - Point-in-time restore between backups

      To perform a point-in-time restore between backups, you must select a restore point other than the latest one and specify a point in time between that backup and the next one, so that the database instance can be brought to the point in time by applying the archive log files from the next backup.

      When performing the point-in-time restore between backups, the control file, database files, and required archive log files are restored.
  - 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 such a restore because the database will not be protected in terms of a complete restore until a new backup is performed.
  - If the control file is lost or needs to be restored due to missing tablespaces, or if the existing archive log files are lost, a point-in-time restore between backups must be performed.
- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- You cannot perform a restore of an application whose retention period specified in the policy has been exceeded (such restore points are grayed out in the HYCU web user interface). However, if required, this can be overridden by setting the `restore.enabled.if.retention.is.up` configuration setting in the HYCU `config.properties` file to `true`. For

details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.

- If you initially performed an Oracle application backup and you then restore the Oracle database by using the Restore VM, Restore Database VM, or Restore vDisks option, the Oracle database may be left in backup mode after the restore. In this case, you need to manually resume the database.

Accessing the Applications panel

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

Procedure

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

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

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

   > **Important** If the backup status for the selected restore point shows that the backup is crash consistent, you cannot use this restore point for restoring the database instance.

3. Click **Restore**. The Restore Oracle Database dialog box opens.

   > **Note** If the Back up and truncate Oracle archive logs option was disabled during the backup, you are prompted that database recovery must be performed after the restore.

4. Select **Restore databases**, and then click **Next**.
5. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**

6. *Only if you want to restore the whole database instance.* Select the **Whole instance** check box.
7. *Only if you want to restore tablespaces.* Select the **Tablespace name** check box to restore all tablespaces or select the tablespaces that you want to restore.
8. *Only if restoring the whole database instance.* Optionally, specify a point in time to which you want to restore data. The database instance will be restored to the state it was in at the specified time.
9. Click **Restore**.
10. *Only if the Back up and truncate Oracle archive logs option was disabled during the backup.* Recover the Oracle databases by applying the archive logs manually.

Restoring PostgreSQL database clusters

With HYCU, you can restore PostgreSQL database cluster data to the original database cluster, or in the case of a standby database cluster, the standby database cluster data to its primary database cluster.

Prerequisite

The system identifier of the database cluster whose data you want to restore must be the same as the system identifier of the destination database cluster.

Limitations

- You can restore only the whole database cluster instance.
- *Only if you selected the latest restore point and plan to restore data from a standby database cluster to its primary database cluster.* The Complete recovery and Point in Time options are not supported.

Considerations

- The point-in-time restore is possible only if the Continuous WAL archiving for point-in-time restore option is enabled for the database cluster. For details, see “Configuring PostgreSQL application backup options” on page 268.
- *Only if you use the systemd service files to manage PostgreSQL services.* After restoring the data to a new database cluster, HYCU does not recreate the

corresponding `systemd` service files. If you want to keep using the `systemd` service files, make sure to manually create them. For instructions, see PostgreSQL documentation.

- After you perform the point-in-time restore of the primary database cluster data, make sure to synchronize the standby database clusters with the primary one. For instructions, see PostgreSQL documentation.

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

Procedure

1. In the Applications panel, click the application whose data you want to restore to open the Detail view. The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open the Detail view.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click **Restore**. The Restore PostgreSQL dialog box opens.
4. Select **Restore Database Cluster**, and then click **Next**.
5. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**
6. *Only if you selected the latest restore point and the Continuous WAL archiving for point-in-time restore option is enabled for the database cluster.* Enable the **Complete recovery** switch if you want the database cluster to be recovered to the latest possible state from the existing WAL segment files.
7. *Optional and only if the Continuous WAL archiving for point-in-time restore option is enabled for the database cluster.* Specify a point in time to which you want to restore the data. The database cluster will be restored to the state it was in at the specified time.
8. Click **Next**.

9. From the Destination Database Cluster drop-down menu, select the database cluster to which you want to restore the data. In the case of a standby database cluster, you must select its primary database cluster as the destination database cluster.
10. Enable the **Overwrite existing database cluster** switch if you want to overwrite the database cluster data when performing the restore.

    If you keep this option disabled, the data is not overwritten—HYCU first adds the `_hycu_<Timestamp>` suffix to the existing contents of the data and tablespace folders, and then performs the restore to the selected database cluster.
11. *Only if you decided not to overwrite the existing database cluster.* Follow these steps:
    a. Enable the **Create new database cluster** switch if you want to create a new database cluster to which the data will be restored.

       Otherwise, HYCU automatically stops the destination database cluster and performs the restore.
    b. *Only if you decided to create a new database cluster.* Do the following:
       i. In the Restore Suffix field, enter the suffix that you want to be added to the restored contents of the data and tablespace folders on the database cluster.

          > ⓘ Important If you leave the default value, the `_hycu_<Timestamp>` suffix is automatically added to the restored contents of the data and tablespace folders.

       ii. Enter the port of the newly created database cluster.
12. Enable the **Leave database cluster in stopped state** switch if you do not want the database cluster to be started automatically after the restore.
13. Click **Restore**.

Chapter 6

Protecting file shares

HYCU enables you to protect your file share data with fast and reliable backup and restore operations. After you back up a file share, you can choose to restore either the whole file share or the files.

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

- "Configuring file share backup options" below
- "Backing up file shares" on page 329
- "Restoring file share data" on page 332

Configuring file share backup options

Before you start protecting file shares, you can adjust file share protection to the needs of your data protection environment by configuring backup options.

Backup option	Description
Exclude files or folders from backup	You can specify any file or folder to be excluded from the file share backup.
Incremental forever backup	If you want all your file share backups after the initial full backup to be only incremental backups, you can configure HYCU to track and store only the changes since the last backup (as incremental backups), allowing you to operate with a single backup chain.

Considerations

- *Only if you exclude files or folders from backup.* Files and folders cannot be excluded from a snapshot. Therefore, if you restore the whole file share from a snapshot, the excluded files and folders will also be restored.

- *Only if you enable the Incremental forever backup option.* To maintain a single backup chain, HYCU analyzes the whole backup chain after every backup (that is, it searches for any restore points with the Backup and Copy tiers whose retention period has expired) and then performs the automatic merge of the relevant restore points, keeping only a limited number of restore points. This means that data from several incremental backups that are marked for expiration is combined and merged into the next restore point in the backup chain. However, you can at any time merge restore points also manually by expiring the restore point tiers. For instructions, see "Managing data retention" on page 501.

> Accessing the Shares panel
>
> To access the Shares panel, in the navigation pane, click **Shares**.

Procedure

1. In the Shares panel, select the file share for which you want to configure backup options.
2. Click **Configuration**.

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

I want to...	Instructions
Exclude one or more file share files or folders from the backup.	In the Exclude Files or Folders from Backup field, do the following: a. Enter the path to the file or the folder that you want to exclude from the backup. The path can be either a full path or a glob pattern. For example, `/backup` or `/*/*/backup`. b. Click ✚ **Add**. c. If required, repeat these steps for each additional file or folder that you want to exclude. The paths to all the files or folders that you excluded from the backup are added to the Exclude File or Folder Paths list. If you want to remove any of them from the exclude list, click ✕.
Enable the incremental forever backup.	Enable the **Incremental forever backup** switch.

4. Click **Save**.

Backing up file shares

With HYCU, you can back up file shares in a fast and efficient way.

Prerequisites

- The backup operator must have the backup operator privileges or full permissions on the file shares that you plan to protect.
- *For Nutanix Files SMB file shares:* If you plan to protect files whose file names are longer than 255 characters, make sure that your file share meets the following requirements:
  - The multi-protocol access must be disabled.
  - The Allow Long Names for Files and Folders option must be enabled.

  For instructions, see Nutanix documentation.

Limitations

- The iSCSI, Nutanix, and tape targets cannot be used for storing file share data.
- Backing up file shares to cloud targets is supported if the file system item names contain only characters in the Unicode Basic Multilingual Plane (BMP).
- If Snapshot is defined as the backup target type in your policy, such a policy cannot be assigned to a file share.
- If you use Smart disaster recovery (DR) for Nutanix Files protection, HYCU enables you to protect replicated file share data. After you add a recovery file server as a source to HYCU, you can back up the corresponding file shares by assigning policies to them, and later also restore them. Keep in mind that you cannot restore data to replicated file shares. For details on how to configure Smart DR, see Nutanix documentation.
- *For NFS file shares:* Backing up files whose file names contain non-UTF-8 multilingual characters (for example, those created by Windows clients) is not supported. Therefore, such files will be skipped during the backup.
- *For NetApp ONTAP file shares:* Backing up qtrees inside an ONTAP volume is not supported.
- *For NetApp ONTAP SMB file shares:* Backing up root volumes is not supported.
- *For Nutanix Files:*
  - Backing up unified namespaces is not supported.
  - Backing up the NFSv4 access control lists is not supported.

Considerations

- You can change the number of incremental file share backups after which a full reindex is performed by customizing the `afs.reindex.interval.count` configuration setting. This allows you to speed up the process of searching for the relevant files when you are restoring them. For details on how to do this, see “Customizing HYCU configuration settings” on page 628.
- By default, if up to 10,000 and 1% of file backups fail during the backup of a file share, the backup status of the file share is marked as Completed with errors (and not as Failed). You can customize these values by editing the following configuration settings:
  - `afs.partial.success.threshold.count`

  - `afs.partial.success.max.fail.fraction`

  For details on how to do this, see "Customizing HYCU configuration settings" on page 628.
- When backing up a file share, HYCU also backs up any nested shares that are inside the selected file share. Keep in mind that backing up nested file shares individually is not supported.
- Backing up file shares with tiered files is supported. However, consider the following:
  - The backup operators or the HYCU instance IP addresses must not be set up as zero users or clients because this could cause backup data corruption on tiered files.
  - Additional fees may apply for backup and restore operations due to data egress.
- The following types of files will not be backed up:
  - Block special
  - Character special
  - FIFO
  - Socket
- *Only if you enabled the Incremental forever backup option for a file share.* When you assign a policy to the file share, the new backup chain settings defined in the policy will be ignored and full backups will not be performed.
- *For Nutanix Files:* Backing up connected file shares is supported. Keep in mind that connected file shares must be backed up individually because the backup of a parent file share does not include the contents of child file shares.
- *For Dell PowerScale OneFS:* When the incremental backup is run, the status of the Create snapshot job shows the values obtained from the Dell PowerScale OneFS file server.
- *For Nutanix Files and Dell PowerScale OneFS*: When backing up file shares, HYCU detects quarantined files and skips them during the backup.

Recommendations

- Using an NFS target for storing file share data requires you to enable public access to the target. For security purposes, it is recommended that you avoid such a configuration.

- *For generic file shares:* Avoid making any changes to your file shares during the backup process.

> Accessing the Shares panel
>
> To access the Shares panel, in the navigation pane, click **Shares**.

Procedure

1. In the Shares panel, select the file shares that you want to back up.

   > Tip You can update the list of file shares by clicking **Refresh**. To narrow down the list of displayed file shares, you can use the filtering options described in "Filtering and sorting data" on page 406.

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

After you assign the policy, the backup is scheduled according to the values that you defined for your policy. If required, you can also perform a manual backup at any time. For details, see "Performing a manual backup" on page 424.

> Tip If you have more than one HYCU instance in your data protection environment, you can see which HYCU instance performed a backup by clicking the preferred backup job in the Jobs panel and checking the HYCU instance IP address in the Detail view.

Restoring file share data

You can restore file share data to the original or a different file server, to a bucket, to an external SMB or NFS file share, or to a local machine. Depending on what you want to restore, you can select one of the following restore methods:

I want to restore...	Restore method	Instructions
A whole file share or individual files from the selected restore point.	Restore Files	"Restoring a whole file share or individual files" on page 335
Specific files by searching for	Search	"Restoring specific files"

I want to restore...	Restore method	Instructions
them and selecting the ones that will be included in the restore. For example, this is useful if a file was deleted some time ago and several restore points were created since then. It helps you identify which restore point contains the most recent version of the file.		on page 338

Prerequisites

- *For restoring data to a different file server share:* The file server with the file share to which you want to restore data must be added to HYCU. For details on how to do this, see "Adding a file server" on page 85.
- *For restoring data to a bucket:* The object server that hosts the bucket to which you want to restore data must be added to HYCU. For details on how to do this, see "Adding an object server" on page 98.
- *Only if you are restoring data that is stored in the archive access tier on an Azure target.* If you plan to restore data to a local machine, you must rehydrate the data. For details on how to do this, see Azure documentation.

Limitations

- Restoring the file share data from a snapshot is not supported for generic file shares.
- The restore of alternate data streams (ADS) is supported only if you are restoring data from one file server SMB share to another file server SMB share.
- *For Azure Files:* The restore of ADS is not supported.
- *Only if restoring Nutanix Files shares that contain ADS in top-level directories to distributed file shares.* Restoring ADS to top-level directories of distributed file shares is not supported. ADS can be restored to subdirectories of distributed file shares or to standard file shares.
- Symbolic links are restored only when restoring data from one NFS file share to another NFS file share, or from one SMB file share to another SMB

file share.

- *Only if restoring files to an external file share.* Restoring files or folders with newlines in their names is supported only for an NFS share set up on Unix.
- *Only if restoring files to a local machine:*
  - The files can be restored only if the size of the uncompressed files is less than or equal to 2 GiB.
  - Restoring the original access control list for the files is not supported.

Considerations

- The file share data can be restored from a target or a snapshot. Restoring data from the snapshot is possible only if the `afs.restore.snapshot.enabled` configuration setting is set to `true` (the default value is `false`). In this case, the restore is always performed from the snapshot if the snapshot is available. Otherwise, the restore is performed from the target. For details on how to customize HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- *Only if restoring a large number of files from the file share backup.* The HYCU instance may require more RAM than is available by default. In this case, increase the default value by using the `afs.instance.memory.mb` configuration setting. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
- *Only if restoring files to a local machine.* The restored files are downloaded as a `.zip` file. To avoid any potential issues with unzipping the restored files and to make sure that the files or folders with newlines in their names are properly restored, always use 7-Zip when extracting the files.
- If the number of files that could not be restored during the file share restore is less than or equal to 100 (the default value), the status of the file share restore is Warning. You can edit this default value by customizing the `afs.restore.partial.success.threshold.count` configuration setting. For details on how to do this, see "Customizing HYCU configuration settings" on page 628.

- *Only if restoring specific files by using the Search restore method.*
  - If the files that you are searching for were not saved in the selected restore point due to being unmodified, HYCU will search other restore points to locate the most recent available versions of the files.
  - If you change your search criteria after moving the files to the right column (which shows the files that are selected for the restore), and you then try to move them back to the left column (which shows all the files that match the updated search), any files that no longer match the updated search will disappear from the left column.
  - If you select to restore more than one version of the same file from different restore points, depending on where you plan to restore the files, consider the following:
    - *To a share or a bucket:* To avoid overwriting the files during the restore, make sure to select the Rename restored option. When this option is selected, the `<RestorePointTimestamp>.restored` suffix is automatically added to each restored file.
    - *To a local machine:* After the restore, each restore point with the restored files will appear as a top-level folder in the downloaded `.zip` file.

Recommendation

For optimal restore performance, it is recommended that you restore data to a file server share or to a bucket instead of an external file share whenever possible.

> ### Accessing the Shares panel
>
> To access the Shares panel, in the navigation pane, click **Shares**.

Restoring a whole file share or individual files

Procedure

1. In the Shares panel, click the file share that contains the files that you want to restore to open the Detail view.

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

2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click **Restore Files**.
4. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**
5. Click **Next**.
6. In the dialog box that opens, select the uppermost check box (the one in front of the icon) if you want to restore the whole file share. Otherwise, from the list of available folders and files, select the ones that you want to restore. Click **Next**.

   > Tip If there are too many files to be displayed on one page, you can move between the pages by clicking < and >. You can also use ⌄ to set the number of files to be displayed per page.

7. Depending on where you want to restore the selected files (to the original or a different file server share, to a bucket, to an external SMB or NFS file share, or the local machine), select the preferred restore option, click **Next**, and then follow the instructions:

Restore option	Instructions
Restore to Share or Bucket	a. From the Restore To drop-down menu, select the file server share or the bucket to which you want to restore the files. b. Select whether you want to restore the files to the original location or an alternate location on the same file server share. If you select an alternate location, in the Path field, specify the part of the path that is appended to the location of the selected share. > Note Under the Path field, the full alternate location path is displayed and updated as you

Restore option	Instructions
	type. c. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file). **Important** If you plan to rename the original files, you must be a file server admin. For all other operations, you can be either a file server or a backup admin. d. *Only if restoring files from one SMB file share to another SMB file share.* Enable the **Restore ACLs** switch if you want to restore the original access control lists for the files. e. Click **Restore**.
Restore to External Share	From the Share Type drop-down menu, select where you want to restore the files, and then provide the required information: • **NFS** a. Enter the path to the NFS shared folder in the following format: `\\server\<Path>` b. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file). c. Click **Restore**. • **SMB** a. Enter the path to the SMB shared folder in the following format:

Restore option	Instructions
	`\\server\<Path>` b. *Optional.* Provide user credentials to access the SMB file share. c. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file). d. *Only if restoring files from an SMB file share to an external SMB file share.* Enable the **Restore ACLs** switch if you want to restore the original access control list for the files. > **Important** If you enable the Restore ACLs switch, keep in mind that the restored files might not be accessible due to the ACLs not being recognized on the destination SMB file share. e. Click **Restore**.
Download	Click **Download** to restore the selected files to the local machine. > **Important** Do not refresh the page or navigate away from the page until the download process job finishes.

Restoring specific files

Procedure

1. In the Shares panel, click the file share that contains the files that you want to restore to open the Detail view.

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

2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click **Search**.
4. In the Search field, enter the name of the file or files that you want to search for, and then press **Enter**. You can also search for the files by using the wildcard characters. For details on wildcard characters, see "Searching for specific files by using wildcard characters" on page 341.

   HYCU scans the backup chain by starting at the selected restore point and proceeds through each restore point that was created before the selected restore point. All the matching files are displayed. If required, click **Load more** to load the remaining files.

   > **Note** If there is a large number of the restore points in the backup chain, searching for the files may take some time to complete.

5. Select one or more files that you want to restore, and then click > to move the selected files to the column on the right that shows which files will be included in the restore.
6. Click **Next**.
7. Depending on where you want to restore the selected files, select the preferred restore option, click **Next**, and then follow the instructions:

Restore option	Instructions
Restore to Share or Bucket	a. From the Restore To drop-down menu, select the file server share or the bucket to which you want to restore the files. b. Select whether you want to restore the files to the original location or an alternate location on the same file server share. If you select an alternate location, in the Path field, specify the part of the path that is appended to the location of the selected share. **Note** Under the Path field, the full alternate location path is displayed and updated as you type. c. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the

Restore option	Instructions
	file, skip the file, rename the original file, or rename the restored file). > ⓘ Important If you plan to rename the original files, you must be a file server admin. For all other operations, you can be either a file server or a backup admin. d. *Only if restoring files from one SMB file share to another SMB file share.* Enable the **Restore ACLs** switch if you want to restore the original access control lists for the files. e. Click **Restore**.
Restore to External Share	From the Share Type drop-down menu, select where you want to restore the files, and then provide the required information: • **NFS** a. Enter the path to the NFS shared folder in the following format: `\\server\<Path>` b. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file). c. Click **Restore**. • **SMB** a. Enter the path to the SMB shared folder in the following format: `\\server\<Path>` b. *Optional.* Provide user credentials to access the SMB file share. c. Specify which action should be performed

Restore option	Instructions
	during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file). d. *Only if restoring files from an SMB file share to an external SMB file share.* Enable the **Restore ACLs** switch if you want to restore the original access control list for the files. > ⓘ Important If you enable the Restore ACLs switch, keep in mind that the restored files might not be accessible due to the ACLs not being recognized on the destination SMB file share. e. Click **Restore**.
Download	Click **Download** to restore the selected files to the local machine. > ⓘ Important Do not refresh the page or navigate away from the page until the download process job finishes.

Searching for specific files by using wildcard characters

You can use the following wildcard characters to search for specific files by glob pattern instead of names:

Wildcard character	Description
Slash (/)	Acts as a separator. A segment is everything that comes between two separators.
Asterisk (*)	Matches zero or more characters within one segment. It is used for globbing the files within one folder. For example, a/*txt would match a/hycu.txt, but not a/b/hycu.txt.
Question mark (?)	Matches a single character within one segment. It can be used for matching the files or folders that differ by only one character. For example, ?ycu.txt would

Wildcard character	Description
	match hycu.txt.
Square brackets ([ ])	Matches one character within the specified set. For example, [a-z]ycu.txt would match hycu.txt.

Chapter 7

Protecting volume groups

HYCU enables you to protect Nutanix volume groups with fast and reliable backup and restore operations. After you back up a volume group, you can choose to restore either the whole volume group or only individual virtual disks by exporting them to an NFS or SMB share.

> **Important** If the volume groups are attached to one or more virtual machines at backup time, they are backed up automatically during the virtual machine backup. For details, see "Protecting virtual machines" on page 154.

For details on how to protect volume groups efficiently, see the following sections:

- "Backing up volume groups" below
- "Restoring volume groups" on the next page

Backing up volume groups

With HYCU, you can back up Nutanix volume groups in a fast and efficient way.

Prerequisite

A Nutanix cluster on which the volume group that you want to protect resides must be added to HYCU. For instructions, see "Adding a Nutanix cluster" on page 74.

Consideration

The volume groups that HYCU creates automatically and uses for data protection purposes are not shown in the Volume Groups panel. The names of these volume groups start with the `NTNX-`, `hycu-vg-`, and `HYCU-` prefixes, therefore make sure not to create your own volume groups with the same prefixes.

> Accessing the Volume Groups panel
>
> To access the Volume Groups panel, in the navigation pane, click **Volume Groups**.

Procedure

1. In the Volume Groups panel, select the volume groups that you want to back up.

   > Tip You can update the list of volume groups by clicking **Refresh**. To narrow down the list of displayed volume groups, you can use the filtering options described in "Filtering and sorting data" on page 406.

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

The backup is scheduled according to the values that you defined for your policy. If required, you can also perform a manual backup of any volume group at any time. For details, see "Performing a manual backup" on page 424.

Restoring volume groups

HYCU enables you to restore either a whole volume group or only individual virtual disks that became corrupted.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for restoring data.
- *Only if you plan to restore data from a QStar tape target.* During the restore procedure, you can use the Tape Info option to view information about the media. If a note indicating that an empty response was received from QStar is displayed, make sure to check whether the data was archived to the media.

Restore options

You can select among the following restore options:

Restore option	Description
Restore Volume Group	Enables you to restore a volume group. Select this option if you want to replace the original volume group with the restored one. For instructions, see “Restoring a volume group” below.
Clone Volume Group	Enables you to restore a volume group by creating its clone. Select this option if you want to keep the original volume group. For instructions, see “Cloning a volume group” on page 347.
Export vDisks	Enables you to restore virtual disks to an NFS or SMB share. Select this option if you want to make the virtual disks available to users with specific access permissions, or if you want to use the virtual disks later to restore data to an environment with a source not supported by HYCU or not added to HYCU. For instructions, see “Exporting virtual disks” on page 348.

Accessing the Volume Groups panel

To access the Volume Groups panel, in the navigation pane, click **Volume Groups**.

Restoring a volume group

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

Consideration

Only if the volume group is attached to one or more virtual machines. The virtual machines to which the volume group is attached must be turned off.

Limitation

You cannot restore volume groups that are attached to an NDB-managed database server. To restore such volume groups, use the Restore vDisks option. See “Restoring virtual disks” on page 233.

Procedure

1. In the Volume Groups panel, click the volume group that you want to restore. The Detail view appears at the bottom of the screen.

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Restore Volume Group**, and then click **Next**.
5. From the Storage container drop-down menu, select where you want to restore the volume group. By default, the original storage container is selected.
6. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**

   > **Note** *Only if you select the Archive tier and the data is stored on a QStar tape target.*
   > To view tape target information, click **Tape Info**. The following information is displayed:
   >
   > - Target name and UUID
   > - Barcode label of the media
   > - Information on whether the media is online or offline
   > - Location of the media
   >
   > You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

7. *Only if the volume group you are restoring is attached to one or more virtual machines.* Enable the **Attach volume group** switch if you want the volume group to be attached to the virtual machines after the restore.
8. Click **Restore**.

Cloning a volume group

You can create a clone of the original volume group by restoring the volume group to its original or a new location. In this case, the original volume group will not be overwritten.

Procedure

1. In the Volume Groups panel, click the volume group that you want to restore. The Detail view appears at the bottom of the screen.

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Clone Volume Group**, and then click **Next**.
5. From the Storage container drop-down menu, select where you want to restore the volume group. By default, the original storage container is selected.
6. In the New volume group name field, specify a new name for the volume group.
7. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**

   > **Note** *Only if you select the Archive tier and the data is stored on a QStar tape target.*
   > To view tape target information, click **Tape Info**. The following information is displayed:
   >
   > - Target name and UUID
   > - Barcode label of the media
   > - Information on whether the media is online or offline
   > - Location of the media

> You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

8. *Only if the volume group you are restoring is attached to one or more virtual machines.* Enable the **Attach volume group** switch if you want the volume group to be attached to the virtual machines after the restore.
9. Click **Restore**.

Exporting virtual disks

You can restore virtual disks to an NFS or SMB share.

Procedure

1. In the Volume Groups panel, click the volume group whose virtual disks you want to restore. The Detail view appears at the bottom of the screen.

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

2. In the Detail view, select the preferred restore point.
3. Click **Restore**.
4. Select **Export vDisks**, and then click **Next**.
5. From the list of virtual disks that are available for the restore, select the ones that you want to restore, and then click **Next**.
6. From the Type drop-down menu, select where you want to restore the virtual disks, and then provide the required information:

Type	Instructions
SMB	a. *Optional.* Enter the domain and user credentials. b. Enter the SMB server host. > **Important** To ensure that HYCU can use Kerberos authentication, you must enter the fully qualified domain name (FQDN). c. Enter the path to the SMB shared folder from the root of the server (for example, `/backups/HYCU`).

Type	Instructions
NFS	a. Enter the NFS server name or IP address. b. Enter the path to the NFS shared folder from the root of the server (for example, `/backups/HYCU`).

7. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**

   > Note *Only if you select the Archive tier and the data is stored on a QStar tape target.*
   > To view tape target information, click **Tape Info**. The following information is displayed:
   >
   > - Target name and UUID
   > - Barcode label of the media
   > - Information on whether the media is online or offline
   > - Location of the media
   >
   > You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

8. Click **Restore**.

Protecting buckets

HYCU enables you to protect your bucket data with fast and reliable backup and restore operations. After you back up a bucket, you can choose to restore either the whole bucket or the objects.

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

- "Configuring bucket backup options" below
- "Backing up buckets" on the next page
- "Restoring bucket data" on page 353

Configuring bucket backup options

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

Backup option	Description
Exclude objects from backup	You can specify any object to be excluded from the bucket backup.
Incremental forever backup	If you want all your bucket backups after the initial full backup to be only incremental backups, you can configure HYCU to track and store only the changes since the last backup (as incremental backups), allowing you to operate with a single backup chain.

Consideration

Only if you enable the Incremental forever backup option. To maintain a single backup chain, HYCU analyzes the whole backup chain after every backup (that is, it searches for any restore points with the Backup and Copy tiers whose retention period has expired) and then performs the automatic merge of the relevant restore points, keeping only a limited number of restore points. This means that data from several incremental backups that are marked for expiration is combined and merged into the next restore point in the backup chain. However, you can at any time merge restore points also manually by

expiring the restore point tiers. For instructions, see “Managing data retention” on page 501.

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

Procedure

1. In the Buckets panel, select the bucket for which you want to configure backup options.
2. Click **Configuration**.
3. Depending on what you want to do, perform the required action:

I want to...	Instructions
Exclude one or more bucket objects from the backup.	In the Exclude Objects from Backup field, do the following: a. Enter the path to the object that you want to exclude from the backup. The path can be either a full path or a glob pattern. For example, `/backup` or `/*/*/backup`. b. Click **Add**. c. If required, repeat these steps for each additional object that you want to exclude. The paths to all the objects that you excluded from the backup are added to the Exclude Object Paths list. If you want to remove any of them from the exclude list, click ✕.
Enable the incremental forever backup.	Enable the **Incremental forever backup** switch.

4. Click **Save**.

Backing up buckets

With HYCU, you can back up buckets in a fast and efficient way.

Limitations

- The iSCSI, Nutanix, and tape targets cannot be used for storing bucket data.
- Backing up buckets to cloud targets is supported if the object names contain only characters in the Unicode Basic Multilingual Plane (BMP).
- If Snapshot is defined as the backup target type in your policy, such a policy cannot be assigned to a bucket.
- HYCU does not support backing up objects whose names contain the leading slash character.

Considerations

- You can change the number of incremental bucket backups after which a full reindex is performed by customizing the `afs.reindex.interval.count` configuration setting. This allows you to speed up the process of searching for the relevant objects when you are restoring them. For details on how to do this, see "Customizing HYCU configuration settings" on page 628.
- By default, if up to 10,000 and 1% of object backups fail during the backup of a bucket, the backup status of the bucket is marked as Completed with errors (and not as Failed). You can customize these values by editing the following configuration settings:
  - `afs.partial.success.threshold.count`
  - `afs.partial.success.max.fail.fraction`

  For details on how to do this, see "Customizing HYCU configuration settings" on page 628.
- Backing up buckets with tiered objects is supported. However, additional fees may apply for backup and restore operations due to data egress.
- A bucket that has WORM enabled is represented by the icon in the list of buckets.
- *Only if you enabled the Incremental forever backup option for a bucket.* When you assign a policy to the bucket, the new backup chain settings defined in the policy will be ignored and full backups will not be performed.

Recommendations

- Using an NFS target for storing bucket data requires you to enable public access to the target. For security purposes, it is recommended that you avoid such a configuration.
- Avoid making any changes to your buckets during the backup process.

Accessing the Buckets panel

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

Procedure

1. In the Buckets panel, select the buckets that you want to back up.

   Tip You can update the list of buckets by clicking **Refresh**. To narrow down the list of displayed buckets, you can use the filtering options described in "Filtering and sorting data" on page 406.

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

After you assign the policy, the backup is scheduled according to the values that you defined for your policy. If required, you can also perform a manual backup at any time. For details, see "Performing a manual backup" on page 424.

Tip If you have more than one HYCU instance in your data protection environment, you can see which HYCU instance performed a backup by clicking the preferred backup job in the Jobs panel and checking the HYCU instance IP address in the Detail view.

Restoring bucket data

You can restore bucket data to the original or a different bucket, to a file server, to an external SMB or NFS file share, or to a local machine. Depending on what you want to restore, you can select one of the following restore methods:

I want to restore...	Restore method	Instructions
A whole bucket or individual objects from the selected restore point.	Restore Objects	"Restoring a whole bucket or individual objects" on page 356
Specific objects by searching for them and selecting the ones that will be included in the restore.	Search	"Restoring specific objects" on page 358

I want to restore...	Restore method	Instructions
For example, this is useful if an object was deleted some time ago and several restore points were created since then. It helps you identify which restore point contains the most recent version of the object.		

Prerequisites

- *For restoring data to a different bucket:* The object server that hosts the bucket to which you want to restore data must be added to HYCU. For details on how to do this, see "Adding an object server" on page 98.
- *For restoring data to a file server share:* The file server with the file share to which you want to restore data must be added to HYCU. For details on how to do this, see "Adding a file server" on page 85.
- *Only if you are restoring data that is stored in the archive access tier on an Azure target.* If you plan to restore data to a local machine, you must rehydrate the data. For details on how to do this, see Azure documentation.

Limitations

- Restoring bucket data from a snapshot is not supported.
- *Only if restoring objects to an external file share.* Restoring objects with newlines in their names is supported only for an NFS share set up on Unix.
- *Only if restoring objects to a local machine:*
  - The objects can be restored only if the size of the uncompressed objects is less than or equal to 2 GiB.
  - Restoring the original access control list for the objects is not supported.
- *Only if restoring objects to an alternate location.* The file destination path must not contain the leading slash character.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for

restoring data.

- *Only if restoring a large number of objects from the bucket backup.* The HYCU instance may require more RAM than is available by default. In this case, increase the default value by using the `afs.instance.memory.mb` configuration setting. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
- *Only if restoring objects to a local machine.* The restored objects are downloaded as a `.zip` file. To avoid any potential issues with unzipping the restored objects and to make sure that the objects with newlines in their names are properly restored, always use 7-Zip when extracting the objects.
- If the number of objects that could not be restored during the bucket restore is less than or equal to 100 (the default value), the status of the bucket restore is Warning. You can edit this default value by customizing the `afs.restore.partial.success.threshold.count` configuration setting. For details on how to do this, see "Customizing HYCU configuration settings" on page 628.
- *Only if restoring specific objects by using the Search restore method.*
  - If the objects that you are searching for were not saved in the selected restore point due to being unmodified, HYCU will search other restore points to locate the most recent available versions of the objects.
  - If you change your search criteria after moving the objects to the right column (which shows the objects that are selected for the restore), and you then try to move them back to the left column (which shows all the objects that match the updated search), any objects that no longer match the updated search will disappear from the left column.
  - If you select to restore more than one version of the same object from different restore points, depending on where you plan to restore the objects, consider the following:
    - *To a bucket or a share:* To avoid overwriting the objects during the restore, make sure to select the Rename restored option. When this option is selected, the `<RestorePointTimestamp>.restored` suffix is automatically added to each restored object.
    - *To a local machine:* After the restore, each restore point with the restored objects will appear as a top-level folder in the downloaded `.zip` file.

Recommendation

For optimal restore performance, it is recommended that you restore data to a bucket or to a file server share instead of an external file share whenever possible.

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

Restoring a whole bucket or individual objects

Procedure

1. In the Buckets panel, click the bucket that contains the objects that you want to restore to open the Detail view.

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

2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click **Restore Objects**.
4. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest restore to the latest state.
   - **Backup**
   - **Copy**
   - **Archive**
5. Click **Next**.
6. In the dialog box that opens, select the uppermost check box (the one in front of the icon) if you want to restore the whole bucket. Otherwise, from the list of available objects, select the ones that you want to restore. Click **Next**.

   > **Tip** If there are too many objects to be displayed on one page, you can move between the pages by clicking and . You can also use to set the number of objects to be displayed per page.

7. Depending on where you want to restore the selected objects (to the original or a different bucket, a file server share, an external SMB or NFS file share,

or the local machine), select the preferred restore option, click **Next**, and then follow the instructions:

Restore option	Instructions
Restore to Bucket or Share	a. From the Restore To drop-down menu, select the bucket or the file server share to which you want to restore the objects. b. Select whether you want to restore the objects to the original location or an alternate location on the same bucket. If you select an alternate location, in the Path field, specify the part of the path that is appended to the location of the selected bucket. > Note Under the Path field, the full alternate location path is displayed and updated as you type. c. Specify which action should be performed during the restore operation if an object with the same name already exists in the selected location (overwrite the object, skip the object, rename the original object, or rename the restored object). d. *Only if restoring objects from one bucket to another bucket.* Enable the **Restore ACLs** switch if you want to restore the original access control lists for the objects. e. *Only if restoring objects from one bucket to another bucket.* Enable the **Restore S3 object tags** switch if you want to restore the original object tags for the objects. f. Click **Restore**.
Restore to External Share	From the Share Type drop-down menu, select where you want to restore the objects, and then provide the required information: • **NFS** a. Enter the path to the NFS shared folder in the following format:

Restore option	Instructions
	```\\server\<Path>``` b. Specify which action should be performed during the restore operation if an object with the same name already exists in the selected location (overwrite the object, skip the object, rename the original object, or rename the restored object). c. Click **Restore**. • **SMB** a. Enter the path to the SMB shared folder in the following format: ```\\server\<Path>``` b. *Optional.* Provide user credentials to access the SMB file share. c. Specify which action should be performed during the restore operation if an object with the same name already exists in the selected location (overwrite the object, skip the object, rename the original object, or rename the restored object). d. Click **Restore**.
Download	Click **Download** to restore the selected objects to the local machine. **Important** Do not refresh the page or navigate away from the page until the download process job finishes.

Restoring specific objects

Procedure

1. In the Buckets panel, click the bucket that contains the objects that you want to restore to open the Detail view.

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

view.

2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click **Search**.
4. In the Search field, enter the name of the object or objects that you want to search for, and then press **Enter**. You can also search for the objects by using the wildcard characters. For details on wildcard characters, see "Searching for specific objects by using wildcard characters" on page 361.

   HYCU scans the backup chain by starting at the selected restore point and proceeds through each restore point that was created before the selected restore point. All the matching objects are displayed. If required, click **Load more** to load the remaining objects.

   > **Note** If there is a large number of the restore points in the backup chain, searching for the objects may take some time to complete.

5. Select one or more objects that you want to restore, and then click > to move the selected objects to the column on the right that shows which objects will be included in the restore.
6. Click **Next**.
7. Depending on where you want to restore the selected objects, select the preferred restore option, click **Next**, and then follow the instructions:

Restore option	Instructions
Restore to Bucket or Share	a. From the Restore To drop-down menu, select the bucket or the file server share to which you want to restore the objects. b. Select whether you want to restore the objects to the original location or an alternate location on the same bucket. If you select an alternate location, in the Path field, specify the part of the path that is appended to the location of the selected bucket. > **Note** Under the Path field, the full alternate location path is displayed and updated as you type. c. Specify which action should be performed during the restore operation if an object with the same

Restore option	Instructions
	name already exists in the selected location (overwrite the object, skip the object, rename the original object, or rename the restored object). d. *Only if restoring objects from one bucket to another bucket.* Enable the **Restore ACLs** switch if you want to restore the original access control lists for the objects. e. *Only if restoring objects from one bucket to another bucket.* Enable the **Restore S3 object tags** switch if you want to restore the original object tags for the objects. f. Click **Restore**.
Restore to External Share	From the Share Type drop-down menu, select where you want to restore the objects, and then provide the required information: • **NFS** a. Enter the path to the NFS shared folder in the following format: `\\server\<Path>` b. Specify which action should be performed during the restore operation if an object with the same name already exists in the selected location (overwrite the object, skip the object, rename the original object, or rename the restored object). c. Click **Restore**. • **SMB** a. Enter the path to the SMB shared folder in the following format: `\\server\<Path>` b. *Optional.* Provide user credentials to access the SMB file share.

Restore option	Instructions
	c. Specify which action should be performed during the restore operation if an object with the same name already exists in the selected location (overwrite the object, skip the object, rename the original object, or rename the restored object). d. Click **Restore**.
Download	Click **Download** to restore the selected objects to the local machine. **Important** Do not refresh the page or navigate away from the page until the download process job finishes.

Searching for specific objects by using wildcard characters

You can use the following wildcard characters to search for specific objects by glob pattern instead of names:

Wildcard character	Description
Slash (/)	Acts as a separator. A segment is everything that comes between two separators.
Asterisk (*)	Matches zero or more characters within one segment. For example, a/*txt would match a/hycu.txt, but not a/b/hycu.txt.
Question mark (?)	Matches a single character within one segment. It can be used for matching the objects that differ by only one character. For example, ?ycu.txt would match hycu.txt.
Square brackets ([ ])	Matches one character within the specified set. For example, [a-z]ycu.txt would match hycu.txt.

Chapter 9

Recovering your data protection environment

If a disaster occurs in your data protection environment and your data is corrupted or unavailable, HYCU provides an effective approach to recovering data by importing targets on which backup data is stored. You can decide to recover the following:

- HYCU backup controller and use it later to restore data
- Virtual machines, applications, file shares, volume groups, and buckets

The procedures described in this section are applicable if your backup data is stored on targets. If you selected Snapshot as the backup target type in your policy and no data archives exist, you can use the snapshots created by HYCU to perform disaster recovery through the management console of the environment in which HYCU is deployed. For details, see your platform documentation.

Considerations

- Azure Local and Hyper-V cluster data can be stored only on targets.
- After performing the disaster recovery, not all your restore points might be used for restoring data because of the expected discrepancy between data in the database and data on targets.
- *For DR-ready virtual machines and applications:* You can recover your HYCU backup controller and protected data to cloud by using the SpinUp functionality. For more information, see “Protecting data across on-premises and cloud environments” on page 586.

Procedures

1. Prepare for disaster recovery. For instructions, see “Preparing for disaster recovery” on the next page.
2. Perform disaster recovery. For instructions, see “Performing disaster recovery” on page 368.

3. *Only if you use HYCU instances:* Create HYCU instances. For more information, see “HYCU instances” on page 25.

Preparing for disaster recovery

Prerequisites

- You must know configuration parameters of the targets that store the backup of your original HYCU backup controller or backups of other entities you want to recover. For details, see “Preparing for disaster recovery of your data” on page 161.
- The targets that store backup data of the entities you want to recover must be accessible to the source where you plan to deploy a recovery HYCU backup controller.
- *Only if the backup of the original HYCU backup controller is stored on an iSCSI target.* The iSCSI storage device must be dedicated to a single HYCU backup controller and no other appliances than HYCU.
- *Only if the backup of the original HYCU backup controller or virtual machines, applications, file shares, volume groups, and buckets that you want to recover is stored on a Google Cloud target.* A Google Cloud service account must be created and added to HYCU. For instructions on how to add a cloud account to HYCU, see “Adding a Google Cloud service account” on page 466.
- *Only if the backup of the original HYCU backup controller or other entities you want to recover is stored on a target with enabled target encryption.* The encryption target key from the original HYCU backup controller must be exported and the file containing the encryption key must be available.

When preparing for disaster recovery, you must perform the following tasks:

Task	Instructions
1. Deploy a recovery HYCU backup controller.	“Deploying a recovery HYCU backup controller” on the next page
2. Import the targets that store the backup of the original HYCU backup controller. The imported targets may also contain backups of virtual machines, applications, file	“Importing targets” on page 366

Task	Instructions
shares, volume groups, and buckets.	
3. Add a source to which you plan to restore your HYCU backup controller. If you plan to restore also virtual machines, applications, file shares, volume groups, and buckets, add the sources to which you plan to restore them.	"Adding sources" on page 73

Deploying a recovery HYCU backup controller

A recovery HYCU backup controller is a HYCU backup controller that you deploy for the purposes of performing disaster recovery. After you successfully import the first target, the HYCU backup controller is automatically put in recovery mode. For the HYCU backup controller in recovery mode, the following applies:

- HYCU automatically refreshes the imported targets every 60 minutes to get the information about the latest restore points (the backups stored on the targets), as well as the information about the targets that are available for importing or that have been removed.

  > **Note** You can at any time refresh the imported targets also manually. To do this, in the Targets panel, click **Refresh**.

- Backup operations are disabled. This means that you cannot assign policies, perform manual backups, or expire backups manually.
- Setting power options is disabled.
- Only limited target options can be edited.
- Adding targets is disabled.

Depending on your data protection needs, you can decide to keep or delete the recovery HYCU backup controller after you perform disaster recovery. If you delete the recovery HYCU backup controller, you will have to deploy a new one the next time you perform disaster recovery.

Limitation

For Nutanix and iSCSI targets: Keeping the recovery HYCU backup controller is not supported. If you want to use such targets for disaster recovery, you must deploy a new recovery HYCU backup controller every time.

Considerations

- For successful target synchronization, the recovery HYCU backup controller must be deployed with HYCU version 4.5.0 or later.
- Deactivated targets are excluded from target synchronization.
- The default automatic target synchronization value can be adjusted to your data protection needs. For details on how to customize HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.

Procedure

1. Sign in to the management console of the environment to which you want to deploy the HYCU backup controller that you will use for restoring the original HYCU backup controller or other entities.
2. Depending on the environment to which you want to deploy the HYCU backup controller, follow the instructions to deploy the HYCU virtual appliance in the HYCU Backup Controller mode:

I want to deploy the HYCU backup controller to...	Instructions
Nutanix AHV cluster	"Deploying HYCU to a Nutanix AHV cluster" on page 45
Nutanix ESXi cluster or vSphere environment	"Deploying HYCU to a Nutanix ESXi cluster or a vSphere environment" on page 49
XenServer environment	"Deploying HYCU to a XenServer environment" on page 52
Azure Local environment	"Deploying HYCU to an Azure Local environment" on page 55
Hyper-V cluster	"Deploying HYCU to a Hyper-V cluster" on page 59
AWS GovCloud (US) environment	"Deploying HYCU to an AWS GovCloud (US) environment" on page 64
Azure environment	"Deploying HYCU to an Azure environment" on page 65

I want to deploy the HYCU backup controller to...	Instructions
Azure Government environment	"Deploying HYCU to an Azure Government environment" on page 67

3. *Only if you plan to restore the HYCU backup controller to a different source.* Enable the creation of a clone of the HYCU backup controller. To do so, in the HYCU `config.properties` file, set the `clone.enabled.for.hycu.dr` configuration setting to `true`. For instructions on how to customize HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.

   > ⚠ **Caution** Make sure that a clone of the HYCU backup controller is not activated while the original HYCU backup controller is still active. Otherwise, data loss may occur.

4. Sign in to the HYCU web user interface of the HYCU backup controller.
5. *Only if the backup of the original HYCU backup controller or backups of the entities you want to restore is stored on a target with enabled target encryption.* Import the encryption key that you have exported from the original HYCU backup controller. For instructions, see "Configuring target encryption" on page 471.

Importing targets

Prerequisites

- The activities on the original HYCU backup controller (if it still exists) must be suspended and no jobs may be running. For instructions, see "Setting power options" on page 509.
- No targets or only imported targets can exist on the recovery HYCU backup controller. Otherwise, importing targets is disabled.
- *For importing iSCSI or Nutanix targets:* The targets must be unmounted on any other powered on HYCU backup controller.
- *For importing iSCSI targets:* The HYCU iSCSI initiator secret must be added on the iSCSI server if you want to enable mutual authentication between HYCU and the iSCSI server.

- *For importing Nutanix Objects or S3 compatible targets:* If you want to provide secure HTTPS access, the CA certificate/chain must be imported to HYCU. For details, see "Importing a custom certificate" on page 527.

Limitation

Backing up data to imported targets is not supported.

Considerations

- The targets you import should contain the complete backup chains of the entities you want to recover.
- Make sure not to make any changes to HYCU until the import job is finished.

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

Procedure

1. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
2. In the Targets panel, click **Import**. The Import Target dialog box opens.
3. From the Type drop-down menu, select the type of target.
4. Specify the values so that they match the original target configuration, and then click **Next**.
5. In the Import Target Catalog dialog box, select the name of the HYCU backup controller whose backup you want to import, and then click **Next**.
6. In the Multiple Targets dialog box, one or more targets that store backup data of the selected HYCU backup controller and other entities is displayed. If any additional targets are found, select them one by one and specify the values so that they match the original target configuration. For each target, click **Validate** to check the configuration.

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

7. After you validated all the targets required for your restore, click **Import**.

   > **Note** It is recommended to import all the targets from the list to ensure that complete backup chains are available for the restore. If you do not import some targets and backup chains are not complete, you can import missing targets later by repeating the import procedure.

After a successful import of targets

- The imported targets are listed in the Targets panel and their mode is set to Read-Only, which prevents you from storing backup data to these targets.
- The HYCU backup controller is listed in the Virtual Machines panel, and its status is Protected deleted.
- For recovering virtual machines, applications, file shares, volume groups, and buckets, consider the following:
  - The self-service groups existing in the original data protection environment are recreated on the recovery HYCU backup controller. The recreated self-service groups do not contain any users. To restore virtual machines, applications, file shares, volume groups, and buckets, you need to create users and add them to the recreated user groups that have ownership over the virtual machines, file shares, volume groups, and buckets that you want to restore. For instructions, see “Setting up a user environment” on page 444.
  - The virtual machines whose backups are stored on the imported targets are listed in the Virtual Machines panel, and their status is Protected deleted. To restore virtual machines other than the HYCU backup controller, see “Restoring virtual machines” on page 184.
  - Applications whose backups are stored on the imported targets are listed in the Applications panel, and their status is Protected deleted. To restore applications, see “Restoring whole applications” on page 272.
  - File shares whose backups are stored on the imported targets are listed in the Shares panel, and their status is Protected deleted. To restore file shares, see “Restoring file share data” on page 332.
  - Volume groups whose backups are stored on the imported targets are listed in the Volume Groups panel, and their status is Protected deleted. To restore volume groups, see “Restoring volume groups” on page 344.
  - Buckets whose backups are stored on the imported targets are listed in the Buckets panel, and their status is Protected deleted. To restore buckets, see “Restoring bucket data” on page 353.

Performing disaster recovery

Perform disaster recovery by using one of the following approaches:

I want to recover…	Instructions
The HYCU backup controller to the original source by using a restore point created with HYCU version 4.0.0 or later.	"Restoring the HYCU backup controller to the original source" below **Important** You cannot restore the HYCU backup controller to an Azure Local environment or to a Hyper-V cluster by using this approach.
The HYCU backup controller to a different source by using a restore point created with HYCU version 4.0.0 or later.	"Restoring the HYCU backup controller to a different source" on page 374
The HYCU backup controller to the original or a different source by using a restore point created with a HYCU version earlier than 4.0.0.	"Exporting virtual disks" on page 238
Virtual machines	"Restoring virtual machines" on page 184
Applications	"Restoring whole applications" on page 272
File shares	"Restoring file share data" on page 332
Volume groups	"Restoring volume groups" on page 344
Buckets	"Restoring bucket data" on page 353

Restoring the HYCU backup controller to the original source

Depending on the source on which the original HYCU backup controller was running, see one of the following sections:

- "Restoring the HYCU backup controller to a Nutanix cluster or a vSphere environment" below
- "Restoring the HYCU backup controller to a XenServer environment" on the next page
- "Restoring the HYCU backup controller to an AWS GovCloud (US) environment" on page 372
- "Restoring the HYCU backup controller to an Azure or Azure Government environment" on page 373

Restoring the HYCU backup controller to a Nutanix cluster or a vSphere environment

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

Prerequisites

- The recovery HYCU backup controller must have network access to the cluster of the original HYCU backup controller.
- *Only if the backup of the original HYCU backup controller is stored on an iSCSI or a Nutanix target:* The target must be deactivated and detached from the recovery HYCU backup controller before you power on the restored HYCU backup controller.

Procedure

1. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
2. In the Virtual Machines panel, select the HYCU backup controller.
3. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

4. Click **Restore**.
5. Select **Virtual machine options**, and then click **Next**.
6. Select **Restore VM**, and then restore the HYCU backup controller by following the instructions described in "Restoring a virtual machine to a Nutanix cluster or a vSphere environment" on page 189.

The activities of the restored HYCU backup controller are suspended automatically.

7. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
8. Sign in to the HYCU web user interface of the restored HYCU backup controller.
9. Resume the activities of the restored HYCU backup controller. For instructions, see "Setting power options" on page 509.
10. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see your platform documentation.
11. *For Nutanix ESXi clusters:* If the original HYCU backup controller does not exist, configure settings for the new network adapter that was assigned to the HYCU backup controller. For instructions, see "Configuring your network" on page 496.

    > **Important** Make sure to enter the original IP address of the HYCU backup controller. After editing the connection, delete the old network adapter.

Restoring the HYCU backup controller to a XenServer environment

Use this procedure when you want to restore the HYCU backup controller to the original XenServer environment.

Prerequisite

The recovery HYCU backup controller must have network access to the source of the original HYCU backup controller.

Procedure

1. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
2. In the Virtual Machines panel, select the HYCU backup controller.
3. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

> open the Detail view.

4. Click **Restore**.
5. Select **Virtual machine options**, and then click **Next**.
6. Select **Restore VM**, and then restore the HYCU backup controller by following the instructions described in "Restoring a virtual machine to a XenServer environment" on page 193.

   The activities of the restored HYCU backup controller are suspended automatically.
7. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
8. Sign in to the HYCU web user interface of the restored HYCU backup controller.
9. Resume the activities of the restored HYCU backup controller. For instructions, see "Setting power options" on page 509.
10. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see your platform documentation.

Restoring the HYCU backup controller to an AWS GovCloud (US) environment

Use this procedure when you want to restore the HYCU backup controller to the original AWS GovCloud (US) region.

Procedure

1. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
2. In the Virtual Machines panel, select the HYCU backup controller.
3. In the Detail view that appears at the bottom of the screen, select the latest restore point.

   > **Note** The Detail view appears only if you click a virtual machine. Selecting the check box before the name of the virtual machine will not open the Detail view.
4. Click **Restore**.
5. Select **Virtual machine options**, and then click **Next**.

6. Select **Restore VM**, and then restore the HYCU backup controller by following the instructions described in "Restoring a virtual machine to an AWS GovCloud (US) environment" on page 196.

   The activities of the restored HYCU backup controller are suspended automatically.
7. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
8. Sign in to the HYCU web user interface of the restored HYCU backup controller.
9. Resume the activities of the restored HYCU backup controller. For instructions, see "Setting power options" on page 509.
10. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see your platform documentation.

Restoring the HYCU backup controller to an Azure or Azure Government environment

Use this procedure when you want to restore the HYCU backup controller to the original Azure or Azure Government subscription.

Procedure

1. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
2. In the Virtual Machines panel, select the HYCU backup controller.
3. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

4. Click **Restore**.
5. Select **Virtual machine options**, and then click **Next**.
6. Select **Restore VM**, and then restore the HYCU backup controller by following the instructions described in "Restoring a virtual machine to an Azure or Azure Government environment" on page 200.

   The activities of the restored HYCU backup controller are suspended automatically.

7. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
8. Sign in to the HYCU web user interface of the restored HYCU backup controller.
9. Resume the activities of the restored HYCU backup controller. For instructions, see “Setting power options” on page 509.
10. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see your platform documentation.

Restoring the HYCU backup controller to a different source

Depending on the source to which you want to restore the HYCU backup controller, see one the following sections:

- “Restoring the HYCU backup controller to a Nutanix cluster, a vSphere environment, an Azure Local environment, or a Hyper-V cluster” below
- “Restoring the HYCU backup controller to a XenServer environment” on page 376
- “Restoring the HYCU backup controller to an AWS GovCloud (US) environment” on page 378
- “Restoring the HYCU backup controller to an Azure or Azure Government environment” on page 379

Restoring the HYCU backup controller to a Nutanix cluster, a vSphere environment, an Azure Local environment, or a Hyper-V cluster

Use this procedure when the source on which the original HYCU backup controller was running is damaged or inoperable, or if you want to relocate the HYCU backup controller.

Prerequisites

- The recovery HYCU backup controller must have network access to the source to which you plan to restore the original HYCU backup controller.
- *Only if the backup of the original HYCU backup controller is stored on an iSCSI or a Nutanix target:* The target must be deactivated and detached from the

recovery HYCU backup controller before you power on the restored HYCU backup controller.

Procedure

1. *Only if the original HYCU backup controller still exists.* Suspend the activities of the original HYCU backup controller.

   > ⚠ Caution Make sure that a clone of the HYCU backup controller is not activated while the original HYCU backup controller is still active. Skipping this step may result in data loss.

   To suspend the activities of the original HYCU backup controller, follow these steps:

   a. *Only if the HYCU backup controller is turned off.* Turn the HYCU backup controller (virtual machine) on.

   b. Sign in to the HYCU web user interface.

   c. Suspend the activities of the HYCU backup controller. For instructions, see "Setting power options" on page 509.

   d. Wait for the running jobs to complete. You can check this by filtering the Jobs list by the Executing job status. For instructions, see "Filtering and sorting data" on page 406.

2. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
3. In the Virtual Machines panel, select the original HYCU backup controller.
4. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

5. Click **Restore**.
6. Select **Virtual machine options**, and then click **Next**.
7. Select **Clone VM**, and then restore the HYCU backup controller by following the instructions described in "Cloning a virtual machine to a Nutanix cluster, a vSphere environment, an Azure Local environment, or a Hyper-V cluster" on page 205.

   The activities of the restored HYCU backup controller are suspended automatically.

8. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
9. Sign in to the HYCU web user interface of the restored HYCU backup controller.
10. Resume the activities of the restored HYCU backup controller. For instructions, see "Setting power options" on page 509.
11. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see your platform documentation.
12. *Only if you want to use network settings of the original HYCU backup controller.* Configure settings for the network adapter of the HYCU backup controller. For instructions, see "Configuring your network" on page 496.

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

Restoring the HYCU backup controller to a XenServer environment

Use this procedure when the source on which the original HYCU backup controller was running is damaged or inoperable, or if you want to relocate the HYCU backup controller.

Prerequisite

The recovery HYCU backup controller must have network access to the source to which you plan to restore the original HYCU backup controller.

Procedure

1. *Only if the original HYCU backup controller still exists.* Suspend the activities of the original HYCU backup controller.

   > **Caution** Make sure that a clone of the HYCU backup controller is not activated while the original HYCU backup controller is still active. Skipping this step may result in data loss.

   To suspend the activities of the original HYCU backup controller, follow these steps:

   a. *Only if the HYCU backup controller is turned off.* Turn the HYCU backup controller virtual machine on.
   b. Sign in to the HYCU web user interface.

c. Suspend the activities of the HYCU backup controller. For instructions, see "Setting power options" on page 509.

d. Wait for the running jobs to complete. You can check this by filtering the Jobs list by the Executing job status. For instructions, see "Filtering and sorting data" on page 406.

2. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
3. In the Virtual Machines panel, select the original HYCU backup controller.
4. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

5. Click **Restore**.
6. Select **Virtual machine options**, and then click **Next**.
7. Select **Clone VM**, and then restore the HYCU backup controller by following the instructions described in "Cloning a virtual machine to a XenServer environment" on page 211.

   The activities of the restored HYCU backup controller are suspended automatically.
8. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
9. Sign in to the HYCU web user interface of the restored HYCU backup controller.
10. Resume the activities of the restored HYCU backup controller. For instructions, see "Setting power options" on page 509.
11. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see your platform documentation.
12. *Only if you want to use network settings of the original HYCU backup controller.* Configure settings for the network adapter of the HYCU backup controller. For instructions, see "Configuring your network" on page 496.

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

Restoring the HYCU backup controller to an AWS GovCloud (US) environment

Use this procedure if you want to relocate the HYCU backup controller.

Procedure

1. *Only if the original HYCU backup controller still exists.* Suspend the activities of the original HYCU backup controller.

   > ⚠ Caution Make sure that a clone of the HYCU backup controller is not activated while the original HYCU backup controller is still active. Skipping this step may result in data loss.

   To suspend the activities of the original HYCU backup controller, follow these steps:

   a. *Only if the HYCU backup controller is turned off.* Turn the HYCU backup controller virtual machine on.

   b. Sign in to the HYCU web user interface.

   c. Suspend the activities of the HYCU backup controller. For instructions, see “Setting power options” on page 509.

   d. Wait for the running jobs to complete. You can check this by filtering the Jobs list by the Executing job status. For instructions, see “Filtering and sorting data” on page 406.

2. Sign in to the HYCU web user interface of the recovery HYCU backup controller.

3. In the Virtual Machines panel, select the original HYCU backup controller.

4. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

5. Click **Restore**.

6. Select **Virtual machine options**, and then click **Next**.

7. Select **Clone VM**, and then restore the HYCU backup controller by following the instructions described in “Cloning a virtual machine to an AWS GovCloud (US) environment” on page 215.

   The activities of the restored HYCU backup controller are suspended automatically.

8. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
9. Sign in to the HYCU web user interface of the restored HYCU backup controller.
10. Resume the activities of the restored HYCU backup controller. For instructions, see “Setting power options” on page 509.
11. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see your platform documentation.

Restoring the HYCU backup controller to an Azure or Azure Government environment

Use this procedure if you want to relocate the HYCU backup controller.

Procedure

1. *Only if the original HYCU backup controller still exists.* Suspend the activities of the original HYCU backup controller.

   > ⚠ Caution Make sure that a clone of the HYCU backup controller is not activated while the original HYCU backup controller is still active. Skipping this step may result in data loss.

   To suspend the activities of the original HYCU backup controller, follow these steps:

   a. *Only if the HYCU backup controller is turned off.* Turn the HYCU backup controller (virtual machine) on.
   b. Sign in to the HYCU web user interface.
   c. Suspend the activities of the HYCU backup controller. For instructions, see “Setting power options” on page 509.
   d. Wait for the running jobs to complete. You can check this by filtering the Jobs list by the Executing job status. For instructions, see “Filtering and sorting data” on page 406.
2. Sign in to the HYCU web user interface of the recovery HYCU backup controller.
3. In the Virtual Machines panel, select the original HYCU backup controller.
4. In the Detail view that appears at the bottom of the screen, select the latest restore point.

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

5. Click **Restore**.
6. Select **Virtual machine options**, and then click **Next**.
7. Select **Clone VM**, and then restore the HYCU backup controller by following the instructions described in "Cloning a virtual machine" on page 205.

   The activities of the restored HYCU backup controller are suspended automatically.
8. Sign out of the HYCU web user interface of the recovery HYCU backup controller.
9. Sign in to the HYCU web user interface of the restored HYCU backup controller.
10. Resume the activities of the restored HYCU backup controller. For instructions, see "Setting power options" on page 509.
11. *Only if you decide not to keep the recovery HYCU backup controller.* Delete the recovery HYCU backup controller from its source. For instructions, see your platform documentation.

Chapter 10

Performing daily tasks

To ensure the secure and reliable performance of the data protection environment, HYCU provides various mechanisms to support your daily activities.

I want to...	Procedure
Get an at-a-glance overview of the data protection environment state, identify eventual bottlenecks, and inspect different areas of the data protection environment.	“Using the HYCU dashboard” on the next page
Track jobs that are running in my environment, get an insight into a specific job status, generate a job report, and cancel a currently running job.	“Managing HYCU jobs” on page 385
View all events that occurred in my environment.	“Managing HYCU events” on page 387
Configure HYCU to send notifications when events occur.	“Configuring event notifications” on page 388
Enable purging of events and jobs.	“Enabling the purge of events and jobs” on page 391
Obtain reports on different aspects of the data protection environment.	“Using HYCU reports” on page 392
View entity details.	“Viewing entity details” on page 399
Narrow down the list of displayed items by applying filters.	“Filtering and sorting data” on page 406
Export data that you can view in a table in any of the panels to a JSON or CSV file.	“Exporting the contents of the panel” on page 416

I want to...	Procedure
View target information, activate or deactivate a target, increase the size of an iSCSI target, or edit or remove a target.	“Managing targets” on page 417
View policy information, or edit or delete a policy.	“Managing policies” on page 422
Back up data manually.	“Performing a manual backup” on page 424
Set up a validation policy and schedule the backup validation.	“Setting up a validation policy” on page 425
Running malware scans manually	“Running malware scans manually” on page 431
Manually override the R-Shield status.	“Overriding the R-Shield status” on page 433
Archive data manually.	“Archiving data manually” on page 435
Recreate a snapshot.	“Recreating snapshots” on page 436

In case of recognized problems in your data protection environment that can degrade the efficiency and reliability of data protection (for example, when storage, vCPU, or memory utilization is exceeded), you can make adjustments to better meet your data protection goals. For details, see “Adjusting the HYCU backup controller resources” on page 438.

You can set up the appearance of the HYCU web user interface. For details, see “Setting up the appearance of your HYCU web user interface” on page 438.

Using the HYCU dashboard

The HYCU dashboard provides you with an at-a-glance overview of the data protection status in your environment. This intuitive dashboard enables you to monitor all data protection activity and to quickly identify areas that need your attention. You can use this dashboard as a starting point for your everyday

tasks because it enables you to easily access the area of interest by simply clicking the corresponding widget.

Accessing the Dashboard panel

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

Important Your user role defines which widgets you are allowed to see and access.

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

Dashboard widget	Description
Virtual Machines	Shows the percentage of compliant and non-compliant virtual machines, the percentage of protected and unprotected virtual machines in the data protection environment, the percentage of R-Shield compliant and non-compliant virtual machines, and the percentage of protected virtual machines that have the DR-ready status. A virtual machine is considered: • Compliant: If the time since the last successful backup is lower than its RPO and the estimated time to recover is lower than its RTO. • Protected: If it has at least one valid backup available and does not have the Exclude policy assigned. • DR-ready: If all backups in the current backup chain are stored on one of the cloud targets. For detailed information on protecting virtual machines, see “Backing up virtual machines” on page 182.
Applications	Shows the percentage of compliant and non-compliant applications, the percentage of protected and unprotected applications in the data protection environment, and the percentage of R-Shield compliant and non-compliant applications. An application is considered: • Compliant: If the time since the last successful backup is lower than its RPO and the estimated time to recover is lower than its RTO.

Dashboard widget	Description
	• Protected: If it has at least one valid backup available and does not have the Exclude policy assigned. For detailed information about protecting applications, see "Backing up applications" on page 270.
Shares	Shows the percentage of compliant and non-compliant file shares and the percentage of protected and unprotected file shares in the data protection environment. A file share is considered: • Compliant: If the time since the last successful backup is lower than its RPO and the estimated time to recover is lower than its RTO. • Protected: If it has at least one valid backup available and does not have the Exclude policy assigned. For detailed information about protecting file shares, see "Protecting file shares" on page 327.
R-Shield	Shows the following: • The number of applications, virtual machines, volume groups, and shares with OK or Suspicious R-Shield statuses. • The number of applications, virtual machines, volume groups, and shares that do not have the R-Shield policy enabled in their assigned policy. • The percentage of entities that were scanned for malware. • The number of entities with Suspicious R-Shield status. For detailed information about R-Shield statuses, see "Viewing the R-Shield status" on page 403.
Targets[a]	Shows the list of all targets in the data protection environment, and the information on how much space is used and available for storing data on each target and on all targets in the data protection environment combined. For detailed information about setting up targets, see "Setting up targets" on page 100.

Dashboard widget	Description
Policies	Shows the number of entities in the data protection environment, the number of entities that have no policy assigned, and the number of entities that are compliant and non-compliant with the RPO and RTO set in their assigned policy. The number of compliant and non-compliant entities for specific policies is also shown. For detailed information about policies, see "Defining your backup strategy" on page 130.
Backups	Shows the number of backups that were performed today and the backup success rate for the last 7 days.
Jobs	Shows the number of jobs in the data protection environment in the last 56 hours according to their status (Success, Warning, Failed, In progress, and Queued). For details on jobs, see "Managing HYCU jobs" below.

[a] Available only for an infrastructure group administrator.

Managing HYCU jobs

In the Jobs panel, you can do the following:

- Check the processes that are currently running.
- Check the completed and stopped processes.
- Check more details about a specific job in the Detail view that appears at the bottom of the screen after you select the job.

  > **Tip** By pausing on the progress bar of a particular task (for example, Backup data), additional information about the task is available, such as how much data has already been backed up and when the progress time has been last updated.

- *For virtual machines with attached volume groups:* Check the backup and restore process statuses of the volume groups attached to the virtual machines. To do so, click the arrow next to the backup or restore job of a virtual machine with attached volume groups, and a list of attached volume group processes and their statuses will be expanded. Keep in mind that

volume group processes will not appear all at once, but one after another, as the job progresses.

- Generate a report about a specific job by selecting it, and then clicking **View Report**. To copy the report to the clipboard, in the Job Report dialog box that opens, click **Copy to Clipboard**.
- Cancel a currently running or queued job by selecting it, and then clicking **Abort Job**.
- Enable purging of jobs. For details, see "Enabling the purge of events and jobs" on page 391.

Considerations

If a backup, backup copy, or archive job fails, HYCU automatically schedules job retries. Consider the following:

- If the backup job fails, the time interval between two successive retries is doubled with each retry until the RPO value is reached (for example, by default, the first retry occurs after 15 minutes, the second one after 30 minutes, the third one after 1 hour, and so on). When the RPO value is reached, the time interval for retrying the backup job becomes the same as the one specified for the RPO.
- If the backup copy job fails, HYCU retries the failed job two times with the time interval of 15 minutes (by default). If these retries fail, the retry job is suspended for 24 hours.
- If the archive job fails, HYCU retries the failed job once after 15 minutes (by default). If this retry fails, the retry job is suspended for 12 hours.

Accessing the Jobs panel

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

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

The following information is available for each job:

Job information	Description
Name	Name of a job that was performed (for example, adding a source, adding a target, running a backup, and so on).
Status	Current status of a job (for example, Queued, a progress bar indicating the Executing status, OK, or Error).

Job information	Description
Created	When a job was created.
Finished	When a job finished.

Managing HYCU events

In the Events panel, you can do the following:

- View all events that occurred in your environment.
- Check details about the selected event.
- List events that match the specified filter.
- Configure HYCU to send notifications when the events occur. For details, see "Configuring event notifications" on the next page.
- Enable purging of events. For details, see "Enabling the purge of events and jobs" on page 391.

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

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

The following information is available for each event:

Event information	Description
Status	Status of the event (Success, Warning, Failed)
Message	Description of the event
Category	Category to which the event belongs (for example, Policies, Backup, Credentials, System for an internal event, and so on)
Timestamp	Event creation time

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

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

Configuring event notifications

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

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

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

- "Setting up email notifications" below
- "Setting up webhook notifications" on the next page

Setting up email notifications

Prerequisite

Email delivery must be configured. For details, see "Configuring email delivery" on page 512.

Procedure

1. In the Notifications dialog box, click the **Email** tab, and then click **New**.
2. In the Subject field, enter a subject for the email notification.
3. From the Category drop-down menu, select one or more categories to which the events belong (for example, Policies, Backup, Credentials, System, and so on). To include all categories, click **Select all**.
4. From the Status drop-down menu, select the status of the events (Success, Warning, Failed). To include all statuses, click **Select all**.
5. From the Language drop-down menu, select the preferred language for email notifications.
6. In the Email address field, enter one or more email addresses to which you want the notifications to be sent. If you are entering more than one email address, make sure to press the spacebar after entering each one.
7. Click **Save**.

Your changes take effect immediately and email notifications are sent to any email address that you specified in the notification settings.

You can later edit settings for existing email notifications (click **Edit** and make the required modifications) or delete the ones that you do not need anymore (click **Delete**).

Setting up webhook notifications

Procedure

1. In the Notifications dialog box, click the **Webhooks** tab, and then click **New**.
2. Enter a name for the webhook notification and, optionally, its description.
3. From the Category drop-down menu, select one or more categories to which the events belong (for example, Policies, Backup, Credentials, System, and so on). To include all categories, click **Select all**.
4. From the Status drop-down menu, select the status of the events (Success, Warning, Failed). To include all statuses, click **Select all**.
5. From the Language drop-down menu, select the preferred language for webhook notifications.
6. In the Post URL field, enter the URL of the endpoint the webhook notifications should be sent to in one of the following formats:

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

   For details on the format of the data that HYCU sends to the specified URL, see “Webhook data format” on the next page.
7. *Only if the receiving endpoint requires sender's identification.* From the Authentication Type drop-down menu, select one of the following authentication types:
   - **Basic authentication**, and then enter the user name and password associated with your webhook endpoint.
   - **Authentication header**, and then do the following:
     a. Enter a custom header key and value that will be used to verify each webhook request.
     b. Click **Add** for the header to be added to the list of headers.

If required, repeat these steps for each additional header that you want to add.

- **Authentication by secret**, and then enter the secret to connect to your webhook endpoint.

8. Click **Next**.
9. *Optional.* Customize the request body that is sent by HYCU. You can click the appropriate fields in the HYCU fields list to easily insert event variables into the body.

   > **Important** Make sure the format you define in the body is supported by the platform to which webhook notifications will be sent.

   For details on the format of the webhook request body, see "Webhook data format" below.
10. Click **Save**.

Your changes take effect immediately and webhook notifications are sent to the URL that you specified in the notification settings.

You can later edit settings for existing webhook notifications (click **Edit** and make the required modifications) or delete the ones that you do not need anymore (click **Delete**).

Webhook data format

The webhook data format is defined by:

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

HTTP request headers

The request headers are sent in the following format:

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

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

HTTP request body

The request body is sent in the following format:

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

Note Null values are ignored.

HTTP response code

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

Enabling the purge of events and jobs

You can configure HYCU to periodically delete events and/or jobs (as well as all associated job reports) that are no longer needed for daily business operations by enabling the purge of data from the HYCU database.

Prerequisite

You must be an infrastructure group administrator.

Consideration

Jobs related to backups, copies of backups, and archives will be deleted only if the corresponding restore points no longer exist or are expired.

Depending on whether you want to purge events or jobs, access one of the following panels:

- #### Accessing the Events panel

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

- Accessing the Jobs panel

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

Procedure

1. In the Events or Jobs panel, click **Purge Configuration**.
2. Depending on your context, use the **Enable purging of events** or **Enable purging of jobs** switch.
3. Specify the number of years, months, weeks, or days to retain the data. Events or jobs that are older than the specified value will be purged. The maximum value is 99 years.
4. Click **Save** to start purging the HYCU database based on the specified value.

   > Important This action cannot be undone. When your event or job data is deleted, you cannot retrieve it.

After you enable purging of events and/or jobs, you can at any later time edit the purge configuration or disable purging.

Using HYCU reports

HYCU reports provide you with a visual presentation of data protection environment resources and jobs. This comprehensive and precise presentation allows you to have an optimum view for analyzing data and therefore making the best decisions when it comes to protecting your data.

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

Consideration

Keep in mind that your user group and user role determine what kind of report data you can view and what report actions you can perform.

After you get familiar with the reports as described in “Getting started with reporting” on the next page, you can continue as follows:

- View reports. For details, see “Viewing reports” on page 395.
- Generate reports. For details, see “Generating reports” on page 396.
- Schedule reports. For details, see “Scheduling reports” on page 397.

> **Note** When scheduling the reports, you can also choose to send them by email.

- Export and import reports. For details, see "Exporting and importing reports" on page 398.

Accessing the Reports panel

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

Getting started with reporting

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

For a list of predefined reports, see "Predefined reports" below. For instructions on how to create reports, see "Creating reports" on the next page.

Predefined reports

The predefined reports represented by the icon enable you to obtain reports on the key aspects of your data protection environment such as data transfer, job status, the number of backups, and the amount of protected data. These reports cannot be edited or deleted.

Predefined report	Description
Backup validation	List of virtual machine backup validations that occurred in the last 24 hours including information such as the status of the backup validation and the reason for the validation failure.
Entity compliance status	List of virtual machines, applications, shares, buckets, and servers that are compliant and non-compliant with backup requirements.
Hourly activities per policy	List of assigned policies with the corresponding number of jobs that were running during each of the last 24 hours.
Hourly activities per target[a]	List of targets with the corresponding number of jobs that were running during each of the last 24 hours.
Protected data	Total amount of protected data calculated on a daily basis.

Predefined report	Description
Protected data per policy	Amount of data protected in the last 24 hours per policy.
Protected data per owner[a]	Total amount of protected data per owner.
Protected data per target[a]	Amount of the data protected in the last 24 hours per target.
Protected data timeline per target[a]	Daily amount of protected data per target.
Protected VM size per target[a]	List of protected virtual machine and servers, and distribution of the corresponding protected data between targets.
R-Shield status	List of backups with the overall R-Shield, anomaly, and malware detection statuses, and manual R-Shield status overrides.
VM backup status	List of backups that occurred in the last 24 hours including information such as status and duration of backups, backup size, and so on.
VM backup status per target[a]	List of targets and related backups that occurred in the last 24 hours including information such as status and duration of backups, backup size, and so on.

[a] Available only to an infrastructure group administrator.

Creating reports

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

Prerequisite

You must have the Administrator user role assigned.

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

I want to...	Procedure
Create a new report from scratch.	1. In the Reports panel, click **New**. The New Report dialog box opens. 2. Enter a report name and, optionally, its description. 3. Select the type of report. 4. Select the aggregation calculation to be used for your report. 5. Select the report tag for which the aggregation calculation should be performed. 6. Specify the time range for the report. You can select one of the predefined time ranges, or select **Custom**, and then use the calendar to select a start date and an end date of the time range. 7. Distribute the report tags for the collected data that you want to include in your report between x-axis and y-axis to determine how the collected data will be presented in the report. 8. Click **Save**.
Edit an existing report and save it as a new report.	1. In the Reports panel, from the list of reports, select the one that you want to edit and save as a new report, and then click **Edit**. The Report Configuration dialog box opens. 2. Enter a new name for the report, and then make the required modifications. 3. Click **Save As**.

You can later edit any of the created reports (click **Edit** and make the required modifications) or delete the ones that you do not need anymore (click **Delete**). You cannot edit or delete the predefined reports represented by the icon.

Viewing reports

You can view the reports on the current state of your data protection environment or the saved reports that were generated either manually or automatically.

I want to...	Procedure
View a report on the current state of my data protection environment.	In the Reports panel, from the list of reports, select the preferred report, and then click **Preview**.
View a saved report.	1. In the Reports panel, from the list of reports, select the preferred report. 2. In the Detail view that appears at the bottom of the screen, select the preferred report version, and then click **View**. For details on how to generate reports manually or automatically, see "Generating reports" below or "Scheduling reports" on the next page.

In the dialog box that opens, besides viewing the report data, you can also do the following:

- Switch between the reports.
- Download the report in the PDF, PNG, or CSV format. To do so, click **Download**, and then select one of the available formats.
- *For users with the Administrator user role assigned:* If you view a report on the current state of the data protection environment, you can save this version of the report by clicking **Generate**. The saved report is added to the list of report versions.

Generating reports

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

Prerequisite

You must have the Administrator user role assigned.

Procedure

1. In the Reports panel, from the list of reports, select the one that you want to generate.

> **Note** If none of the available reports meets your reporting requirements, you can create a new report. For details, see "Creating reports" on page 394.

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

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

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

You can later do the following:

- View the saved reports. For details, see "Viewing reports" on page 395.
- Delete the saved reports that you do not need anymore. To do so, select the preferred report version, and then click **Delete**.

Scheduling reports

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

Prerequisites

- You must have the Administrator user role assigned.
- *For sending reports by email:* An SMTP server must be configured. For details, see "Configuring email delivery" on page 512.

Procedure

1. In the Reports panel, from the list of reports, select the one that you want to be generated on a regular basis, and then click **Scheduler**. The Report Scheduler dialog box opens.

   > **Note** If none of the available reports meets your reporting requirements, you can create a new report. For details, see "Creating reports" on page 394.

2. In the Schedule date field, specify the date and the time of day when you want the report generation to begin.
3. From the Interval drop-down menu, select how often you want the reports to be generated (daily, weekly, or monthly).
4. Use the **Send report to email** switch if you want to schedule the automatic delivery of the reports to email recipients, and then do the following:
   a. From the Report format drop-down menu, select a file format for your report (PDF, PNG, or CSV).
   b. In the Email address field, enter one or more email recipients that should receive the reports. If you are entering more than one email address, make sure to press the spacebar after entering each one.
   c. In the Email text field, enter the custom information to be included in the report email. For example, you can specify the purpose of the report, list further actions to be done by the recipient, or add contact information for additional report details.
5. Click **Save**.

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

You can later do the following:

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

Exporting and importing reports

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

> Important Your permissions determine what kind of reports you can view and edit, and therefore also define a different level of access to the reports, which you should consider before copying reports from one HYCU deployment to another.

Exporting reports

Procedure

In the Reports panel, from the list of all reports, select the one that you want to export, and then click **Export**. The selected report will be exported to a JSON file and saved to the download location on your system.

Importing reports

Procedure

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

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

Viewing entity details

You can view the details about each virtual machine, discovered application, file share, server, volume group, and bucket in the Detail view of the Virtual Machines, Applications, Shares, Volume Groups, or Buckets panel. The following details are available:

Summary	Shows detailed information about the selected entity (for example, the source UUID, the time since the last successful backup, and the R-Shield status). **Note** The R-Shield status is shown only if you enabled the R-Shield policy option in the assigned policy.
Restore point	You can view the following information about each restore point: • Date and time when the restore point was created. • Tiers: ◦ BCKP Backup: Available by default unless a backup is expired. ▪ FULL Full: Visible if a full backup was performed.

	▪ INCR Incremental: Visible if an incremental backup was performed. ◦ Archive: ▪ D ARCH Daily archive: Available if a daily data archive was created. ▪ W ARCH Weekly archive: Available if a weekly data archive was created. ▪ M ARCH Monthly archive: Available if a monthly data archive was created. ▪ Y ARCH Yearly archive: Available if a yearly data archive was created. By pausing on the archive icon, you can see the total number of data archives and the archive expiration time. If any of the archive jobs failed, the number of failed archive jobs is shown. ◦ COPY Copy: Available if a copy of backup data was created. By pausing on the icon, you can see the total number of backup copies and the backup copy expiration time. If any of the backup copy jobs failed, the number of failed backup copy jobs is shown. ◦ SNAP Snapshot: Available if the source contains a local snapshot that enables you to perform a fast restore. By pausing on the icon, you can see whether the snapshot was recreated and its expiration time. If any virtual disks were excluded from a backup, the corresponding tier label is marked with a line in the upper left corner. For example, FULL. **Important** If any of the tiers is colored red, it cannot be used for a restore.
Compliance	Shows the compliance status of an entity: • Success • Failure • Undefined An entity is considered to be compliant with backup

	requirements if the time since the last successful backup is lower than the RPO set in the HYCU policy and the estimated time to recover is lower than the RTO set in the HYCU policy. By pausing on the compliance status indicated by a respective icon, additional information about the backup is available. You can see backup frequency, the elapsed time since the last successful backup, the time limit you set for the restore, and the estimated time required for the restore. In addition, if the compliance status of your entity is Failure, this list will also include a reason why it is not compliant.
Backup status	For details, see "Viewing the backup status of entities" below.
R-Shield status	For details, see "Viewing the R-Shield status" on page 403.
Restore status	Shows a progress bar indicating the progress of the entity restore. **Tip** If you double-click a progress bar, you are directed to the Jobs panel where you can check details about the related job.
Backup data size change rate	Shows changes in full and incremental backup data sizes over time. For details, see "Viewing the backup data size change rate" on page 404.
Tier status	Shows backup statuses of individual tiers. For details, see "Viewing tier statuses" on page 405.

Tip If there are too many items to be displayed on one page, you can move between the pages by clicking > and <. You can also use ⌄ to set the number of items to be displayed per page.

Viewing the backup status of entities

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

Limitation

For virtual machines with attached volume groups: The Completed with errors backup status is available only for virtual machines that have volume groups attached directly.

Backup status of the entity	Restore a VM, a VG, or vDisks?	Restore VM files?	Restore an application?	Restore a file share or a bucket?
Completed successfully	✓	✓	✓	✓
Completed with warnings	✓	✓	✓[a]	✓
Completed with errors	✓[b]	✓[c]	✓[d]	✓[e]
Failed	✗	✗	✗	✗
Expired	✗	✗	✗	✗
Skipped[f]	✓	✓	✗	N/A

[a] You cannot specify a point in time to which you want to restore data. This backup status may occur because disk mapping failed or a virtual machine does not have an NIC, or, in case of applications, at least one database log backup failed (whereas all other databases are in a consistent state).
[b] Because not all virtual machine disk files were backed up successfully, the virtual machine can be partially restored. It may not be possible to turn it on if one of the system disks was not backed up.
[c] Because not all virtual machine disk files were backed up successfully, the individual files can be partially restored (only the files that are displayed in the Restore Files dialog box).
[d] An application can be partially restored (only the databases that are displayed in the respective restore dialog boxes).
[e] Because not all files or objects were backed up successfully, the file share or bucket can be partially restored. The files or objects whose backup was unsuccessful are listed in the Job Report in their corresponding subtasks.
[f] Applicable only for backups of passive nodes of failover clusters with shared storage.

> **Note** By pausing on the backup status indicated by an icon, additional information about the backup is available. You can see the backup type, backup consistency, the duration and size of the backup, which target was used, and the backup UUID. For volume groups, you can also see if the

> volume group has been backed up both as part of the virtual machine backup and by assigning a policy directly to it.
> If you double-click a backup status icon, you are directed to the Jobs panel where you can check details about the related jobs.

Viewing the R-Shield status

The R-Shield status shows if there are changes in the data size between backups and if the backup data contains suspicious items that are detected by the malware scans. The R-Shield status is available for the entities that have the R-Shield option enabled in the assigned policy. For details, see "Creating an R-Shield policy" on page 148.

The R-Shield status is displayed for entities and their restore points. For the entities, the R-Shield status is determined by the R-Shield status of the last restore point with a successfully completed backup. For the restore points, the R-Shield status is determined by taking the least favorable of the anomaly detection and malware detection statuses. For example, if the anomaly detection status is OK, and the malware detection status is Suspicious, the R-Shield status will be Suspicious.

If after reviewing the anomaly or malware detection statuses of any protected entity or restore point, you determine that it is not a malicious anomaly or a suspicious item, you can manually override these statuses. For details, see "Overriding the R-Shield status" on page 433. For the manually overridden R-Shield statuses, the icons are marked with the following symbol: Ⓜ.

Considerations

- The malware detection status of restore points is determined based only on the valid restore points.
- If any of the files in the backup cannot be scanned and no suspicious items are found in the remaining scanned files, the R-Shield status is OK. The Warning status of the Scan VM Image job in the detail view of the relevant malware scan job in the Jobs panel indicates that some files could not be scanned.
- *Only if you enabled anomaly detection in the R-Shield policy.* At least five consecutive backups of the same type (full or incremental) must be available in the detection range before the anomaly detection status is indicated. Before that, the R-Shield status is Pending.

R-Shield status	Description
OK	OK status is displayed if all of the following conditions apply: • No detected anomaly in the backup data size and no suspicious items were found for the selected restore point. • No suspicious anomaly or malware detection statuses persist in the following interval: ◦ *For anomaly detection:* Within the configured detection range. ◦ *For malware scanning:* Since the last scan of all files. OK status is displayed also if the anomaly and malware detection statuses were manually overridden to OK.
Suspicious	Detected anomaly in the backup data size or suspicious items found for the selected restore point.
Warning	No detected anomaly in the backup data size and no suspicious items were found for the selected restore point. However, suspicious anomaly or malware detection statuses persist in the following interval: • *For anomaly detection:* Within the configured detection range. • *For malware scanning:* Since the last scan of all files.
Undefined	The entity does not have a policy with the R-Shield option assigned.

By hovering over the R-Shield status in the Detail view, you can see the anomaly and malware detection statuses and their details.

Viewing the backup data size change rate

A chronological representation of the full and incremental backup data size changes is available in the Change Rate graph, which you can access by clicking **Average change rate** in the Summary.

If the R-Shield option is enabled in the policy that is assigned to the entity, in addition to the data size changes, the following information is also displayed in the graph:

- The incremental anomaly detection threshold. For details on how to set the threshold for the change in the standard deviation of data size between the backups, see “Creating an R-Shield policy” on page 148.
- Standard deviation of the incremental backup data size.
- *For the entities with Suspicious or Warning R-Shield status:* By hovering over a restore point, you can see the number of detected malware threats for that restore point.

Limitation

If Snapshot is defined as the backup target type in the policy that is assigned to the selected entity, the Change Rate graph is not available.

Viewing tier statuses

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

Tier status	Restore an entity?
BCKP or B (Done)	✓
BCKP or B (Done with warnings)	✓ For details on what data can be restored if one of these backup statuses is shown, see “Viewing the backup status of entities” on page 401.
BCKP or B (Done with errors)	
BCKP or B (Inaccessible on source)	✗
BCKP or B (Deleted from the source)	✗
BCKP or B (Failed)	✗
BCKP or B (Expired)	✗

Filtering and sorting data

HYCU enables you to filter data in the panels so you can easily find what you need. You can apply two types of filters—the main view filter (to focus on certain aspects of your data protection environment) or the detail view filter (to focus on the information about the backup and restore data of the selected entity). After you apply any of the filters, only data that matches the filter criteria is displayed and you can easily find what you need.

In addition, to make it easier to work with the tables in the panels that have a large number of columns, you can also sort the data in ascending or descending order.

Depending on whether you want to filter or sort your data, see one of the following sections:

- "Filtering data in panels" below
- "Sorting data in panels" on page 415

Filtering data in panels

I want to apply...	Available in panels	Instructions
Main view filter	HYCU Controllers, Applications, Virtual Machines, Volume Groups, Shares, Buckets, Policies, Targets, Jobs, Events, and Self-Service	"Applying the main view filter" below
Detail view filter	Applications, Virtual Machines, Volume Groups, Shares, and Buckets	"Applying the detail view filter" on page 414

Applying the main view filter

Apply the main view filter when you want to focus on certain aspects of your data protection environment (for example, filtering data in the Virtual Machines panel helps you to focus only on the virtual machines that you are interested in or responsible for).

> **Note** You can filter the items also by using the Search field on the left side of the panel. Typing text in this field automatically filters and displays only the matching items.

Procedure

1. In the selected panel, click **Filters**.
2. In the side panel that opens, select your filter criteria.
3. Click **Apply Filters**.

For the details about the available main view filtering options, see one of the following sections:

- "Main view filtering options in the HYCU Controllers panel" below
- "Main view filtering options in the Applications panel" below
- "Main view filtering options in the Virtual Machines panel" on page 409
- "Main view filtering options in the Volume Groups panel" on page 410
- "Main view filtering options in the Shares panel" on page 411
- "Main view filtering options in the Buckets panel" on page 412
- "Main view filtering option in the Policies panel" on page 413
- "Main view filtering options in the Targets panel" on page 413
- "Main view filtering options in the Jobs panel" on page 413
- "Main view filtering options in the Events panel" on page 414
- "Main view filtering option in the Self-Service panel" on page 414

Main view filtering options in the HYCU Controllers panel

Filtering option	Filter HYCU controllers by...
Type	Type of the HYCU controller (Hybrid Cloud Edition, Azure, Microsoft 365/Google Workspace).
Status	Status of the HYCU controller (Active, Unavailable, Suspended, or Undefined).

Main view filtering options in the Applications panel

Filtering option	Filter applications by...
Sources	Sources that host the virtual machines on which the applications are running.
Policy Assignment	Policies assigned to the applications (Unassigned, Assigned, and/or Specific policies). Note If you filter applications by the Assigned option, the ones to which the Exclude policy is

Filtering option	Filter applications by...
	assigned will not be listed.
Owners	Owners that are assigned to the virtual machines on which the applications are running.
Application Types	Application types.
Compliance	Compliance statuses of the applications: • Success • Failure • Undefined: The Exclude policy is assigned to the applications or the applications do not have a policy assigned.
Protection	Protection statuses of the applications (Protected, Unprotected, and/or Protected deleted).
R-Shield Detection Status	Overall R-Shield status that is based on the anomaly and malware detection statuses of the applications (OK, Warning, Suspicious, Pending, Undefined).
Anomaly Detection Status	Anomaly detection status that is based on the backup data size comparison between multiple backups or across time intervals (OK, Warning, Suspicious, Pending, Undefined).
Malware Detection Status	Malware detection status that is based on the results of backup data scans for potential malware and ransomware (OK, Warning, Suspicious, Pending, Undefined).
Discovery	Discovery statuses of the applications: • Success • Failure • Warning: Application discovery failed because the virtual machine is offline or is not reachable.
Nutanix NDB support	NDB-managed databases.

Main view filtering options in the Virtual Machines panel

Filtering option	Filter virtual machines by...
Sources	Sources that host the virtual machines.
Credential Groups	Credential groups assigned to the virtual machines.
Policy Assignment	Policies assigned to the virtual machines (Unassigned, Assigned, and/or Specific policies). **Note** If you filter virtual machines by the Assigned option, the ones to which the Exclude policy is assigned will not be listed.
Validation Policy Assignment	Validation policies assigned to the virtual machines (Unassigned, Assigned, and/or Specific policies).
Owners	Owners assigned to the virtual machines.
Compliance	Compliance statuses of the virtual machines: • Success • Failure • Undefined: The Exclude policy is assigned to the virtual machines or the virtual machines do not have a policy assigned.
Discovery	Discovery statuses of the applications running on the virtual machines: • Success • Failure • Warning: Application discovery failed because the virtual is offline or is not reachable. • Undefined: The information about the application discovery status is not available.
Protection	Protection statuses of the virtual machines (Protected, Unprotected, and/or Protected deleted).
Validation Status	Backup validation statuses of the virtual machines.
R-Shield Detection status	Overall R-Shield status that is based on the anomaly and malware detection statuses of the virtual machines (OK, Warning, Suspicious, Pending, Undefined).

Filtering option	Filter virtual machines by...
Anomaly Detection Status	Anomaly detection status that is based on the virtual machine backup data size comparison between multiple backups or across time intervals (OK, Warning, Suspicious, Pending, Undefined).
Malware Detection Status	Malware detection status that is based on the results of virtual machine backup data scans for potential malware and ransomware (OK, Warning, Suspicious, Pending, Undefined).
Disaster Recovery Readiness	DR readiness statuses of the virtual machines.
Nutanix NDB Support	NDB-managed database server virtual machines.

Main view filtering options in the Volume Groups panel

Filtering option	Filter volume groups by...
Sources	Sources that host the volume groups.
Policy Assignment	Policies assigned to the volume groups (Unassigned, Assigned, and/or Specific policies). **Note** If you filter volume groups by the Assigned option, the ones to which the Exclude policy is assigned will not be listed.
Owners	Owners assigned to the volume groups.
Compliance	Compliance statuses of the volume groups: • Success • Failure • Undefined: The Exclude policy is assigned to the volume groups or the volume groups do not have a policy assigned.
Protection	Protection statuses of the volume groups (Protected, Unprotected, and/or Protected deleted).
R-Shield Detection Status	Overall R-Shield status that is based on the anomaly and malware detection statuses of the volume groups (OK,

Filtering option	Filter volume groups by...
	Warning, Suspicious, Pending, Undefined).
Anomaly Detection Status	Anomaly detection status that is based on the volume group backup data size comparison between multiple backups or across time intervals (OK, Warning, Suspicious, Pending, Undefined).
Malware Detection Status	Malware detection status that is based on the results of volume group backup data scans for potential malware and ransomware (OK, Warning, Suspicious, Pending, Undefined).

Main view filtering options in the Shares panel

Filtering option	Filter file shares by...
Sources	Sources that host the file shares.
Protocol	Protocols of the file shares (SMB or NFS).
Policy Assignment	Policies assigned to the file shares (Unassigned, Assigned, and/or Specific policies). **Note** If you filter file shares by the Assigned option, the ones to which the Exclude policy is assigned will not be listed.
Owners	Owners assigned to the file shares.
Compliance	Compliance statuses of the file shares: • Success • Failure • Undefined: The Exclude policy is assigned to the file shares or the file shares do not have a policy assigned.
Protection	Protection statuses of the file shares (Protected, Unprotected, and/or Protected deleted).
R-Shield Detection Status	Overall R-Shield status that is based on the anomaly and malware detection statuses of the file shares (OK, Warning, Suspicious, Pending, Undefined).
Anomaly Detection Status	Anomaly detection status that is based on the file share backup data size comparison between multiple backups

Filtering option	Filter file shares by...
	or across time intervals (OK, Warning, Suspicious, Pending, Undefined).
Malware Detection Status	Malware detection status that is based on the results of file share backup data scans for potential malware and ransomware (OK, Warning, Suspicious, Pending, Undefined).
Incremental Forever Backup	Enabled or disabled Incremental forever backup option.

Main view filtering options in the Buckets panel

Filtering option	Filter buckets by...
Sources	Sources that host the buckets.
Policy Assignment	Policies assigned to the buckets (Unassigned, Assigned, and/or Specific policies). **Note** If you filter buckets by the Assigned option, the ones to which the Exclude policy is assigned will not be listed.
Owners	Owners assigned to the buckets.
Compliance	Compliance statuses of the buckets: • Success • Failure • Undefined: The Exclude policy is assigned to the buckets or the buckets do not have a policy assigned.
Protection	Protection statuses of the buckets (Protected, Unprotected, and/or Protected deleted).
R-Shield Detection Status	Overall R-Shield status that is based on the anomaly and malware detection statuses of the buckets (OK, Warning, Suspicious, Pending, Undefined).
Anomaly Detection Status	Anomaly detection status that is based on the bucket backup data size comparison between multiple backups or across time intervals (OK, Warning, Suspicious, Pending, Undefined).

Filtering option	Filter buckets by...
Malware Detection Status	Malware detection status that is based on the results of bucket backup data scans for potential malware and ransomware (OK, Warning, Suspicious, Pending, Undefined).
Incremental Forever Backup	Enabled or disabled Incremental forever backup option.

Main view filtering option in the Policies panel

Filtering option	Filter policies by...
Compliance	Compliance statuses of the policies: • Success: All entities to which the policies are assigned are compliant with the policy settings. • Failure: Not all entities to which the policies are assigned are compliant with the policy settings. • Undefined: The Exclude policy is assigned to the entities or the entities do not have a policy assigned.

Main view filtering options in the Targets panel

Filtering option	Filter targets by...
Target Type	Target types.
Health	Health statuses of the targets (OK, Warning, Error, and/or Undefined).

Main view filtering options in the Jobs panel

Filtering option	Filter jobs by...
Status	Job statuses (OK, Warning, Queued, Executing, Aborted, and/or Failed).
Time Range	Time ranges: You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, and/or Last week), or use the calendar to select a start date and hour and an end date and hour of the time range for the jobs to be displayed.

Main view filtering options in the Events panel

Filtering option	Filter events by...
Category	Event categories.
Username	User names to include only the events started by the selected users and/or user groups.
Status	Event statuses (Success, Warning, and/or Failed).
Time Range	Time ranges: You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, and/or Last week), or use the calendar to select a start date and hour and an end date and hour of the time range for the events to be displayed.

Main view filtering option in the Self-Service panel

Filtering option	Filter users and/or user groups by...
Status	User or user group statuses—which users or user groups are allowed to sign in to HYCU and which are not (Active or Inactive).

Applying the detail view filter

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

Procedure

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

> Tip If there are too many filtered items to be displayed on one page, you can move between the pages by clicking > and <. You can also use ⌄ to set the number of filtered items to be displayed per page.

The following detail view filtering options are available:

Filtering option	Filter entities by...
Tiers	Restore point tiers of the entities (Archive daily, Archive weekly, Archive monthly, Archive yearly, Backup, Copy, and/or Snapshot).
Backup Type	Entity backup types (Incremental or Full).
Restore Point Date	Time when the entity restore points were created. You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, and/or Last week), or use the calendar to select a start date and hour and an end date and hour of the time range within which the restore points were created.
Backup Status	Entity backup statuses (Success, Failure, Expired, and/or Warning).
Compliance	Compliance statuses of the entities: • Success • Failure • Undefined: The Exclude policy is assigned to the entities or the entities do not have a policy assigned.

Sorting data in panels

In each panel, you can sort the data in the table columns in ascending or descending order, specify which columns are shown in the table, and arrange the order of the columns in the table to your preference. Depending on what you want to do, in the selected panel, follow these instructions:

I want to...	Instructions
Sort the data in the table columns in ascending or descending order.	1. Click the column header of the column that you want to sort. The column is sorted in ascending order, which is indicated by the ^ icon. 2. Click the column header again to sort the data in descending order, which is indicated by the v icon.
Specify which columns are	1. In the top right corner of the table, click

I want to...	Instructions
shown in the table.	⋮ icon. A list of table columns that can be shown or hidden is displayed. 2. On the list of the table columns, make sure that only the columns that you want to be shown are selected.
Arrange the order of the columns in the table.	1. Hover over the column header of the column that you want to move. If the column can be moved, this is indicated by the ✣ icon. 2. Click the ✣ icon, and then drag the column to the preferred position in the table.

Tip You can at any time reset the table columns to the default settings. To do so, in the top right corner of the table, click the ⋮ icon, and then click **Reset to defaults**.

Exporting the contents of the panel

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

Consideration

If you want to export only specific data, click **Filters**, select your filter criteria based on what kind of data you want to export to a file, and then click **Apply Filters**.

Procedure

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

Option	Description
Export to JSON	Exports the current table page to a JSON file.

Option	Description
(Current)	
Export to JSON (All)	Exports all table data to a JSON file.
Export to CSV (Current)	Exports the current table page to a CSV file.
Export to CSV (All)	Exports all table data to a CSV file.

Managing targets

If you have the proper permissions, you can view target information, edit target properties, activate or deactivate a target, or remove a target if you do not want to use it for storing protected data anymore.

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

Viewing target information

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

Target information	Description
Name	Name of the target.
Type	Type of target (NFS, SMB, Nutanix, Nutanix Objects, iSCSI, Amazon S3 / S3 Compatible, Azure, Google Cloud, QStar NFS, QStar SMB, or Data Domain). **Note** A tape target is represented by the icon, and a target that has WORM enabled is represented by the icon.
Health	Health status of the target: • Gray: Shows the initial target status before a health test. It also indicates an inactive target.

Target information	Description
	• Green: The target is in a healthy state with target utilization of less than the configured value (by default, 90%). • Yellow: Target utilization is over the configured value (by default, 90%). • Red: Target utilization is over the configured value (by default, 95%). It also indicates a target error state after a test task (for example, an I/O error occurred, the target is not accessible, the permission is denied, and so on). HYCU calculates if there is enough space on the target for storing backup data based on the following: • *If no previous backup is stored on the target:* The total provisioned space of all disks included in the virtual machine backup, regardless of whether the backup is full or incremental. • *If a previous backup is stored on the target:* The size of the last incremental backup for incremental backups, or the size of the last full backup for full backups or incremental backups if no previous incremental backup exists.
Size	Estimation of the amount of storage space that should be reserved for the backup files (in MiB, GiB, or TiB).
Utilization	Percentage of the specified target size that is already used for storing protected data.
Mode	Mode of the target: • Read/Write: You can use this target for backing up and restoring data. • Read Only: You can use this target only for restoring data. **Important** The Read-Only mode is automatically set on an imported target to prevent you from performing backups. Make sure not to change the mode of the imported targets.
Status	Status of the target:

Target information	Description
	• Active: You can use this target for backing up and restoring data. • Inactive: You cannot use this target for backing up and restoring data. This status indicates that the target is deactivated due to maintenance tasks (for example, adding new disks). For details on how to change the status of the target, see "Activating or deactivating a target" on page 421.

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

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

Editing a target

> ⚠ Caution Making any changes to the target location may result in data loss. Therefore, before specifying a new target location, make sure you have already moved the existing backup data to this new location on the same or a different server.

Prerequisites

Only if you plan to increase the size of an iSCSI target.

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

Considerations

- If you change the target settings in the policy assigned to the HYCU backup controller, make sure to update the note of the target's configuration.

- *For QStar tape targets:* If the status of the Integral Volume set is offline, the corresponding tape target is automatically deactivated in HYCU. When the Integral Volume set is remounted in QStar, make sure to activate the target. For details on how to do this, see "Activating or deactivating a target" on the next page.

Procedure

1. In the Targets panel, select the target that you want to edit, and then click **Edit**. The Edit Target dialog box opens.
2. Edit the selected target as required. For detailed information about target properties, see "Setting up targets" on page 100.

   > **Important** If you want to make specific changes to the NFS, SMB, Nutanix, iSCSI, or tape target, make sure you first detach the storage. For a list of possible changes and instructions, see "Detaching storage and changing target data" below.

3. *Only if you are editing an iSCSI target and you want to increase its size.* Select the **Extend target** check box.
4. Click **Save**.

After you edit the target, you will receive a message indicating whether the target was edited successfully. If you increased the size of the iSCSI target, you will also receive a message indicating whether the increase of the target size completed successfully.

Detaching storage and changing target data

If you want to change data for the NFS, SMB, Nutanix, iSCSI, or tape target, make sure that the storage is first detached from the HYCU backup controller to be able to perform the required changes:

Target type	Possible changes
NFS	Server name, IP address, or path to the shared folder
SMB	Server name, IP address, or path to the shared folder
Nutanix	URL
iSCSI	Portal IP address
Tape (QStar NFS and QStar SMB)	Web service endpoint

Procedure

1. Deactivate the target and detach the storage from the HYCU backup controller as described in "Activating or deactivating a target" below.
2. Make the required changes first on the server where the target is located, and then also in the HYCU web user interface as described in "Editing a target" on page 419.
3. Activate the target as described in "Activating or deactivating a target" below.

Activating or deactivating a target

Procedure

1. In the Targets panel, select the target that you want to activate or deactivate.
2. Change the status of the selected target by clicking ✓ **Activate** or ⊘ **Deactivate**.
3. If you are deactivating the target to change the data related to the NFS, SMB, Nutanix, iSCSI, or tape target, enable the **Detach storage** switch. For details on detaching storage from the HYCU backup controller, see "Detaching storage and changing target data" on the previous page.
4. *For target deactivation:* Click **Deactivate** to confirm that you want to deactivate the selected target.

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

Removing a target

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

Procedure

1. In the Targets panel, select the target that you want to remove, and then click 🗑 **Remove**.

   > 📋 **Note** If the target that you want to remove is used for archiving, make sure that no data archive with the specified archive target is used by any policy.

2. Click **Remove** to confirm that you want to remove the selected target.

Managing policies

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

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

Viewing policy information

You can view information about each policy in the list of policies in the Policies panel. This allows you to have an overview of the general status of the policies.

Consideration

The values for the backup RPO, RTO, and retention period that are defined in a policy are rounded to days, weeks, months, or years for display, but are stored and used internally as defined. For example, 30 days are rounded to one month in the HYCU web user interface.

The following information is available for each policy:

Policy information	Description
Name	Name of the policy.
Compliance	Compliance status of the policy: • Success • Failure • Undefined A policy is considered compliant if all entities to which this policy is assigned are compliant with the policy settings. For detailed information about the compliance status of entities, see “Viewing entity details” on page 399.
VM Count	Total number of virtual machines and servers that have the particular policy assigned to them.
App Count	Total number of applications that have the particular policy

Policy information	Description
	assigned to them.
Description	Description of the policy (how often backup and restore jobs are performed).

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

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

Editing a policy

Considerations

- If you edit a policy in such a way that you enable the Copy option, the next backup of the virtual machines and volume groups to which this policy is assigned will be a full backup.
- The changes that you make when editing a policy are applicable for future backups and do not affect existing backups.

Procedure

1. In the Policies panel, select the policy that you want to edit, and then click ✎ **Edit**. The Edit Policy dialog box opens.
2. Edit the selected policy as required. For detailed information about policy properties, see “Creating a policy” on page 134.
3. Click **Save**.

Deleting a policy

Considerations

- A policy that is assigned to one or more entities for which backups are scheduled cannot be deleted. If you want to delete such a policy, you must first abort the scheduled backups. For details on how to abort queued jobs, see “Managing HYCU jobs” on page 385.

- If you delete a policy that is assigned to one or more entities, keep in mind that no further backups will be performed for these entities.
- You cannot delete the Exclude policy.

Procedure

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

Performing a manual backup

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

Prerequisite

Only if backing up a volume group manually. Make sure a policy is assigned directly to the volume group. If the policy is assigned only to the virtual machine to which the volume group is attached, performing a manual backup for the selected volume group is not possible.

Considerations

- You can prevent your manual backups from interfering with the scheduled backups determined by the RPO specified in the policy. To do so, set the `exclude.manually.run.backups.regarding.rpo` configuration setting to `true`. This is especially important if you define backup windows because performing a manual backup can prevent the backup scheduled in the backup window from starting, which can result in data not being protected until the next backup window or the next manual backup. For details on how to customize HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
- *Only if you enabled the Incremental forever backup option for a file share or a bucket.* Make sure not to enable the **Force full backup** switch for such a file share or bucket. Otherwise, a full backup will be performed instead of an incremental one.

Procedure

1. In the Applications, Virtual Machines, Volume Groups, Shares, or Buckets panel, select which entities you want to back up.
2. Click **Backup** to perform the backup of the selected entities.
3. *Only if you selected Target as the backup target type in your policy.* Enable the **Force full backup** switch if you want to perform a full backup. Otherwise, HYCU will perform a full or incremental backup based on the settings defined in your policy.
4. Click **Yes** to confirm that you want to start the manual backup.

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

Setting up a validation policy

As an alternative to manually performing the backup validation for a virtual machine and verifying that the virtual machine has no corrupted backups, you can set up a validation policy and schedule the backup validation according to the values that you define in your validation policy. For details on how to validate the virtual machine backup by creating a virtual machine clone, see "Validating the virtual machine backup" on page 228.

> Important HYCU automatically creates a clone of the virtual machine while performing the backup validation.

Prerequisites

- If you plan to select a vSphere storage container for the virtual machine copy, the latest version of VMware Tools must be installed on the virtual machine.
- *Only if you plan to specify the Advanced validation type.*
  - Credentials must be assigned to the virtual machine. For prerequisites, limitations, considerations, and instructions, see "Enabling access to application data" on page 252.
  - A network card must be added to the virtual machine.
- *For Nutanix AHV clusters:* The Nutanix cluster that hosts the virtual machine for which you want to set up the validation policy must be registered with Prism Central and your Prism Central user must have a role with sufficient

permissions assigned. For details, see “Configuring Prism Central user permissions” on page 556.

Limitations

- Performing the backup validation is not supported for the following:
  - The HYCU backup controller and virtual machines running in AWS GovCloud (US), Azure, or Azure Government environments.
  - Virtual machines that have volume groups attached by using iSCSI.
- If you assign a validation policy to a virtual machine that has only the Snapshot tier available, the backup validation will not be performed. However, you can manually validate the backup of such a virtual machine. For instructions, see “Validating the virtual machine backup” on page 228.

Considerations

- Network conflicts may occur during the backup validation if the virtual machine is configured with a static IP address, resulting in unreliable backup validation data.
- *Only if you plan to specify the Advanced validation type when performing the backup validation for a Windows virtual machine.* Checking for disk errors may fail in some cases, which does not mean that your virtual machine is corrupted. However, it is highly recommended that you check the status of such a virtual machine manually.
- After the backup validation is performed, consider the following:
  - You can view the backup validation status of a virtual machine in the Validation column in the Virtual Machines panel (represented by an icon). By pausing on the icon, you can also see which validation policy is assigned to the virtual machine.
  - The Exclude policy is automatically assigned to the cloned virtual machine.

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

Procedure

1. In the Virtual Machines panel, select one or more virtual machines for which you want the backup validation to be performed.

> **Tip** You can update the list of virtual machines by clicking **Refresh**. To narrow down the list of displayed virtual machines, you can use the filtering options described in "Filtering and sorting data" on page 406.

2. Click **Validation**. The Validation Policies dialog box opens.
3. Click **New**.
4. Enter a name for your validation policy and, optionally, its description.
5. From the Storage container drop-down menu, select where you want to clone the virtual machine for which you are performing the backup validation.
6. From the Restore from drop-down menu, select which tier you want to use for the backup validation. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**
   - **Backup**
   - **Copy**
   - **Archive**

> **Note** When selecting the restore point tier, consider the following:
> - If you select Automatic, the tier for the backup validation is by default selected in the following priority order: Backup > Copy > Archive. This means that HYCU will always use the first available tier in the specified order for the backup validation. However, you can at any time change this default behavior by customizing the `backup.validation.restore.source.priority.order` configuration setting in the HYCU `config.properties` file and adjusting the tier order to your data protection needs. For details on how to customize HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
> - *Only if you select the Archive tier and the data is stored on a QStar tape target.*
>
> To view tape target information, click **Tape Info**. The following information is displayed:
> - Target name and UUID
> - Barcode label of the media
> - Information on whether the media is online or offline
> - Location of the media

> You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to a JSON or CSV file.

7. From the Keep VM after validation drop-down menu, depending on whether you want to keep the virtual machine after the backup validation is performed, select one of the following options:

Option	Description
Always	The virtual machine will be kept after the backup validation is performed.
On validation error	The virtual machine will be kept after the backup validation is performed only if a validation error occurs during the validation.
Never	The virtual machine will be automatically deleted after the backup validation is performed.

8. From the Validation type drop-down menu, select one of the following types:

Validation type	Description
Basic	During the backup validation, the following tasks will be performed: • The virtual machine will be cloned and turned on. • The guest OS will be shut down.
Advanced	During the backup validation, the following tasks will be performed: • The virtual machine will be cloned and turned on. • Any applications running on the virtual machine will be discovered. • Virtual disks will be validated, which includes checking the virtual machine file system and existing disks on the virtual machine. For Windows virtual machines, checking for disk errors is also performed. • The custom scripts will be run, if specified.

Validation type	Description
	• The guest OS will be shut down.

9. *Only if you selected the Advanced validation type.* Do the following:
   a. Enable the **Run custom script** switch if you want the custom script to be run on the virtual machine as part of the backup validation process, and then make sure that the proper path to the script is specified.

      > **Note** The script returns an exit code of 0 for success and any other value for failure.

   b. From the Network drop-down menu, select the network for the virtual machine.
10. Click **Next**.
11. Depending on whether you want backup validation for the virtual machine to be performed on a daily, weekly, monthly, and/or yearly basis, add any of the preferred backup validation options to the list of the enabled options by clicking it:
    - **Daily**
    - **Weekly**
    - **Monthly**
    - **Yearly**
12. In the Start at fields, specify the hour and the minute when the backup validation job should start.
13. From the Time zone drop-down menu, select the appropriate time zone for the backup validation job.

    > **Note** All backup validation jobs are by default started based on the HYCU backup controller time zone.

14. Depending on the selected backup validation options, specify at what intervals you want backup validation to be performed:

Backup validation option	Instructions
Daily	a. In the Recur every field, specify whether you want backup validation to be performed every day or every few days.

Backup validation option	Instructions
	b. Use the **Apply only on weekdays** switch if you want backup validation to be performed only on weekdays.
Weekly	a. In the Recur every field, specify whether you want backup validation to be performed every week or every few weeks. b. Select one or more days of the week on which you want backup validation to be performed.
Monthly	a. In the Recur every field, specify whether you want backup validation to be performed every month or every few months. b. Select whether you want backup validation to be performed on the same day of the month (for example, on the fifth day of the month), or on a specific day of the month (for example, on the second Friday of the month).
Yearly	a. In the Recur every field, specify whether you want backup validation to be performed every year or every few years. b. Select whether you want backup validation to be performed on the same day of the preferred month (for example, on the fifth day of January), or on a specific day of the preferred month (for example, on the second Friday of April).

15. Click **Save**.
16. Click **Assign**.

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

Running malware scans manually

HYCU runs malware scans automatically after you create an R-Shield policy with enabled malware scanning and enable the R-Shield option in a policy that is assigned to an entity. However, you can at any time decide to manually run a malware scan of all files in the backup or only of the files that change between the selected restore point and the last successful backup. For example, this is useful if you want to run a malware scan for a specific restore point or if a malware scan job fails.

Prerequisites

- You must have the Administrator user role assigned.
- The R-Shield option must be enabled in the policy that is assigned to the entity and malware scanning must be enabled in the corresponding R-Shield policy. For details, see "Creating an R-Shield policy" on page 148.
- *Only if you plan to run malware scanning for file shares:*
  - The selected restore point must contain the Snapshot tier.
  - A HYCU instance is required. For details, see "HYCU instances" on page 25.

Limitations

- You can run malware scanning for one restore point at a time.
- If any malware scans or restore jobs are already in progress for the selected restore point, you cannot manually run a malware scan.
- *For file shares:* If the selected restore point contains a full backup, you can only scan all files.
- *For the volume groups that are not attached to a virtual machine, for generic file shares, and for buckets:* Malware scanning is not supported.

Depending on the type of data that you want to scan for malware, access one of the following panels:

- Accessing the Applications panel

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

- Accessing the Virtual Machines panel

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

- Accessing the Shares panel

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

- Accessing the Volume Groups panel

  To access the Volume Groups panel, in the navigation pane, click **Volume Groups**.

Procedure

1. In the relevant panel, click the entity whose data you want to scan for malware.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

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

3. Click the ⋮ icon, and then click **Run Malware Scan**.
4. From the Scan From drop-down menu, select which tier you want to use for the scan. Your restore point can contain one or more tiers among which you can select:
   - **Archive**
   - **Backup**
   - **Copy**
   - **Snapshot**
5. Depending on the preferred scope of scan, do one of the following:
   - Select **Scan all files** to scan all of the files that the selected restore point contains.
   - Select **Scan changed files** to scan the files that have been created or modified since the last successful backup.
6. Click **Run**.

Overriding the R-Shield status

If after reviewing the detection statuses of any protected entity or restore point, you determine that it is not a malicious anomaly or a suspicious item, you can override the anomaly detection or the malware detection status and set it to OK. For details on the R-Shield statuses, see "Viewing entity details" on page 399.

The R-Shield status can be manually overridden in the following ways:

- You can override the R-Shield status for one or more entities. By doing so, you set the R-Shield status of the latest available restore point for the entity.
- You can manually override the R-Shield status for one or more available restore points.

Considerations

- After you manually override the R-Shield status for an entity or restore point, the following applies:
  - You cannot remove the manually overridden status.
  - If you manually override the status from Suspicious to OK after a suspicious item is detected, the status will stay OK for the future restore points until a new suspicious item is detected.
- You can manually set the R-Shield status for the protected entities that do not have the R-Shield option enabled in the assigned policies. In this case, at least one restore point must be available for each protected entity.
- You cannot manually set or override the malware detection status for the entities for which malware scanning is not supported. For details, see "Creating an R-Shield policy" on page 148.

Depending on the entity for which you want to override the R-Shield status, access one of the following panels:

- Accessing the Applications panel

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

- Accessing the Virtual Machines panel

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

- Accessing the Shares panel

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

- Accessing the Volume Groups panel

  To access the Volume Groups panel, in the navigation pane, click **Volume Groups**.

- Accessing the Buckets panel

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

Procedure

1. In the relevant panel, click one or more entities whose R-Shield status you want to override manually.
2. Depending on how you want to override the R-Shield status, do one of the following:
   - *For virtual machines:* In the toolbar, click the icon, and then click **R-Shield Status**.
   - *For other entities:* Click **R-Shield Status**.
   - *For restore points of a virtual machine:*
     a. In the Detail view, select one or more restore points.
     b. Click the icon, and then click **R-Shield Status**.
   - *For restore points of other entities:* In the Detail view, select one or more restore points, and then click **R-Shield Status**.
3. Depending on which detected R-Shield status you want to override, do one of the following:
   - Under Anomaly Detection, select the preferred anomaly detection status, and then enter the reason for the manual status override.
   - Under Malware Detection, select the preferred malware detection status, and then enter the reason for the manual status override.
4. Click **Save**.

The detected R-Shield status is overridden.

Archiving data manually

HYCU archives your data automatically after you enable the Archiving policy option. However, you can archive data manually at any time (for example, if you want to archive data for a specific restore point or if an archiving job fails).

Prerequisites

- You must have the Administrator, Backup and Restore Operator, or Backup Operator user role assigned.
- The Archiving option must be specified in the assigned policy and a data archive must be created.

Considerations

- Retention time for archives is calculated from the date and time when the restore point for the entity whose data you are archiving was created.
- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for archiving data manually.
- *Only if you selected Snapshot as the backup target type in your policy.* The configuration settings that HYCU uses for archiving are the ones that the virtual machine has at the time when archiving starts.

Depending on the type of data that you want to archive, access one of the following panels:

- #### Accessing the Applications panel

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

- #### Accessing the Virtual Machines panel

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

- #### Accessing the Shares panel

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

- #### Accessing the Volume Groups panel

To access the Volume Groups panel, in the navigation pane, click **Volume Groups**.

- Accessing the Buckets panel

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

Procedure

1. In the relevant panel, click the entity whose data you want to archive.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

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

3. Depending on the type of entity whose data you want to archive, do one of the following:
   - *For a virtual machine:* Click the ⋮ icon, and then click **Run Archiving**.
   - *For other entities:* Click **Run Archiving**.
4. Select the preferred archiving option.
5. Click **Run**.

Recreating snapshots

Recreating snapshots is required in the following scenarios:

- If you plan to restore files from a snapshot (and not directly from a target) and no snapshot is available for the selected virtual machine restore point.
- If you plan to restore applications, export virtual disks, or restore files that are stored in the archive access tier on an Azure target.

Limitation

Recreating snapshots is not supported for XenServer, Azure Local, Hyper-V, AWS GovCloud (US), Azure, and Azure Government.

Consideration

If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives

missing or being stored on a deactivated target), you cannot use this tier for recreating snapshots.

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

Procedure

1. In the Virtual Machines panel, select the virtual machine whose snapshot you want to recreate.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

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

3. Click the ⋮ icon, and then click **Recreate Snapshot**. The Recreate Snapshot dialog box opens.
4. From the Storage container drop-down menu, select where you want to recreate the snapshot.
5. From the Restore from drop-down menu, select which tier you want to use for recreating the snapshot. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest snapshot creation.
   - **Backup**
   - **Copy**
   - **Archive**

   > **Note** *Only if you select the Archive tier and the data is stored on a QStar tape target.*
   >
   > To view tape target information, click **Tape Info**. The following information is displayed:
   >
   > - Target name and UUID
   > - Barcode label of the media
   > - Information on whether the media is online or offline
   > - Location of the media
   >
   > You can also export tape target information by clicking **Export**, and then selecting if you want to export all or only offline tape information to

a JSON or CSV file.

6. Click **Recreate**.

Adjusting the HYCU backup controller resources

When storage, vCPU, or memory utilization is exceeded (that is, when the utilization of any of these resources is greater than 90 percent), their values that are indicated by circles become red in the HYCU Controller widget in the Dashboard panel. You can adjust these resources to better meet your data protection goals.

Procedure

1. Sign in to the management console of the environment in which HYCU is deployed.
2. Shut down the HYCU backup controller.
3. Modify the storage, vCPU, or memory configuration as required.
4. Turn on the HYCU backup controller.

For instructions on how to perform these steps, see your platform documentation.

Setting up the appearance of your HYCU web user interface

When setting up the appearance of your HYCU web user interface, you can do the following:

- Adjust the table density to determine how close or far apart the rows in the tables should be.
- Choose to show or hide the separator line between the rows in the tables.
- Switch your HYCU web user interface to light or dark mode. HYCU by default uses the color mode exposed by your browser.

Accessing the Appearance dialog box

To access the Appearance dialog box, click at the upper right of the screen, and then select **Appearance**.

Procedure

1. In the Appearance dialog box, do the following:
   - Under Table density, select **Default density** or **High density** depending on how close or far apart you want the rows in the tables to be.
   - Enable the **Row dividers** switch if you want to show the separator line between the rows in the tables.
   - Do one of the following:
     - *If you want to use the dark mode:* Enable the **Dark mode** switch.
     - *If you want to use the light mode:* Disable the **Dark mode** switch.
2. Click **Close**.

The changes take place immediately without the need to sign out and sign in again to the HYCU web user interface. The preferred appearance is remembered for the next time you sign in to HYCU.

Chapter 11

Managing users

The HYCU user management system provides security mechanisms to help prevent unauthorized users from accessing protected data. Only users that are given specific rights have access to the data protection environment. These users can be authenticated either by HYCU or any of the supported identity providers. For details on identity providers, see “Integrating HYCU with identity providers” on page 471.

Each user that signs in to HYCU must belong to one of the HYCU groups—an infrastructure group or a self-service group—and have a user role assigned.

For details on HYCU groups and user roles, see “HYCU groups” below and “User roles” on page 442.

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

HYCU groups

For a consolidated user management experience, HYCU provides two types of groups to which users can belong.

Group	Description
Infrastructure group	Created by default during the deployment of the HYCU virtual appliance and already includes a built-in user with the Administrator user role assigned (represented by)—cannot be edited, deactivated, and deleted. Users can be added to this group by an infrastructure group administrator (an infrastructure group user with the Administrator user role assigned).
Self-service	Must be created by an infrastructure group administrator

Group	Description
group	and represents a customer or a department responsible for a specific set of entities in the data protection environment. Users can be added to this group by an infrastructure group administrator. **Important** If a specific self-service group is deleted, all data that is backed up by this group is deleted from the database.

You can manage users only if you have an Administrator role assigned. However, keep in mind that the scope of user management actions that you can perform differs depending on whether you belong to the infrastructure or self-service group. As an infrastructure group administrator, you can manage users and groups throughout the whole data protection environment, whereas as a self-service group administrator, you can manage only the group you belong to. The following diagram shows which user-related actions you can perform:

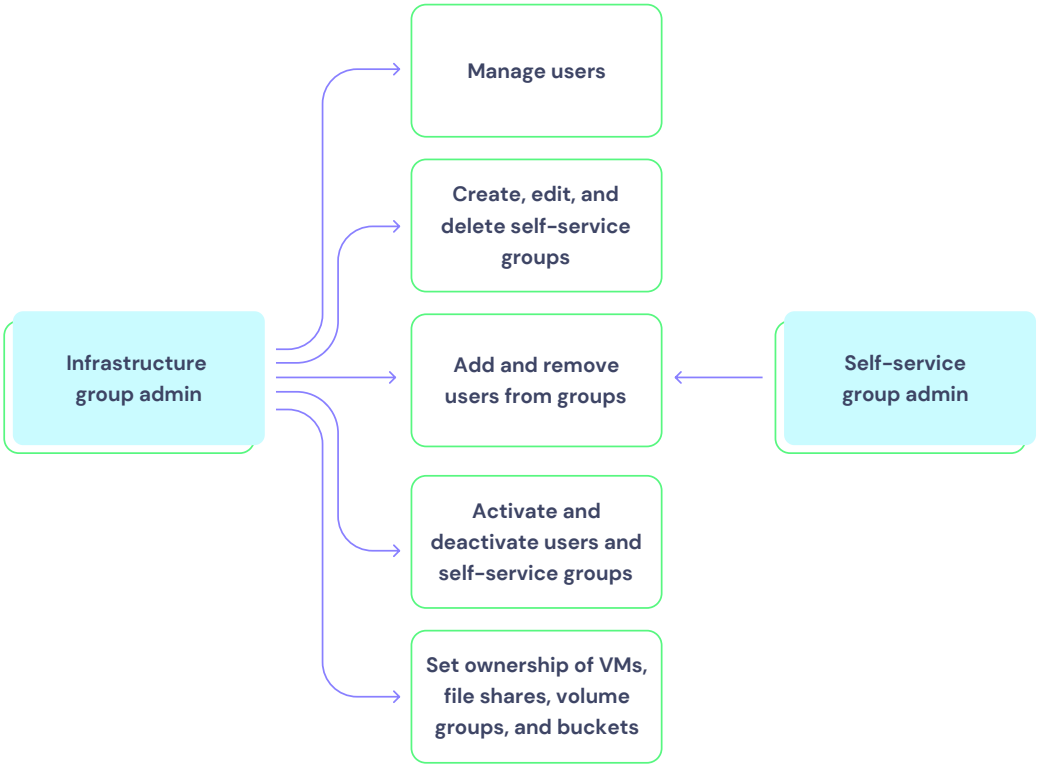

Figure 11-1: User management actions performed by the infrastructure and self-service group administrators

Depending on the HYCU group to which you belong and the assigned user role, you can perform only specific actions in the data protection environment. For details on user roles, see “User roles” below.

User roles

Each user in a group has an assigned role that determines the scope of actions the user can perform in the data protection environment. This means that access to data and information within the data protection environment is limited based on the role that the user has assigned. If a user is a member of multiple groups, this user can have different roles assigned in different groups, depending on the business needs, and can switch between these groups while being signed in to HYCU.

Depending on the group to which a user belongs, the user can perform the following actions:

Role	Infrastructure group	Self-service group
Administrator	• Perform all actions in the data protection environment.	• Assign policies. • Back up and restore virtual machines, applications, file shares, volume groups, and buckets. • Manage data retention. • Perform virtual machine backup validation by using the Validate VM backup option. • Assign and unassign validation policies. • Add and remove users from groups. • Perform all report management actions. • Add, edit, and remove cloud accounts.

Role	Infrastructure group	Self-service group
Viewer	• View information about applications, virtual machines, file shares, volume groups, buckets, policies, targets, jobs, events, users, generated report versions, and settings available through the ⚙ **Administration** menu in the data protection environment.	• View information about applications, virtual machines, file shares, volume groups, buckets, policies, jobs, events, and generated report versions in the data protection environment.
Backup Operator	• View the same information as Viewer. • Define a backup strategy. • Back up virtual machines, file shares, volume groups, and buckets that are not owned by any self-service group, and back up applications.	• View the same information as Viewer. • Assign policies. • Back up virtual machines, applications, file shares, volume groups, and buckets.
Restore Operator	• View the same information as Viewer. • Restore virtual machines, file shares, volume groups, and buckets that are not owned by any self-service group, and restore applications. • Perform virtual machine backup validation by using the Validate VM backup option. • Assign and unassign validation policies.	• View the same information as Viewer. • Restore virtual machines, applications, file shares, volume groups, and buckets. • Perform virtual machine backup validation by using the Validate VM backup option. • Assign and unassign validation policies.

Role	Infrastructure group	Self-service group
Backup and Restore Operator	• View the same information as Viewer. • Define a backup strategy. • Back up and restore virtual machines, file shares, volume groups, and buckets that are not owned by any self-service group, and back up and restore applications. • Perform virtual machine backup validation by using the Validate VM backup option. • Assign and unassign validation policies.	• View the same information as Viewer. • Assign policies. • Back up and restore virtual machines, applications, file shares, volume groups, and buckets. • Perform virtual machine backup validation by using the Validate VM backup option. • Assign and unassign validation policies.

Setting up a user environment

Before users can start using HYCU for data protection, you must give them rights to access data within the data protection environment. By creating a user and adding the user to a group, you allow the user to access only the defined data protection environment and to perform a set of actions specified by the assigned role:

Task	Performed by...	Instructions
1. Create a new user.	An infrastructure group administrator	“Creating a user” on the next page
2. Add a user to a user group.	An infrastructure or a self-service group administrator	“Adding a user to a group” on page 450

While setting up a user environment, you can tailor it to the user's needs by performing one or more of the following tasks:

Task	Performed by...	Instructions
Create a new self-service group.	An infrastructure group administrator	“Creating a self-service group” on page 451
Set ownership of virtual machines, file shares, volume groups, and buckets.	An infrastructure group administrator	“Setting ownership” on page 452
Enable or disable specific groups or users from signing in to HYCU.	An infrastructure group administrator	“Activating or deactivating users or self-service groups” on page 457

Accessing the Self-Service panel

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

Creating a user

Prerequisites

- *For using two-factor authentication*: An appropriate authenticator must be set up. Depending on the authentication method:
  - A time-based one-time password (OTP) authentication application, such as Google Authenticator on your mobile phone.
  - A FIDO-compatible authenticator, such as a hardware key, fingerprint reader, or similar.
- *For integrating HYCU with identity providers:* In an identity provider environment, HYCU must be assigned as an application to users for whom you want to enable signing in to HYCU by using the identity provider. For detailed instructions on how to integrate HYCU with identity providers, see “Integrating HYCU with identity providers” on page 471.

Limitations

- You cannot add the Active Directory primary group (usually the Domain Users group) as an AD group.
- If certificate authentication is enabled, setting up two-factor authentication for AD users is not supported.

Considerations

- The members of an identity provider group (Active Directory, OpenID Connect (OIDC), or LDAP) are listed as individual users. For AD and LDAP users, this allows you to enable two-factor authentication and set the preferred language for each user.
- *Only if you plan to add an LDAP group.* If the LDAP group that you add contains a user with the at sign (@) in their name, such a user will not be able to sign in to HYCU.

Procedure

1. In the Self-Service panel, click **Manage Users**, and then click **New**.
2. Depending on what kind of user you are adding, enter one of the following:
   - *For a HYCU user, an AD user, an OIDC user, an OIDC group, or an LDAP user:* User name

     > **Important** Keep in mind the following:
     > - *For an AD user:* When entering a name, make sure it complies with the SAM account name limitations—name length may not exceed 20 characters and contain any of the following characters: "/ \ [ ] : ; | = , + * ? < >. In addition, HYCU does not allow the at sign (@) in the name.
     >
     >   If your environment requires it, these limitations can be overridden by editing the `ad.username.filter.regex` configuration setting. However, this is not supported and could cause authentication issues. For details on how to customize HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
     > - *For an LDAP user:* HYCU does not allow the at sign (@) in the name.

   - *For an AD group or an LDAP group:* Common name
3. From the Authentication Type drop-down menu, select one of the following authentication types, and then follow the instructions:

Authentication type	Instructions
HYCU	a. From the Language drop-down menu, select the preferred language for the user.

Authentication type	Instructions
	b. In the Name field, enter a display name for the user. c. *Optional.* In the Email field, enter the email address of the user. d. In the Password field, enter the user password. > Note The minimum password length is six characters.
OIDC User	a. From the Language drop-down menu, select the preferred language for the user. b. From the Identity Provider drop-down menu, select the identity provider. c. In the Identity Provider User ID field, enter the ID of the identity provider user. > Note Depending on your identity provider, the user ID corresponds to the following: > • *Active Directory Federation Services:* Object GUID > • *Google:* User email address > • *Keycloak*: User ID > • *Microsoft:* Object ID > • *Okta:* Part of the URL when you navigate to the user's profile > For details, see the respective identity provider documentation.
OIDC Group	a. From the Language drop-down menu, select the preferred language for the group. b. From the Identity Provider drop-down menu, select the identity provider. c. In the Identity Provider Group ID field, enter the ID of the identity provider group. > Note Depending on your identity provider, the group ID corresponds to the

Authentication type	Instructions
	following: • *Active Directory Federation Services:* Object GUID • *Keycloak*: Group ID • *Microsoft:* Group Object IDs • *Okta:* Group name For details, see the respective identity provider documentation.
AD User	a. From the Language drop-down menu, select the preferred language for the user. b. From the Identity Provider drop-down menu, select the Active Directory the AD user belongs to.
AD Group	a. From the Language drop-down menu, select the preferred language for the user. b. From the Identity Provider drop-down menu, select the Active Directory the AD group belongs to.
LDAP User	a. From the Language drop-down menu, select the preferred language for the user. b. From the Identity Provider drop-down menu, select the identity provider that the LDAP user belongs to.
LDAP Group	a. From the Language drop-down menu, select the preferred language for the group. b. From the Identity Provider drop-down menu, select the identity provider that the LDAP group belongs to.

4. *Only if you are adding a HYCU user, an AD user, an AD group, an LDAP user, or an LDAP group.* Use the **Two-factor authentication** switch if you want to enable two-factor authentication for the user, and then select one of the following two-factor authentication methods:

- **Time-based one-time password**

  This option enables the use of a time-based one-time password (OTP) generated by an OTP application. The user needs to set up an OTP during the first sign-in after two-factor authentication is enabled.

- **FIDO**

  This option enables the use of an authenticator complying with FIDO protocols (FIDO authenticator). The user needs to register a FIDO authenticator. For details, see "Managing FIDO authenticators" on page 544.

5. *Only if you enabled two-factor authentication.* To prevent the user from disabling two-factor authentication, make sure the **User cannot disable two-factor authentication** check box is selected. If you clear the check box, the user can disable two-factor authentication. An infrastructure group administrator can disable two-factor authentication even if this option is enabled.

   > **Note** If a user disables two-factor authentication, the administrator is notified with a security warning.

6. Click **Next**.
7. *Only if you want to add the user to a user group.* Do the following:
   a. From the User Group drop-down menu, select the user group to which you want to add the user.

      > **Important** You can add the user to the user group also later by following the procedure described in "Adding a user to a group" on the next page. In any case, keep in mind that the user must be added to at least one user group before they can sign in to HYCU.

   b. From the User Role drop-down menu, select the role that you want to assign to the user (**Administrator**, **Backup and Restore Operator**, **Restore Operator**, **Backup Operator**, or **Viewer**).
8. Click **Save**.

The user is added to the list of all users.

You can later do the following:

- Edit any of the existing HYCU or identity provider users by clicking **Edit** and making the required modifications. Keep in mind that the built-in user, AD users, and AD groups cannot be edited.

- Enable or disable specific users from signing in to HYCU. For details, see “Activating or deactivating a user” on page 457.
- Delete any of the existing users by clicking **Delete**. Keep in mind that the built-in user cannot be deleted.

Adding a user to a group

Prerequisite

Only if you want to add a user to a self-service group. The self-service group must be created. For instructions, see “Creating a self-service group” on the next page.

Considerations

- You can add a user to multiple groups in which the user can have different user roles assigned. For details on user roles, see “User roles” on page 442.
- If an AD, LDAP, or OIDC group is added to a group in HYCU, all AD, LDAP, or OIDC group users acquire the role that is assigned to the AD, LDAP, or OIDC group.
- If an AD, LDAP, or OIDC user is a member of multiple AD, LDAP, or OIDC groups that are added to the same group in HYCU, the user acquires the role with the highest privilege level. User roles are prioritized in the following order: Administrator > Backup and Restore Operator > Restore Operator > Backup Operator > Viewer. However, keep in mind that a role assigned to an AD, LDAP, or OIDC user independently of an AD, LDAP, or OIDC group always takes precedence over a role within a group.
- If an AD, LDAP, or OIDC group user is a member of multiple AD, LDAP, or OIDC groups with different two-factor authentication methods, the method with the highest priority among all the groups will be used for two-factor authentication. The methods are prioritized in the following order: FIDO > OTP > none.
- If any of the AD, LDAP, or OIDC users that you are adding to a group share the same user name, they are all added to the selected group. Make sure to remove any duplicate user names that are not required.
- AD, LDAP, or OIDC group users cannot be added to a group separately, but only as part of the AD, LDAP, or OIDC group.

Procedure

1. In the Self-Service panel, in the Detail view, select the group to which you want to add a user.
2. Click ✚ **Add to Group**.

   > **Note** You can add the user to the infrastructure group that is created by default or a self-service group that you must create yourself.

3. Provide the username of the user that you want to add to the selected group.

   > **Important** *For AD users, AD groups, LDAP users, and LDAP groups:* If you are signed in to HYCU as a self-service group member, enter the user name or the common name in one the following formats: `<Name>@<Domain>` or `<Domain>\<Name>`. For LDAP, `<Domain>` is the domain that was configured when the identity provider was added to HYCU.

4. From the User Role drop-down menu, select the role that you want to assign to the user (**Administrator**, **Backup and Restore Operator**, **Restore Operator**, **Backup Operator**, or **Viewer**).
5. Click **Add User**.

You can later do the following:

- Assign a different role to a user by selecting the user whose role you want to change, and then clicking ✎ **Change Role in Group**.
- Remove a user from a group by selecting the user that you want to remove from the group, and then clicking 🗑 **Remove from Group**.

> **Note** Assigning a different role to a user or removing a user from a group cannot be done for the following types of users:
> - Built-in administrator in the infrastructure group.
> - User that you are currently signed in as.

Creating a self-service group

Prerequisite

Only if you plan to use the Advanced option that allows you to automatically assign virtual machines to the self-service group that you are creating. Tags must be assigned to virtual machines on their source. For details, see “Setting up automatic virtual machine assignment by using tags” on page 453.

Procedure

1. In the Self-Service panel, click **New**.
2. Enter a self-service group name and, optionally, its description.
3. *Only if you want virtual machines to be automatically assigned to the self-service group you are creating.* Click **Advanced**, enter a key and a value, and then click **Add**. If required, repeat this step for all the keys and the values that you want to add.

   For details, see "Setting up automatic virtual machine assignment by using tags" on the next page.
4. Click **Save**.

You can later do the following:

- Add users to groups. For details, see "Adding a user to a group" on page 450.
- Edit any of the existing self-service groups by clicking **Edit** and making the required modifications.
- Allow users belonging to a specific self-service group to see only policies whose names start with their group name followed by an underscore (for example, HYCUGroup_Policy1) and the Exclude policy (alongside of other policies already assigned to the virtual machines, file shares, volume groups, and buckets whose owners they are). To do so, in the HYCU `config.properties` file, set the `policies.group.specific.synchronized` configuration setting to `true`. Keep in mind that such policies can be edited or deleted only if they are not assigned to any entity. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
- Enable or disable specific self-service groups from signing in to HYCU. For details, see "Activating or deactivating a self-service group" on page 458.
- Delete any of the existing self-service groups by clicking **Delete**.

Setting ownership

By setting ownership of virtual machines, file shares, volume groups, and buckets, you enable specific groups to protect only the assigned virtual machines, file shares, volume groups, and buckets. Depending on the entity to which you want to assign an owner, see one of the following sections:

- "Setting ownership of virtual machines" on the next page
- "Setting ownership of file shares" on page 455

- “Setting ownership of volume groups” on page 456
- “Setting ownership of buckets” on page 456

Setting ownership of virtual machines

When setting ownership of virtual machines, you can use one of the following approaches:

Approach to setting ownership	Instructions
Assign tags to virtual machines on their source, and then specify the corresponding keys and values in HYCU. **Important** You cannot set ownership of Azure Local or Hyper-V virtual machines by using this approach.	“Setting up automatic virtual machine assignment by using tags” below
Select virtual machines in HYCU, and then assign these virtual machines to a specific self-service group.	“Setting up virtual machine assignment by using the Assign option” on the next page

Considerations

- Assigning virtual machines to self-service groups by using tags takes precedence over assigning policies by using the Assign option. This means that the tag added to the virtual machine defines to which self-service group the virtual machine will be assigned (even if this virtual machine is already assigned to another self-service group by using the Assign option).
- A virtual machine is not unassigned from its self-service group if a tag is removed from the virtual machine.

Setting up automatic virtual machine assignment by using tags

By setting up automatic virtual machine assignment, you ensure that virtual machines are automatically assigned to self-service groups.

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

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

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

Prerequisite

For Nutanix AHV clusters: The Nutanix cluster that hosts the virtual machines for which you want to set up automatic virtual machine assignment must be registered with Prism Central and your Prism Central user must have a role with sufficient permissions assigned. For details, see "Configuring Prism Central user permissions" on page 556.

Procedure

1. Sign in to the management console of your data protection environment.
2. Assign tags to virtual machines for which you want to set up automatic assignment. For instructions, see your platform documentation.
3. Sign in to the HYCU web user interface.
4. Specify the matching keys and values for the specific self-service group as described in "Creating a self-service group" on page 451.

   > **Important** Depending on your data protection environment, the key and the value that you should enter represent the following:
   >
   > - *For Nutanix AHV clusters:* The name and the value of the category.
   > - *For Nutanix ESXi clusters or vSphere environments:* The tag name and the category of the tag, or the attribute and the value of the custom attribute.
   > - *For XenServer environments:* The name of the tag for both the key and the value.
   > - *For AWS GovCloud (US), Azure, or Azure Government environments:* The name and the value of the tag.

Setting up virtual machine assignment by using the Assign option

Consideration

When changing ownership of virtual machines, you can choose whether you want data protected by a specific owner to be kept or deleted. If you choose to keep data protected by the specific owner, such virtual machines will be kept in

HYCU with the Protected deleted status. Restoring these virtual machines by using the Restore VM option is not possible.

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

Procedure

1. In the Virtual Machines panel, select the virtual machines to which you want to assign an owner.
2. In the toolbar, click the ⋮ icon, and then click **Owner**.
3. From the list of groups, select which group you want to assign as the owner of the selected virtual machines, and then click **Assign**.

   > Important If a virtual machine or an application has backup or restore jobs in progress, or a scheduled backup task in the queue, you cannot assign a new group to the relevant virtual machine.

Depending on the needs of a specific data protection environment, you can at any time remove the owner from the virtual machines by selecting the virtual machines from which you want to remove the owner, and then clicking **Owner** followed by **Unassign**.

Setting ownership of file shares

Consideration

When changing ownership of file shares, you can choose whether you want data protected by specific owners to be kept or deleted. If you choose to keep data protected by the specific owner, such file shares will be kept in HYCU with the Protected deleted status.

> Accessing the Shares panel
>
> To access the Shares panel, in the navigation pane, click **Shares**.

Procedure

1. In the Shares panel, select file shares to which you want to assign an owner, and then click **Owner**.
2. From the list of groups, select which group you want to assign as an owner of the selected file shares, and then click **Assign**.

> **Important** If any backup or restore job for a file share is already in progress, or a scheduled backup task is in the queue, you cannot assign a new group to this file share.

Depending on the needs of a specific data protection environment, you can at any time remove an owner from the file shares by selecting the file shares from which you want to remove the owner, and then clicking **Owner** followed by **Unassign**.

Setting ownership of volume groups

Consideration

When changing ownership of volume groups, you can choose whether you want data protected by specific owners to be kept or deleted. If you choose to keep data protected by the specific owner, such volume groups will be kept in HYCU with the Protected deleted status.

> Accessing the Volume Groups panel
>
> To access the Volume Groups panel, in the navigation pane, click **Volume Groups**.

Procedure

1. In the Volume Groups panel, select volume groups to which you want to assign an owner, and then click **Owner**.
2. From the list of groups, select which group you want to assign as an owner of the selected volume groups, and then click **Assign**.

   > **Important** If any backup or restore job for a volume group is already in progress, or a scheduled backup task is in the queue, you cannot assign a new group to this volume group.

Depending on the needs of a specific data protection environment, you can at any time remove an owner from the volume groups by selecting the volume groups from which you want to remove the owner, and then clicking **Owner** followed by **Unassign**.

Setting ownership of buckets

Consideration

When changing ownership of buckets, you can choose whether you want data protected by specific owners to be kept or deleted. If you choose to keep data

protected by the specific owner, such buckets will be kept in HYCU with the Protected deleted status.

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

Procedure

1. In the Buckets panel, select buckets to which you want to assign an owner, and then click **Owner**.
2. From the list of groups, select which group you want to assign as an owner of the selected buckets, and then click **Assign**.

   > **Important** If any backup or restore job for a bucket is already in progress, or a scheduled backup task is in the queue, you cannot assign a new group to this bucket.

Depending on the needs of a specific data protection environment, you can at any time remove an owner from buckets by selecting the buckets from which you want to remove the owner, and then clicking **Owner** followed by **Unassign**.

Activating or deactivating users or self-service groups

Depending on the nature of your business, you can at any time enable or disable specific users or self-service groups from signing in to HYCU by activating or deactivating them. By activating or deactivating a self-service group, you enable or disable all users belonging to the specific self-service group from signing in to HYCU as members of that group.

Activating or deactivating a user

Procedure

1. In the Self-Service panel, click **Manage Users**.
2. From the list of all users, select the one whose status you want to change.
3. Depending on the status of the user, do one of the following:
   - If the status of the selected user is Inactive and you want to activate it, click **Activate**.

- If the status of the selected user is Active and you want to deactivate it, click ⊘ **Deactivate**.

Activating or deactivating a self-service group

Procedure

1. In the Self-Service panel, from the list of self-service groups, select the one whose status you want to change.
2. Depending on the status of the self-service group, do one of the following:
   - If the status of the selected self-service group is Inactive and you want to activate it, click ✓ **Activate**.
   - If the status of the selected self-service group is Active and you want to deactivate it, click ⊘ **Deactivate**.

> **Note** If a user is a member of several self-service groups and at least one of these groups has the Active status, the user is automatically switched to it. If there is more than one group with the Active status to which the user belongs, the user is automatically switched to the one that was created first.

Switching to another group

As a user you can belong to one or more groups and sign in to HYCU with all the permissions associated with the group to which you belong. If you are a member of more than one group, you can at any time switch to another group (provided that its status is Active) while being signed in to HYCU. This means that you can select any of the groups to which you belong and use it for a session.

Procedure

1. Click the group under which you are currently signed in to HYCU at the upper right of the screen.

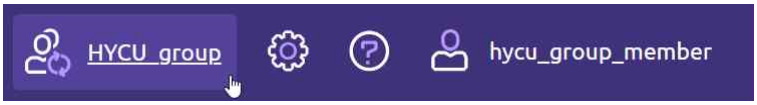

Figure 11-2: Example of a self-service group, HYCU_group, under which a user, hycu_group_member, is signed in to HYCU

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

The group under which you are currently signed in to HYCU has ✓ next to it.

3. Click **Switch**.

You are automatically switched to the group you selected.

Updating your user profile

You can edit your name, email address, preferred language, and authentication settings by using the Update Profile option.

Consideration

As a user with the Administrator role assigned, you can edit other users' information through the Self-Service panel. For details, see “Creating a user” on page 445.

> #### Accessing the Update Profile dialog box
>
> To access your Update Profile dialog box, click 👤 at the upper right of the screen and then select **Update Profile**.

Procedure

1. In the Name field, specify a new name.
2. In the Email field, enter the email address that you want to be associated with your user profile.
3. From the Language drop-down menu, select the preferred language.
4. *Optional*. Enable two-factor authentication by using the **Two-Factor Authentication** switch. Select the two-factor authentication method:
   - **Time-based one-time password**

     This option enables the use of a time-based one-time password (OTP) generated by an OTP application.
   - **FIDO**

     This option enables the use of an authenticator complying with FIDO protocols (FIDO authenticator).
5. Click **Save**.
6. *Only if you enabled two-factor authentication*. Perform the initial two-factor authentication setup:

- *For a time-based one-time password*: The Configure Two-Factor Authentication dialog box opens. Do the following:
  a. Scan the QR code with an appropriate OTP application (for example Google Authenticator on a mobile phone) or alternatively enter the OTP backup code in the application manually.
  b. In the Authentication code field, enter the generated six-digit code, and then click **Confirm** to finish the setup process.

  > Note If you do not set up a one-time password, the Configure Two-Factor Authentication dialog box opens during your next sign-in.

- *For FIDO*: The FIDO Authenticators dialog box opens. Do the following:
  a. Follow the wizard to register the authenticator (for example, a security key or Windows Hello with a fingerprint reader). The process depends on the type of authenticator you select and the operating system version.
  b. Enter a name for the authenticator, and then click **Register**.

  > Note If you do not complete the registration of at least one authenticator, you are prompted to register one at the first sign-in after two-factor authentication is enabled.
  > Later you can also add additional authenticators or revoke existing ones. For more details, see “Managing FIDO authenticators” on page 544.

Chapter 12

Administering

After you deploy HYCU, you can perform various administration tasks through the ⚙ **Administration** menu to customize HYCU for your data protection environment.

I want to...	Procedure
Add cloud accounts to HYCU.	“Adding a cloud account” on the next page
Configure encryption for targets.	“Configuring target encryption” on page 471
Integrate HYCU with identity providers.	“Integrating HYCU with identity providers” on page 471
Manage HYCU instances.	“Managing HYCU instances” on page 484
Set the iSCSI initiator secret.	“Setting the iSCSI initiator secret” on page 487
Obtain a permanent HYCU license.	“Licensing” on page 488
Configure log file settings to troubleshoot problems if HYCU does not perform as expected.	“Setting up logging” on page 493
Change network settings or enable network bandwidth throttling.	“Configuring your network” on page 496
Set up HYCU to be monitored with an external monitoring tool.	“Enabling external monitoring of HYCU” on page 500
Set data retention for restore point tiers.	“Managing data retention” on page 501
Set power options.	“Setting power options” on page 509

I want to...	Procedure
Securely store, access, and manage my credentials (secrets) by employing the Conjur secrets management solution.	“Managing secrets” on page 509
Configure email delivery settings in HYCU.	“Configuring email delivery” on page 512
Upgrade HYCU to a new available version.	“Upgrading HYCU” on page 515
Apply a HYCU update.	“Applying HYCU updates” on page 519
Configure the SSL certificate.	“Configuring SSL certificates” on page 524
Share telemetry diagnostic data with HYCU.	“Sharing telemetry data with HYCU” on page 530

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 531.

Adding a cloud account

You must add one or more cloud accounts to HYCU before you can perform data protection tasks. The type of cloud account that you must add to HYCU depends on which data protection tasks you want to perform and which HYCU backup and recovery service you use to protect your data, and whether you want to present data protection information in HYCU R-Cloud. For details on the different types of cloud accounts, see the following sections:

- “Adding a cloud provider account” below
- “Adding a HYCU account” on page 470

Adding a cloud provider account

Before performing certain data protection tasks, you must add one or more cloud provider accounts to HYCU.

Consideration

Restoring virtual machines to a different cloud source, migrating virtual machines to cloud, performing disaster recovery of data to cloud, and monitoring cloud data protection environments are supported only if you own a premium tier Platform license.

I want to...	Cloud provider account	Instructions
• Migrate data protected with HYCU to AWS. • Migrate AWS data protected with HYCU R-Cloud to the on-premises environment. • Perform disaster recovery of data to AWS.	AWS user account	“Adding an AWS user account” on the next page
• Add an AWS GovCloud (US) region to HYCU to be able to protect AWS GovCloud (US) virtual machines and applications.	AWS GovCloud (US) account	“Adding an AWS GovCloud (US) account” on page 465
• Store data to a Google Cloud target. • Migrate data protected with HYCU to Google Cloud. • Migrate Google Cloud data protected with HYCU R-Cloud to the on-premises environment. • Perform disaster recovery of data to Google Cloud.	Google Cloud service account	“Adding a Google Cloud service account” on page 466
• Add an Azure subscription to HYCU. • Migrate data protected with HYCU to Azure. • Migrate data protected with HYCU R-Cloud or HYCU for Azure to the on-premises environment. • Perform disaster recovery of data	Azure service principal	“Adding an Azure service principal” on page 468

I want to...	Cloud provider account	Instructions
to Azure. • Monitor my HYCU for Azure data protection environment in HYCU Manager. • Configure HYCU email delivery for Microsoft 365.[a]		
• Add an Azure Government subscription to HYCU. • Migrate data protected with HYCU to Azure Government. • Perform disaster recovery of data to Azure Government. • Configure HYCU email delivery for Microsoft 365.[a]	Azure Government service principal	"Adding an Azure Government service principal" on page 469

[a] Depending on your data protection environment, you must add either an Azure or Azure Government service principal to HYCU, not both.

Adding an AWS user account

Prerequisite

A user account must be created in AWS and it must have permissions to perform the following actions in the S3 service: `ListBucket`, `CreateBucket`, `DeleteBucket`, `GetObject`, `PutObject`, `DeleteObject`, and `PutBucketPublicAccessBlock`. In addition, you must set the Resources value to **All resources** for these actions. For more information about S3 permissions, see AWS documentation.

> **Accessing the Cloud Accounts dialog box**
>
> To access the Cloud Accounts dialog box, click **Administration**, and then select **Cloud Accounts**.

Procedure

1. In the Cloud Accounts dialog box, click the **Cloud Provider Accounts** tab, and then click **Add**. The Select Cloud dialog box opens.

2. Select **Add AWS user account**, and then click **Next**. The AWS Authentication dialog box opens.
3. In the Name field, enter a name for your AWS user account.
4. In the Access key ID field, enter the access key ID of your AWS user account.
5. In the Secret access key field, enter the secret access key of your AWS user account.

   > Note The access key ID and the secret access key are used to authenticate AWS API service calls.

6. Click **Save**.

You can later edit any of the existing cloud accounts (click **Edit** and make the required modifications) or remove the ones that you do not need anymore (click **Remove**).

Adding an AWS GovCloud (US) account

Prerequisite

Your AWS GovCloud (US) account must have permissions to perform the following actions in the S3 service: `ListBucket`, `CreateBucket`, `DeleteBucket`, `GetObject`, `PutObject`, `DeleteObject`, and `PutBucketPublicAccessBlock`. In addition, you must set the Resources value to **All resources** for these actions. For more information about S3 permissions, see AWS documentation.

> Accessing the Cloud Accounts dialog box
>
> To access the Cloud Accounts dialog box, click **Administration**, and then select **Cloud Accounts**.

Procedure

1. In the Cloud Accounts dialog box, click the **Cloud Provider Accounts** tab, and then click **Add**.
2. Select **Add AWS GovCloud (US) account**, and then click **Next**.
3. In the Name field, enter a name for your AWS GovCloud (US) account.
4. In the Access key ID field, enter the access key ID of your AWS GovCloud (US) account.
5. In the Secret access key field, enter the secret access key of your AWS GovCloud (US) account.

> **Note** The access key ID and the secret access key are used to authenticate AWS GovCloud (US) API service calls.

6. Click **Save**.

You can later edit any of the existing cloud accounts (click **Edit** and make the required modifications) or remove the ones that you do not need anymore (click **Remove**).

Adding a Google Cloud service account

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

> **Important** You must always add a dedicated service account that you have created yourself to HYCU.

I want to...	Service account to add
Store data to a Google Cloud target.	An account that has access to the buckets where you want to store your backup data.
Migrate data protected with HYCU R-Cloud from Google Cloud to the on-premises environment.	An account that is added to HYCU R-Cloud and has the Storage Admin role assigned on the projects containing the instances.
Migrate data protected with HYCU from the on-premises environment to Google Cloud.	An account that is added to HYCU R-Cloud, and has the Compute Admin, Storage Admin, Service Account User, and Service Account Token Creator roles assigned on the projects where you want to migrate your virtual machines.
Perform disaster recovery of data to Google Cloud in the event of a disaster.	An account that is added to HYCU R-Cloud, and has the Compute Admin, Storage Admin, Service Account User, and Service Account Token Creator roles assigned on the projects where you want to perform disaster recovery.

Prerequisites

- The service account must be configured in Google Cloud.
- The following APIs must be enabled on the Google Cloud project on which the service account was created:
  - Cloud Resource Manager API
  - Compute Engine API
  - Cloud Storage API
  - Identity and Access Management API

  For instructions on how to enable the listed APIs, see Google Cloud documentation.
- You must have access to a valid JSON file that stores the service account information, including its private key.

> #### Accessing the Cloud Accounts dialog box
>
> To access the Cloud Accounts dialog box, click **Administration**, and then select **Cloud Accounts**.

Procedure

1. In the Cloud Accounts dialog box, click the **Cloud Provider Accounts** tab, and then click **Add**. The Select Cloud dialog box opens.
2. Select **Add Google Cloud service account**, and then click **Next**. The Google Cloud Authentication dialog box opens.
3. Browse for the JSON file with the service account information. In the Service account authentication field, the file name is displayed.

   > **Note** *Only if you are signed in to HYCU as a self-service group administrator.* If you use Conjur for managing your HYCU secrets, you can enable the **Retrieve values from secrets manager** switch if you want to provide the secret instead of browsing for the file. For details on managing secrets, see "Managing secrets" on page 509.

4. In the Name field, enter a name for your service account.
5. Click **Upload** or **Save**.

You can later edit any of the existing cloud accounts (click **Edit** and make the required modifications) or remove the ones that you do not need anymore (click **Remove**). Keep in mind that you cannot remove a cloud account in the following cases:

- A Google Cloud target uses this account.
- A protection set that is monitored in HYCU Manager uses this account.

Adding an Azure service principal

Prerequisites

- A service principal must be created in Azure.
- The service principal must be assigned the Contributor role at the subscription level.
- *Only if you plan to use HYCU R-Cloud or HYCU for Azure to migrate data from or to the on-premises environment, or to perform a disaster recovery:*
  - The service principal must be added to HYCU R-Cloud or HYCU for Azure. For details, see HYCU R-Cloud or HYCU for Azure documentation.

    > ⓘ Important You must always add a dedicated service principal that you created yourself to HYCU and not use the default one that HYCU R-Cloud or HYCU for Azure automatically creates for you when you start using the service.

  - You need to assign the Storage Blob Data Contributor role assigned at the subscription, resource group, or storage account level.
- *Only if you plan to monitor your HYCU for Azure environment in HYCU Manager*:
  - The service principal must be added to HYCU for Azure. For details, see HYCU for Azure documentation.

    > ⓘ Important You must always add a dedicated service principal that you created yourself to HYCU and not use the default one that HYCU for Azure automatically creates for you when you start using the service.

> Accessing the Cloud Accounts dialog box
>
> To access the Cloud Accounts dialog box, click ⚙ **Administration**, and then select **Cloud Accounts**.

Procedure

1. In the Cloud Accounts dialog box, click the **Cloud Provider Accounts** tab, and then click ✚ **Add**. The Select Cloud dialog box opens.
2. Select **Add Azure Service Principal**, and then click **Next**. The Azure Authentication dialog box opens.

3. In the Name field, enter the name for your service principal.
4. In the Tenant ID field, enter your tenant ID.
5. In the Application ID field, enter the HYCU application registration ID from Microsoft Entra ID.
6. In the Client secret value field, enter the client secret's value.
7. Click **Save**.

You can later edit any of the existing service principals (click **Edit** and make the required modifications) or remove the ones that you do not need anymore (click **Remove**). If you are monitoring your HYCU for Azure data protection environment in HYCU Manager, keep in mind that you cannot remove a service principal if the protection set that is monitored in HYCU Manager uses this cloud account.

Adding an Azure Government service principal

Prerequisites

- The service principal must be created in Azure Government.
- The service principal must be assigned the Contributor role at the subscription level.

> #### Accessing the Cloud Accounts dialog box
>
> To access the Cloud Accounts dialog box, click **Administration**, and then select **Cloud Accounts**.

Procedure

1. In the Cloud Accounts dialog box, click the **Cloud Provider Accounts** tab, and then click **Add**. The Select Cloud dialog box opens.
2. Select **Add Azure Government service principal**, and then click **Next**. The Azure Government Authentication dialog box opens.
3. In the Name field, enter the name for your service principal.
4. In the Tenant ID field, enter your tenant ID.
5. In the Application ID field, enter the HYCU application registration ID from Microsoft Entra ID.
6. In the Client secret value field, enter the client secret's value.
7. Click **Save**.

You can later edit any of the existing service principals (click **Edit** and make the required modifications) or remove the ones that you do not need anymore (click **Remove**).

Adding a HYCU account

You must add one or more HYCU accounts to HYCU if you want to use HYCU R-Cloud to:

- Perform any of the following data protection tasks:
  - Migrating data protected with HYCU from the on-premises environment to cloud.
  - Migrating data protected with HYCU R-Cloud from cloud to the on-premises environment.
  - Performing disaster recovery of data to cloud.
- Present data protection information from the HYCU backup controller in HYCU R-Cloud by enabling publishing to HYCU R-Cloud.

Consideration

If you use HYCU for Azure to protect your data, adding a HYCU account is not required.

> Accessing the Cloud Accounts dialog box
>
> To access the Cloud Accounts dialog box, click **Administration**, and then select **Cloud Accounts**.

Procedure

1. In the Cloud Accounts dialog box, click the **HYCU Accounts** tab, and then click **Add**.
2. In the Name field, enter a name for your HYCU account.
3. In the HYCU account ID field, enter the HYCU account ID that you received when you subscribed to HYCU R-Cloud.
4. In the Username and Password fields, enter the credentials of the user account that you use for accessing HYCU R-Cloud.
5. *Only if you want to allow HYCU R-Cloud to collect data protection information from the HYCU backup controller and to present it in HYCU R-Cloud dashboard.* Enable the **Use for publishing to R-Cloud** switch.
6. Click **Save**.

You can later edit any of the existing HYCU accounts (click **Edit** and make the required modifications) or remove the ones that you do not need anymore (click **Remove**).

Configuring target encryption

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

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

Exporting an encryption key

Procedure

1. In the Encryption dialog box, click **Export**.
2. Save the exported file to a safe location.

Importing an encryption key

Procedure

1. In the Encryption dialog box, click **Import**.
2. In the Import dialog box, browse for the file containing the encryption key, and then click **Import**.

Integrating HYCU with identity providers

You can integrate HYCU with the Active Directory identity provider, identity providers that use LDAP, and identity providers that support the OpenID Connect authentication protocol, such as Google, Keycloak, Microsoft, Okta, and Active Directory Federation Services. This gives users the possibility to

securely sign in to HYCU by using these identity providers, without the need to maintain dedicated credentials for HYCU.

When integrating HYCU with an identity provider, you must complete the following tasks:

Task	Instructions
1. Add an identity provider to HYCU to be able to authenticate users.	Follow the procedure described in "Adding an identity provider to HYCU" below.
2. Create a user for whom you want to enable signing in by using the identity provider, and then add this user to a user group.	Follow the procedures described in "Creating a user" on page 445 and "Adding a user to a group" on page 450.

Adding an identity provider to HYCU

Prerequisites

- *Only when adding identity providers that support the OpenID Connect authentication protocol.* HYCU must be registered as a web application within the identity provider that you plan to add to HYCU. When registering HYCU, depending on the identity provider, make sure the following is done:
  - *Active Directory Federation Services*:
    - In Active Directory Federation Services, you must select **Server application accessing a web API** as the client-server application, and **openid** and **allatclaims** when configuring application permissions. Also, make sure that the outgoing claim type to which you map the ObjectGUID attribute is ObjectGUID when configuring rules.
    - *Only if you are adding identity provider groups.* In Active Directory Federation Services, create an issuance transform rule, and then set the groups claim.
  - *Keycloak*:
    - *Only if you are adding identity provider groups.* In Keycloak, when configuring the client scopes, add a mapper of the type Group Membership, with the token claim named `group_mem`. Make sure to disable the **Full group path** switch.

  - *Microsoft*:
    - In Azure, HYCU must be given access permissions to the following Azure API: Microsoft Graph with delegated permissions for User.Read.
    - *Only if you are adding identity provider groups.* In Azure, under Token configuration, add a groups claim.
  - *Okta*:
    - In Okta, you must select **Authorization Code** under Client acting on behalf of a user as the grant type.
    - *Only if you are adding identity provider groups.* In Okta, add a new scope named `groups`, and then add a new claim.

  For instructions on how to register an application, see the respective identity provider documentation.
- *Only if you plan to use LDAPS with Active Directory or with identity providers that use LDAP:* LDAPS authentication must be set up. For details, see "Setting up LDAPS authentication" on page 541.

Considerations

- To increase the security of user accounts further, you can also configure multi-factor authentication within the identity providers. For instructions on how to do this, see the respective identity provider documentation.
- If you use Active Directory as your authentication source in HYCU, you can also enable certificate authentication to allow users to sign in to the HYCU web user interface with a client certificate or a smart card. For instructions, see "Enabling certificate authentication" on page 483.

> #### Accessing the Identity Providers dialog box
>
> To access the Identity Providers dialog box, click ⚙ **Administration**, and then select **Identity Providers**.

Procedure

1. In the Identity Providers dialog box, click ✚ **Add**.
2. In the Name field, enter a name for the identity provider.
3. From the Type drop-down menu, select one of the following types of identity providers, and then follow the instructions:

Identity provider type	Instructions
Active Directory	a. In the Domain field, enter the FQDN or domain alias name of the Active Directory. If you plan to use AD groups, it is mandatory to enter the FQDN. For example, if you enter `mycompany.com` as the FQDN and `mc` as the alias domain name, the user will be able to sign in to HYCU with `<Username>@mycompany.com` or `mc\<Username>`. **Note** You can enter more than one FQDN or domain alias name. In this case, press the Spacebar after entering each one. b. In the Provider URL field, enter the URL of the corresponding LDAP server in one of the following formats: • `ldap://<LDAPServerHostnameorIPAddress>:<Port>` When using the LDAP protocol, the default port is 389. Entering the port is optional if the default value is used. • *Only if LDAPS authentication is set up.* `ldaps://<LDAPServerHostname>:<Port>` **Important** Make sure that the LDAP server host name matches the DNS entry specified in the Subject Alternative Name (SAN) extension of the LDAP server's certificate. Otherwise, connection to the LDAP server will fail. When using the LDAPS protocol, the default port is 636. Entering the port is optional if the default value is used. **Note** You can enter more than one URL. In this case, press the Spacebar after entering each one. c. *Only if LDAPS authentication is set up.* Click

Identity provider type	Instructions
	SSL Certificates to view the information about your SSL certificates. If you want to import a new certificate, see "Importing a custom certificate" on page 527 for instructions. d. *Only if you plan to enable certificate authentication.* Enable the **Use service account** switch, and then enter the user name and password of the service account that HYCU will use to sign in to the Active Directory and authorize users. e. *Optional.* If you want to check whether the identity provider is configured correctly and the users will be able to sign in to HYCU, do the following: ii. Click **Test Active Directory Configuration**. iii. Depending on whether or not you enabled the Use service account option, do one of the following: • If you enabled the Use service account option: Enter the user name of a user that will be signing in to HYCU with a client certificate or a smart card, and then click **Run Test**. • If you did not enable the Use service account option: Enter the user name and the password of a user that will be signing in to HYCU, and then click **Run Test**. Note The user credentials that you provide will only be used to test the identity provider configuration and will not be saved. If the test is successfully completed, the list of groups that the user is a member of is displayed. Review the list to make sure that the correct groups are displayed. Any missing groups could indicate an issue with the identity provider configuration in

Identity provider type	Instructions
	HYCU or on the Active Directory server. iv. Click **Close**.

Active Directory Federation Services	a. In the Client ID field, enter the application ID that is generated by the identity provider.
	b. In the Client Secret field, enter the application secret that is associated with the client ID and generated by the identity provider.
	c. *Only if PKCE is configured for the identity provider.* From the Proof Key for Code Exchange (PKCE) drop-down menu, select the required code challenge method:
	• **S256**
	• **Plain**
	d. In the Redirect URI field, enter the URL to which the user will be redirected after authentication. The format is as follows:
	`https://<ServerName>:8443`
	In this instance, `<ServerName>` is the fully qualified domain name of the HYCU server.
	For example:
	`https://hycu.example.com:8443`
	e. *Optional.* In the Well-Known Configuration URL field, enter the well-known configuration URL to enable HYCU to automatically retrieve the necessary configuration details for your identity provider. Click **Fetch Data**.
	Note After HYCU retrieves the data, the remaining configuration fields are populated automatically.
	f. In the Issuer field, enter the URL of the issuer of the

Identity provider type	Instructions
	identity provider.

Google	a. In the Client ID field, enter the application ID that is generated by the identity provider.
	b. In the Client Secret field, enter the application secret that is associated with the client ID and generated by the identity provider.
	c. *Only if PKCE is configured for the identity provider.* From the Proof Key for Code Exchange (PKCE) drop-down menu, select the required code challenge method:
	• **S256**
	• **Plain**
	d. In the Redirect URI field, enter the URL to which the user will be redirected after authentication. The format is as follows:
	`https://<ServerName>:8443`
	In this instance, `<ServerName>` is the fully qualified domain name of the HYCU server.
	For example:
	`https://hycu.example.com:8443`

Keycloak	a. In the Client ID field, enter the application ID that is generated by the identity provider. b. In the Client Secret field, enter the application secret that is associated with the client ID and generated by the identity provider. c. *Only if PKCE is configured for the identity provider.* From the Proof Key for Code Exchange (PKCE) drop-down menu, select the required code challenge method: • **S256** • **Plain**

Identity provider type	Instructions

	d. In the Redirect URI field, enter the URL to which the user will be redirected after authentication. The format is as follows:
	`https://<ServerName>:8443`
	In this instance, *<ServerName>* is the fully qualified domain name of the HYCU server.
	For example:
	`https://hycu.example.com:8443`
	e. *Optional.* In the Well-Known Configuration URL field, enter the well-known configuration URL to enable HYCU to automatically retrieve the necessary configuration details for your identity provider. Click **Fetch Data**.
	Note After HYCU retrieves the data, the remaining configuration fields are populated automatically.
	f. In the Issuer field, enter the URL of the issuer of the identity provider.

Microsoft	a. In the Client ID field, enter the application ID that is generated by the identity provider. b. In the Client Secret field, enter the application secret that is associated with the client ID and generated by the identity provider. c. *Only if PKCE is configured for the identity provider.* From the Proof Key for Code Exchange (PKCE) drop-down menu, select the required code challenge method: • **S256** • **Plain** d. In the Redirect URI field, enter the URL to which the

Identity provider type	Instructions

	user will be redirected after authentication. The format is as follows:
	`https://<ServerName>:8443`
	In this instance, *<ServerName>* is the fully qualified domain name of the HYCU server.
	For example:
	`https://hycu.example.com:8443`

Okta	a. In the Client ID field, enter the application ID that is generated by the identity provider.
	b. In the Client Secret field, enter the application secret that is associated with the client ID and generated by the identity provider.
	c. In the Authorization Server field, enter the ID of the authorization server.
	d. *Only if PKCE is configured for the identity provider.* From the Proof Key for Code Exchange (PKCE) drop-down menu, select the required code challenge method:
	• **S256**
	• **Plain**
	e. In the Redirect URI field, enter the URL to which the user will be redirected after authentication. The format is as follows:
	`https://<ServerName>:8443`
	In this instance, *<ServerName>* is the fully qualified domain name of the HYCU server.
	For example:
	`https://hycu.example.com:8443`

Identity provider type	Instructions
	f. *Optional.* In the Well-Known Configuration URL field, enter the well-known configuration URL to enable HYCU to automatically retrieve the necessary configuration details for your identity provider. Click **Fetch Data**. > Note After HYCU retrieves the data, the remaining configuration fields are populated automatically. g. In the Issuer field, enter the URL of the issuer of the identity provider.

OpenID Connect IdP	a. In the Client ID field, enter the application ID that is generated by the identity provider.
	b. In the Client Secret field, enter the application secret that is associated with the client ID and generated by the identity provider.
	c. *Only if PKCE is configured for the identity provider.* From the Proof Key for Code Exchange (PKCE) drop-down menu, select the required code challenge method:
	• **S256**
	• **Plain**
	d. In the Redirect URI field, enter the URL to which the user will be redirected after authentication. The format is as follows:
	`https://<ServerName>:8443`
	In this instance, `<ServerName>` is the fully qualified domain name of the HYCU server.
	For example:
	`https://hycu.example.com:8443`
	e. *Optional.* In the Well-Known Configuration URL field,

Identity provider type	Instructions
	enter the well-known configuration URL to enable HYCU to automatically retrieve the necessary configuration details for your identity provider. Click **Fetch Data**. > **Note** After HYCU retrieves the data, the remaining configuration fields are populated automatically. f. In the Issuer field, enter the URL of the issuer of the identity provider. g. In the Authorization Endpoint field, enter the authorization endpoint of the identity provider. h. In the Token Endpoint field, enter the token endpoint of the identity provider. i. In the JWKS Endpoint field, enter the JSON web key set endpoint of the identity provider. j. In the UserInfo Endpoint field, enter the UserInfo endpoint of the identity provider.
LDAP	a. In the Domain field, enter a domain for your identity provider. > **Important** The domain that you configure must be provided by users when signing in to HYCU. For details about signing in to HYCU, see "Signing in to HYCU" on page 68. b. In the Provider URL field, enter the URL of the LDAP server. > **Note** You can enter more than one URL. In this case, press the Spacebar after entering each one. c. *Only if LDAPS authentication is set up.* Click **SSL Certificates** to view the information about your SSL certificates. If you want to import a new certificate, see "Importing a custom certificate" on page 527 for

Identity provider type	Instructions
	instructions. d. In the User Object Class field, enter the LDAP object class that represents the user. e. In the User Search Base field, enter the distinguished name (DN) in the directory tree from which HYCU will search for users. f. In the Username Attribute field, enter the attribute that HYCU will use to uniquely identify the user. g. *Only if you plan to authenticate users by using LDAP groups.* Make sure the **Group configuration** switch is enabled, and then do the following: i. In the Group Object Class field, enter the LDAP object class that represents the group. ii. In the Group Search Base field, enter the DN in the directory tree from which HYCU should search for groups. iii. In the Group Member Attribute field, enter the attribute that lists members of a group. iv. In the Group Member Attribute Value field, enter the user attribute that is used as the value that determines group membership. h. *Optional.* If you want to check whether the identity provider is configured correctly and the users will be able to sign in to HYCU, do the following: i. Click **Test LDAP Configuration**. ii. Enter the user name and the password of a user that will be signing in to HYCU, and then click **Run Test**. Note The user credentials that you provide will only be used to test the identity provider configuration and will not be saved. iii. *Only if group configuration is enabled.* If the test is

Identity provider type	Instructions
	successfully completed, the list of groups that the user is a member of is displayed. Review the list to make sure that the correct groups are displayed. Any missing groups could indicate an issue with the identity provider configuration in HYCU or on the LDAP server. iv. Click **Close**.

4. Click **Save**.

You can later do the following:

- Edit information about any of the existing identity providers by clicking **Edit** and making the required modifications.

  > **Note** The Redirect URI field shows to which URL the user will be redirected after authentication (for example, `https://hycu.example.com:8443`). The prepopulated host name is the host name of the HYCU backup controller to which you are authenticating user access.

- Remove any of the existing identity providers by clicking **Remove**.

Enabling certificate authentication

By enabling certificate authentication, you allow Active Directory users to sign in to the HYCU web interface by using a client certificate or a smart card, without having to enter a password.

Prerequisites

- At least one Active Directory with a configured service account must be added to HYCU.
- A CA-signed certificate must be imported to HYCU. For details on how to do this, see “Importing a custom certificate” on page 527.

Limitation

If certificate authentication is enabled, setting up two-factor authentication for AD users is not supported.

Consideration

When you enable or disable certificate authentication, all affected users that are signed in the HYCU web user interface will lose their connections and will be required to sign in again.

Procedure

1. In the Identity Providers dialog box, use the **Enable certificate authentication** switch if you want to enable certificate authentication.
2. From the CA Certificate drop-down menu, select the CA-signed certificate for verifying the client certificate.

   > Tip By clicking **Manage**, you are automatically directed to the SSL Certificates dialog box where you can manage your certificates.

Managing HYCU instances

All existing HYCU instances in your data protection environment are listed in the HYCU Instances dialog box. Besides viewing all the existing HYCU instances, you can use this dialog box also to create new HYCU instances, view information about each HYCU instance, delete HYCU instances, and apply updates to HYCU instances.

For details on HYCU instances, see "HYCU instances" on page 25. For instructions on how to apply an update to HYCU instance, see "Applying an update to a HYCU instance" on page 522.

> Accessing the HYCU Instances dialog box
>
> To access the HYCU Instances dialog box, click ⚙ **Administration**, and then select **HYCU Instances**.

Creating a HYCU instance by using the HYCU web user interface

You can create a HYCU instance by using the HYCU web user interface as an alternative to creating it by deploying the HYCU virtual appliance in the HYCU Instance mode.

Prerequisites

- The Nutanix cluster, the vCenter Server, or the XenServer where you want to create your HYCU instance must be added as a source to HYCU. For instructions, see "Adding sources" on page 73.
- *For Nutanix AHV clusters:* The HYCU virtual appliance image must be present on the Nutanix cluster in the following format:

  `hycu-<Version>-<Revision>`

  For example, `hycu-5.3.0-3634`.
- *For Nutanix ESXi clusters or vSphere environments:*
  - A user with specific privileges for vCenter Servers must be specified. For details on which privileges must be assigned to a vSphere user, see "Assigning privileges to a vSphere user" on page 552.
  - The HYCU OVF package must be imported to the vCenter Server content library and its format must be as follows:

    `hycu-<Version>-<Revision>`

    For example, `hycu-5.3.0-3634`.
- *For XenServer environments:* The HYCU virtual machine disk image must be imported to the XenServer. To do so, follow these steps:
  1. Create a virtual machine as described in "Creating a virtual machine for HYCU deployment in a XenServer environment" on page 53. Make sure that you specify the virtual machine name in the following format:

     `hycu-<Version>-<Revision>`

     For example, `hycu-5.3.0-3634`.
  2. Convert the created virtual machine to a template. For instructions, see XenServer documentation.

Limitation

For Azure Local environments and Hyper-V clusters: You cannot create a HYCU instance by using the HYCU web user interface. The only supported method to create a HYCU instance is by deploying the HYCU virtual appliance in the HYCU Instance mode. For instructions, see "Deploying HYCU to an Azure Local environment" on page 55 or "Deploying HYCU to a Hyper-V cluster" on page 59.

Procedure

1. In the HYCU Instances dialog box, click ✚ **New**.
2. In the General section, enter a name for the HYCU instance.

3. In the Network configuration section, do the following:
   a. Enter a host name for the HYCU instance.

      > **Important** Make sure that you enter a unique host name for each HYCU instance that you create and follow these rules:
      > - The host name contains only letters, numbers, hyphens (-), and periods. The maximum number of characters is 253 and at least one of the characters is a letter.
      > - The maximum number of characters in each host name segment is 63. A host name segment cannot begin or end with a hyphen.
      > - The top-level domain cannot begin or end with a number.

   b. Use the **DHCP** switch if you want a dynamic IP address to be assigned to the HYCU instance.

      Otherwise, specify the IP address, the netmask, and the gateway. You can click **Auto** for the netmask and the gateway if you want HYCU to automatically set them for you.
4. In the Deployment section, do the following:
   a. From the Destination drop-down menu, select the source where you want to create your HYCU instance.
   b. From the Storage container drop-down menu, select the storage container where you want to create your HYCU instance.

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

   c. From the Network drop-down menu, select a VLAN.
5. Click **Save**.

Viewing HYCU instance information

You can view the following information about each HYCU instance:

HYCU instance information	Description
VM name	Name of the HYCU instance, if known.
Hostname	Host name of the HYCU instance.
Source	Nutanix cluster, the vCenter Server, or the XenServer on which the HYCU instance resides.

HYCU instance information	Description
Status	Shows if the HYCU instance is up and running, and communicating with the HYCU backup controller.
Version	Version of the HYCU instance (for example, 5.3.0-3634).
IP address	IP address currently assigned to the HYCU instance.

Deleting a HYCU instance

Procedure

1. In the HYCU Instances dialog box, from the list of HYCU instances, select the one that you want to delete, and then click **Delete**.

   > **Important** The selected HYCU instance will be deleted from both HYCU and the source.

2. In the Delete HYCU Instance dialog box, click **Delete** to confirm that you want to delete the selected HYCU instance.

Setting the iSCSI initiator secret

During the HYCU deployment, the HYCU iSCSI client, referred to as the iSCSI initiator, is set up so that HYCU can use iSCSI targets for storing data.

If you want to configure mutual CHAP authentication between the iSCSI initiator and the iSCSI target, you must specify the iSCSI initiator secret (the security key). For details on how to enable mutual authentication, see "Setting up targets" on page 100.

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

Procedure

1. In the iSCSI Initiator dialog box, enable the **Specify iSCSI initiator secret** switch.

2. In the Secret field, enter the iSCSI initiator secret.
3. Click **Save**.

You can at any time remove the iSCSI initiator secret by disabling the **Specify iSCSI initiator secret** switch, and then clicking **Save**.

Licensing

After you deploy the HYCU virtual appliance, you can start using HYCU immediately with a trial license. This license expires automatically after 30 days and cannot be reused. Therefore, make sure to obtain a valid license within this 30-day period.

The HYCU license is linked to the HYCU backup controller and you can decide on the license type or a combination of license types that best suits your data protection environment.

Platform licenses

License type	Description
• Socket-based • Core-based	Licenses are based on the number of CPU sockets or CPU cores on the sources (Nutanix clusters, vCenter Servers, XenServers, Azure Local environments, Hyper-V clusters, or servers) that you plan to protect by using HYCU. Socket-based and core-based licenses cannot be used to protect cloud environments (AWS GovCloud (US) regions, Azure subscriptions, or Azure Government subscriptions). **Note** Core-based licenses and socket-based licenses cannot be used in combination.
VM-based	Licenses are based on the number of virtual machines on all sources (Nutanix clusters, vCenter Servers, XenServers, Azure Local environments, Hyper-V clusters, or cloud environments) and servers that you plan to protect by using HYCU.

> **Important** For each Platform license, the standard and premium tiers are available. A standard tier Platform license is sufficient for most data protection scenarios. A premium tier Platform license is required for the following data protection scenarios:
> - Protecting NDB-managed database server virtual machines
> - Restoring a virtual machine to a different cloud source
> - Monitoring HYCU for Azure data protection environments
> - Migrating data to cloud
> - Performing disaster recovery of data to cloud
> - Performing malware scans

File and object server licenses

License type	Description
• Socket-based • Core-based	Licenses are based on the number of CPU sockets or CPU cores on all Nutanix clusters where the Nutanix Files servers that you plan to protect by using HYCU reside. Socket-based and core-based licenses are reserved exclusively for Nutanix Files servers. **Note** Core-based licenses and socket-based licenses cannot be used in combination.
Capacity-based	Licenses are based on the capacity of file shares and buckets, which is calculated automatically as an overall size (in TiB) of all protected file shares and buckets.

Pay-as-you-go license

The license is based on the capacity of the protected resources on all sources (Nutanix clusters, vCenter Servers, XenServers, Azure Local environments, Hyper-V clusters, or cloud environments) and servers. The license use is reported in real time by using telemetry.

Considerations

- When verifying that your license is valid, HYCU takes into account only the sources containing the entities with the Protected or Protected deleted

status.

- The protection of the HYCU backup controller does not require any license.
- Depending on your data protection environment needs, you can combine Platform licenses with File and object server licenses. However, the Pay-as-you-go license can only be used alone.
- When a Pay-as-you-go license is applied to HYCU, sharing telemetry data with HYCU is enabled by default and cannot be disabled.
- *For Nutanix Community Edition (CE) environment:* No HYCU licenses are required.

Procedure

1. Purchase the required number of HYCU licenses. To discuss the options, contact your HYCU Sales representative.
2. Create a license request. For details, see "Creating a license request" below.
3. Request and obtain licenses from the web licensing portal. For details, see "Requesting and retrieving licenses" on the next page.
4. Activate the licenses to start using HYCU. For details, see "Activating licenses" on page 492.

> 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 purchased the required number of HYCU licenses.
- You added sources that you want to protect to the data protection environment. For instructions, see "Adding sources" on page 73.

Procedure

1. In the Licensing dialog box, click **Download Request**.
2. Save the license request file to a temporary location.

> **Example**
> `license.req` file:
>
> ```
> CN myCompany
> ND C0F90A56-3FCC-4437-A49C-EFBA9BD8FC0F
> VER V2
> PID nutanixbackup
> NRP 3
> CORES 112
> QTY 127
> AFS 12
> AFSCORES 112
> AFSCAP 1
> CAPSTD 1
> NRPALL 12
> CORESALL 112
> QTYALL 167
> HYCUVER 5.1.0-2014
> ZD ...
> HSUD 970798D0272032CCB142F6F34...
> NEXT NODE
> ```

Requesting and retrieving licenses

After you create a license request file, you can obtain the licenses from the licensing portal.

Procedure

1. Connect to the web licensing portal at:
   https://licensing.hycu.com/
2. If you already have a licensing portal account, enter your email address or user name and the password, and then click **Sign In**. Otherwise, contact HYCU Support.
3. Click the **Activate or Renew Licenses** link.
4. Click **Choose a File** to browse for the license request file, and then click **Request License**.
5. Click **Next**.
6. From the Expiration Date drop-down menu, select the expiration date for your license.

> **Note** The available expiration dates are based on the end dates of the active purchase orders. For details, see the web licensing portal Overview page.

7. *Only if you purchased Platform licenses or File and object server licenses.* Select the required number of licenses.
8. Click **Activate**.

   Within a few minutes, you should receive an email with a license file `license.dat` attached.

   Example

   `license.dat` file:

   ```
   CN myCompany
   ND C0F90A56-3FCC-4437-A49C-EFBA9BD8FC0F
   VER V2
   PID nutanixbackup
   EXP 23.10.2024
   CORES 128
   AFSCAP 23
   AFSCORES 128
   LK
   302C0214077EC128B5EF4961CF64DB...
   21FE123412341234
   NEXT NODE
   ```

9. Save the license file locally.

Activating licenses

After you submit your license request for the HYCU licenses to the web licensing portal, you get an email with a product license file attached.

Procedure

1. In the Licensing dialog box, click **Upload License**.
2. Browse for the license file that you received by email, and then click **Upload**.

After the licenses are activated, the information related to licensing is updated. On the Licensing tabs, you can view the following information related to licensing:

- Status
- License type
- HYCU backup controller ID
- License expiration date
- Discovered, protected, and licensed number of virtual machines and servers
- Discovered, protected, and licensed number of sockets or cores for virtual machines and servers
- Protected and licensed file server and object server capacity

> **Note** You can always add new licenses for your growing environment. Contact your HYCU Sales representative.

Applying Platform licenses manually

If you purchase a combination of different Platform license types (socket-based and VM-based, or core-based and VM-based), HYCU automatically selects how the individual sources in your data protection environment are licensed. However, you can also manually select which license type is applied to which source.

> **Note** File and object server licenses are always automatically applied by HYCU.

Procedure

1. In the Licensing dialog box, on the Platform tab, disable the **Optimize auto selection** switch.
2. Select the preferred license type for each source.
3. Click **Apply**.

Setting up logging

You can set up logging to log information at various levels to help you analyze and troubleshoot the entire HYCU operation and diagnose issues with backup and restore performance.

Prerequisite

For sending log files to HYCU Support: Sharing telemetry data with HYCU must be enabled. For instructions, see "Sharing telemetry data with HYCU" on page 530.

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

In the Logging dialog box, you can do the following:

- Download and view the existing log files by clicking **Get Logs**.

  You download log files with the level that was specified at the time they were recorded. If logging is not set up, the log files are downloaded with the default settings. The changed logging level is applied only to the log files that are recorded after you save new logging settings.

  After you extract the `zip` file, check the log files at the following location:

  `/opt/grizzly/logs/`
- *Only if Sharing telemetry data with HYCU is enabled.* Send the existing log files to HYCU Support by clicking **Send Logs**.

  You send log files with the level that was specified at the time they were recorded. If logging is not set up, the log files are uploaded with the default settings. The changed logging level is applied only to the log files that are recorded after you save new logging settings.

- Set up logging. To do so, follow these steps:
  1. Specify values for the following logging settings:

Logging setting	Description
Maximum log file size (MiB)	The maximum size of a log file. The default log file size is 200 MiB, whereas the maximum log file size is 2047 MiB.
Number of log files	The number of log files. The default number is 9.
Level	The following logging levels are available: ◦ Informational (default): Informational messages about the operation of HYCU are recorded to log files. ◦ Detailed: All activity is recorded to log files.
Outbound REST call level *(Available only if the Detailed logging level is selected.)*	The following levels are available: ◦ Off (default): Outbound REST call logs are not recorded to log files. ◦ Informational: Informational messages about the operations related to outbound REST calls are recorded to log files. ◦ Detailed: All activity related to outbound REST calls is recorded to log files.
Inbound REST call level *(Available only if the Detailed logging level is selected.)*	The following levels are available: ◦ Off (default): Inbound REST call logs are not recorded to log files. ◦ Informational: Informational messages about operations related to inbound REST calls are recorded to log files. ◦ Detailed: All activity related to inbound REST calls is recorded to log files.

  2. Use the **Keep settings after upgrade** switch if you want the custom logging settings to remain the same after you upgrade HYCU. As you usually set logging for troubleshooting purposes and do not need the

same logging level for regular use of the product, by default, this switch is turned off.

3. Click **Save**.

> **Note** Keep in mind that the changed logging level is applied only to the log files that are recorded after you save new logging settings.

You can later modify the settings by specifying new values and clicking **Save**, or set the default values by clicking **Default**.

Configuring your network

When configuring your network, you can change network settings such as the IP address and the HYCU listening port number, or enable network bandwidth throttling. Depending on what you want to do, see one of the following sections:

- “Changing network settings” below
- “Limiting network bandwidth” on page 498

> Accessing the Networks dialog box
>
> To access the Networks dialog box, click **Administration**, and then select **Networks**.

Changing network settings

Changing network settings allows you to configure your network to suit the needs of your environment.

> **Important** After you make any changes to the HYCU network settings, you will be signed out automatically and your session will restart.

Limitations

- Multiple network adapters on the same network are not supported.
- *Only if you plan to make HYCU accessible through an IPv6 connection.*
  - The HYCU web user interface and REST API can be accessed through IPv6. Other IPv6 connections with HYCU are not supported.
  - The communication with HYCU instances uses only the IPv4 connections.

- The IPv6 address cannot be configured during the initial deployment of HYCU.
- Only static IPv6 network configuration is supported.
- The pure IPv6 network configuration without an existing IPv4 configuration is not supported.

Consideration

The network that you specified during the HYCU deployment is set to main and is represented by the icon. If you later connect your HYCU backup controller to more than one network by adding a new network interface, you can use another network as the main network. To do so, make sure that a listening port and an SSL certificate are specified for the preferred network, select this network, and then click **Set Main**.

Procedure

1. In the Networks dialog box, on the General tab, the host name of your HYCU backup controller and the networks to which it is connected are displayed. Select the network for which you want to change settings, and then click **Edit**.
2. Configure the IPv4 settings as required: the address, the gateway, and the netmask.
3. *Only if you want to make HYCU accessible through an IPv6 connection.* Use the **Enable IPv6 configuration** switch, and then configure the IPv6 settings as required: the address, the gateway, and the prefix.
4. *Optional.* Configure the domain name and the DNS as required.
5. *Only if your HYCU backup controller is connected to more than one network.* Use the **Enable listening on this port** switch if you want to use the selected network to access the HYCU web user interface. For the network that you specified during the HYCU deployment, the switch is enabled by default.

   > **Note** You cannot disable this switch if the selected network is configured as the main network in HYCU.

6. *Only if the Enable listening on this port switch is enabled.* In the Listening port field, enter the port that you want to use to access the HYCU web user interface (by default, 8443).

   > **Important** If a firewall is configured in your infrastructure, make sure that the port you specified is open.

7. From the SSL certificate drop-down menu, select the SSL certificate that you want to use for this network.
8. Click **Save**.

Limiting network bandwidth

Network bandwidth throttling allows you to limit the bandwidth that is available to HYCU. By defining sites with limited bandwidth, you ensure that enough bandwidth is available for all the network operations in your environment.

Limitation

You can limit network bandwidth only for traffic that is outbound from HYCU.

Considerations

- Network bandwidth throttling is not available in HYCU Manager.
- If the IP address of the storage container to which you plan to restore data is defined in a site for which you want to limit bandwidth, restore performance may be affected.
- Cloud, iSCSI, or SMB targets may utilize multiple IP addresses. Make sure to enter all the utilized IP addresses when defining a site. For details on IP ranges used by public clouds, see respective cloud documentation.
- Throttling network bandwidth for AWS IP addresses also affects telemetry data sharing. Sending log files may take longer.
- *Only if you use HYCU instances.* If you enable network bandwidth throttling, the limit you set applies also to HYCU instances.

Recommendation

It is not recommended to throttle network bandwidth for NFS targets.

Procedure

1. In the Networks dialog box, click the **Throttling** tab, and then click ✙ **New**.
2. Enter a name for the site for which you want to limit bandwidth and, optionally, its description.
3. In the Bandwidth limit field, specify the maximum speed (in KiBps, MiBps, or GiBps) that can be used to transfer data from HYCU to the site.

4. In the IP address/range list field, enter the IP addresses or IP ranges of the sites for which you want to limit bandwidth. You can enter the IP addresses or IP ranges in the following form:
   - Single IPv4 address: 192.0.2.1
   - IPv4 subnet with CIDR prefix: 192.0.2.0/24
   - IPv4 range: 192.0.2.3-192.0.2.100
5. *Only if you want to specify a throttling window.* Select the **Use throttling window** check box, and then, from the Throttling window drop-down menu, select the throttling window that you want to be used for limiting bandwidth. For details on how to create a throttling window, see "Creating a throttling window" below.

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

   > ⓘ Important If you define multiple sites with the same IP addresses, make sure the throttling windows that you assign to these sites do not overlap.

6. Click **Save**.

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

Creating a throttling window

HYCU enables you to define time frames for network bandwidth throttling. If you use a throttling window, network bandwidth is limited only within the specified hours. For example, you can limit network bandwidth during peak production hours when there is more activity on the network.

Procedure

1. In the Networks dialog box, click the **Throttling** tab, and then click **Throttling Windows**.
2. Click **New**.
3. In the Name field, enter a name for the throttling window.
4. From the Time zone drop-down menu, specify the time zone for the throttling window. You can also click one of the displayed time zones (your local time zone or your HYCU backup controller time zone).

5. Select the weekdays and hours during which you want network bandwidth to be limited.

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

6. Click **Save**.

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

Enabling external monitoring of HYCU

You can set up HYCU to be monitored with an external monitoring tool. Monitoring HYCU metrics, such as memory usage and available data disk space, helps you ensure stable functioning of your HYCU backup controller.

After you enable monitoring in the HYCU web user interface, HYCU makes Prometheus metrics available through the `/metrics` endpoint. The endpoint can be accessed at the following URL:

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

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

For example:

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

For instructions on how to enable monitoring in HYCU, see the procedure described in this section. For instructions on how to configure your monitoring tool to access the HYCU metrics, see your monitoring tool documentation.

Prerequisite

The monitoring tool that you use must be able to process metrics in the Prometheus text-based format.

Limitation

The monitoring tool that you use may not collect metrics more often than every 60 seconds.

> Accessing the Monitoring dialog box
>
> To access the Monitoring dialog box, click ⚙ **Administration**, and then select **Monitoring**.

Procedure

1. In the Monitoring dialog box, enable the **Enable metrics** switch.
2. *Only if you want to limit access to the HYCU metrics to specific IP addresses or IP ranges.* Do the following:
   a. Enable the **Enable IP access control** switch.
   b. In the Allowed IPs field, enter one or more IP addresses or IP ranges. If you want to enter more than one IP address or IP range, press **Enter** after each one.
3. *Only if you want to enable basic http authentication for accessing the HYCU metrics.* Enable the **Enable basic http authentication** switch, and then enter the preferred user name and password.
4. Click **Save**.

Managing data retention

When data reaches the end of the retention period defined in the policy, it is automatically expired and deleted from targets by the HYCU cleaning process. For details on automatic data expiration, see “Expiring data automatically” on page 507.

However, if there is a restore point tier (Backup, Copy, Snapshot, and/or Archive) that you do not want to use for restoring data anymore, HYCU enables you to expire it also manually at any time. You can do this by choosing one of the following methods:

Expiration method	Prerequisite	Instructions
Select restore point tiers for expiration from the Retention Management dialog box.	You must be familiar with the information related to expiring data described in this section.	“Procedure” on page 504
Select restore point tiers for expiration from the Virtual Machines, Applications, Shares,		“Procedure” on page 508

Expiration method	Prerequisite	Instructions
Volume Groups, or Buckets panel.		

As part of data retention management, HYCU also allows you to set the preferred expiration date and time for restore points tiers, as well as to specify how long restore point tiers should be retained.

> **Tip** You can check the backup, copy, snapshot, and archive expiration time in the Detail view of the Virtual Machines, Applications, Shares, Volume Groups, or Buckets panel. For details, see "Viewing entity details" on page 399.

When managing data retention, you can select among the following options:

Option	Enables you to...
Expire restore point tiers	Expire the selected restore point tiers. All subsequent restore point tiers in the same backup chains will also be expired (except for the file share and bucket restore point tiers with incremental forever backups in the last backup chain). For details on incremental forever backups for file shares, see "Configuring file share backup options" on page 327. For details on incremental forever backups for buckets, see "Configuring bucket backup options" on page 350.
Set expiration date	Set the date and time at which the selected restore point tiers will expire.
Set retention period	Specify how long the selected restore point tiers should be retained. The expiration time is calculated based on the restore point creation time.

Prerequisite

You must have an Administrator user role assigned.

Considerations

- Your user role determines the restore point tiers that are visible to you. If you are a member of the infrastructure group, you can view all restore point tiers in the data protection environment. If you are a self-service group

administrator, you can view only restore point tiers related to the entities whose owners belong to your self-service group.

- *Only if expiring restore point tiers.*
  - An expire action cannot be undone.
  - If you expire all restore point tiers belonging to a restore point, the backup status is shown as Expired (🔘). This indicates that the restore point cannot be used for restoring data anymore.
  - Restore point tiers other than Backup that have the Failed backup status are not available for expiration.
  - If you selected Target as the backup target type in your policy, the following applies:
    - *For all entities except file shares and buckets with the enabled Incremental forever backup option:*
      - If the most recent restore point is expired, the next backup will be a full backup.
      - When a restore point is expired, any subsequent incremental backups within the same backup chain will also be expired unless the status of the selected restore point is Failed. In this case, only the selected restore point is expired and not the whole backup chain.
      - If you expire the Backup restore point tier, the associated snapshot is also expired, if there is one.
    - *For file shares and buckets with the enabled Incremental forever backup option:* If you are expiring restore point tiers belonging to the last backup chain, consider the following:
      - When a restore point is expired, any subsequent incremental backups within the same backup chain are not expired and the restore point is merged into the next restore point in the backup chain.
      - If you want to expire the entire backup chain, all restore points in the backup chain must be expired. In this case, the next backup will be a full backup.
      - The most recent restore point is protected and cannot be expired separately, but only as part of expiring the entire backup chain. In addition, merging other restore points into the most recent one is not possible.

For details on the Incremental forever backup option, see "Configuring file share backup options" on page 327 or "Configuring bucket backup options" on page 350.

- *Only if expiring data for a volume group that has been backed up both as part of the virtual machine backup and by assigning a policy directly to it.* Before expiring data, make sure that the data is not being used by any virtual machine to which the volume group might be attached.
- *Only if expiring data for a file share or bucket with the enabled Incremental forever backup option.* If your restore point tier data is stored on a target that has WORM protection enabled, removing the data from the target will be performed when the data is no longer WORM protected.

- You cannot set the expiration date or the retention period for restore point tiers whose backup status is Failed, for the Snapshot restore point tiers created for the fast restore, or for the latest Snapshot restore point tiers.
- *Only if setting the retention period for restore point tiers with active WORM retention.* The retention period of restore point tiers with active WORM retention is not modified if the WORM target retention period is longer than the specified one.

Accessing the Retention Management dialog box

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

Procedure

In the Retention Management dialog box, select the action that you want to perform, click **Continue**, and then follow the instructions:

Option	Instructions
Expire Restore Point Tiers	1. From the list of all restore point tiers, select the ones that you want to expire, and then click **Continue**. > Tip You can use filtering options to display only specific types of restore point tiers. For details, see "Filtering restore point tiers" on page 506. 2. Review the expiration information, and then click **Expire**. 3. Click **Expire** to confirm that you want the selected restore point tiers to be expired.

Option	Instructions
Set Expiration Date	1. From the list of all restore point tiers, select the ones for which you want to set an expiration date, and then click **Continue**. > Tip You can use filtering options to display only specific types of restore point tiers. For details, see "Filtering restore point tiers" on the next page. 2. In the Expiration date and time field, specify the date and time when you want the selected restore point tiers to be expired. 3. Review the expiration information, and then click **Execute**.
Set Retention Period	1. From the list of all restore point tiers, select the ones for which you want to set a retention period, and then click **Continue**. > Tip You can use filtering options to display only specific types of restore point tiers. For details, see "Filtering restore point tiers" on the next page. 2. Depending on whether you want to set a new retention period or to modify the retention period for the selected restore point tiers, select one of the following options, and then do as required: • **Set New Retention Period:** Set a retention period for the data (in months, weeks, days, or hours). • **Adjust Current Retention Period**: Select **Increase** or **Decrease**, and then set for how long you want to increase or decrease the retention period of the data (in months, weeks, days, or hours). > Important Keep in mind that setting a new retention period or modifying the retention period for the selected restore point tiers can affect the retention period of other restore point tiers in the same backup chains. 3. Review retention information, and then click **Execute**.

The next HYCU cleaning process removes data from the targets.

Filtering restore point tiers

The restore point tier filtering options allow you to filter and focus only on specific types of restore point tiers (for example, the restore point tiers whose data is stored on a particular target). After you apply any of the filters, only restore point tiers that match the filter criteria are displayed and you can easily find what you need.

> Tip When you specify any of the filtering options and you want to include all available categories, click **Select all**.

Procedure

1. In the side panel of the Select Restore Point Tiers dialog box that opens after you select the preferred retention management action, select one or more of the following filtering options:

Filtering option	Filter by...
Tiers	Restore point tiers (Backup, Copy , Snapshot, Archive daily, Archive weekly, Archive monthly, and Archive yearly).
Backup type	Type of backup (Incremental and Full).
Restore point creation date	Time range within which the restore point was created.
Restore point tier expiration date	Time range within which the restore point tier is marked for expiration.
Source	Sources that host the entities with the restore point tiers.
Target	Targets on which restore point tier data is stored.
Type	Entities with the restore point tiers.
Owner	Owners that are assigned to the entities with the restore point tiers.
Backup status	Backup statuses (Success, Failure, and Warning).

2. Click **Apply Filters**.

Expiring data automatically

When any of the restore point tiers reaches its retention period, it is grayed out in the HYCU web user interface. Depending on which backup target type you selected in your policy, tiers are expired as follows:

Backup target type	Conditions for tier expiration
Target	*For all entities except file shares and buckets with the enabled Incremental forever backup option:* A tier is expired when the last tier in the backup chain reaches its retention period. This means that this data is not removed from HYCU or the target until all tiers in the backup chain are expired. However, if there is a restore point that contains the Archive tier, this restore point is kept although the rest of the backup chain is expired. In addition, if this restore point is an incremental backup, it is changed to full. *For file shares and buckets with the enabled Incremental forever backup option:* A tier is expired when it reaches its retention period and is merged into the next restore point. However, if there is a restore point that contains the Archive tier, this restore point is kept.
Snapshot	A tier is expired when the snapshot reaches its retention period. However, if there is a restore point that contains the Archive tier, this restore point is kept although the snapshot is expired.

Considerations

- Changing the retention period in the policy does not affect existing backups.
- HYCU automatically expires the last backup chain of an unprotected entity (the one from which a policy was unassigned or whose policy was deleted), whereas the last backup chain of a protected entity is never expired automatically.

Expiring data from the entity panel

As an alternative to selecting restore point tiers for expiration from the Retention Management dialog box, you can select the preferred restore point

tiers also from the Virtual Machines, Applications, Shares, Volume Groups, or Buckets panel.

Depending on the entity for which you want to expire data, access one of the following panels:

- Accessing the Virtual Machines panel

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

- Accessing the Applications panel

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

- Accessing the Shares panel

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

- Accessing the Volume Groups panel

  To access the Volume Groups panel, in the navigation pane, click **Volume Groups**.

- Accessing the Buckets panel

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

Procedure

1. In the Virtual Machines, Applications, Shares, Volume Groups, or Buckets panel, select the entity for which you want to expire data.
2. Select one or more restore points that contain the restore point tiers that you want to expire, and then click **Expire**.
3. From the list of available restore point tiers, select the ones that you want to expire, and then click **Continue**. You are redirected to the Retention Management > Expire dialog box.
4. Review the expiration information, and then click **Expire**.
5. Click **Expire** to confirm that you want the selected restore point tiers to be expired.

Setting power options

You can set power options for the HYCU backup controller so that its activities are suspended or resumed.

> Accessing the Power Options dialog box
>
> To access the Power Options dialog box, click ⚙ **Administration**, and then select **Power Options**.

Power option	Description
Suspend all	Pauses all HYCU backup controller activities. If you want the HYCU backup controller activities to automatically resume after a specified amount of time, in the Auto resume after field, specify the number of hours (1–168) to pass before the activities are resumed. All currently running jobs are allowed to complete normally. All jobs that are in the queue will start when the HYCU backup controller is resumed. While activities are paused, you cannot start any new backup jobs.
Suspend cleanup	Pauses the cleanup of targets and, if enabled, the purge of events and jobs. The snapshot cleanup is not affected.
Resume	Allows HYCU backup controller activities to continue.

Managing secrets

HYCU enables you to securely store, access, and manage your credentials (secrets) by employing the Conjur secrets management solution. After you store your HYCU secrets in Conjur as one or more Conjur configurations (that is, a set of one or more policies where you define your security rules), you can easily manage them and be confident that your resources can be accessed only by authorized parties.

Prerequisites

- Make sure you set up your Conjur environment and stored HYCU secrets as a set of one or more policies. For instructions, see Conjur documentation.
- The SSL certificate of the Conjur server must be imported to HYCU by an infrastructure administrator. For instructions, see "Configuring SSL certificates" on page 524.

Limitations

- HYCU credentials that you plan to store in Conjur may not start with `${` and end with `}`.
- HYCU users cannot be managed by using Conjur. For details on HYCU users, see "Managing users" on page 440.

Considerations

- Mixed mode is possible. This means that you do not have to store all your HYCU secrets in Conjur to be able to benefit from this integration.
- When providing secrets that are stored in Conjur, make sure to use the following syntax in HYCU:

  `${<PathtoSecret>}`

Recommendation

Only if you plan to change the names of secrets in Conjur. Every time you change the name of a secret in Conjur, it is recommended that you clear the cache in HYCU. To do so, in the Secrets Management dialog box, click **Clear Cache**. This is done also automatically by HYCU every 24 hours, but for the continuity of your business processes, it is recommended that you do it manually.

> **Tip** All the fields in the HYCU web user interface that support entering values stored in Conjur have the following icon next to them: for infrastructure group configuration and for private configuration.

Accessing the Secrets Management dialog box

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

Adding a Conjur configuration

Consideration

For each HYCU data protection environment, one infrastructure group Conjur configuration and one private Conjur configuration for each self-service group can be added.

Procedure

1. In the Secrets Management dialog box, depending on which type of Conjur configuration you want to perform, click one of the following buttons:

Conjur configuration type	Description
Add Infrastructure Group Configuration	*Available only if you are an infrastructure group administrator.* Enables you to provide secrets stored in Conjur when performing all data protection and administrative tasks. For example, when adding sources and targets, adding identity providers, and so on.
Add Private Configuration	*Available if you are an infrastructure or a self-service group administrator.* Enables you to provide secrets stored in Conjur when performing the following tasks: • Adding cloud accounts. • Assigning credential groups to virtual machines. • Setting up webhook notifications.

2. In the Appliance URL field, enter the URL of the Conjur server that you are connecting to.
3. In the Account filed, enter the name of the account that you specified during the Conjur environment setup.
4. In the Authentication login field, enter the Conjur host user name. For example:

```
host/HycuPolicy/hycuBckupController
```

In this example, `host` is the type of user, `HycuPolicy` is the name of the policy to which the user belongs, and `hycuBckupController` is the user name.

5. In the Authentication API key field, enter the API key that corresponds to the Conjur host user name.
6. *Only if you are an infrastructure group administrator.* When performing one type of Conjur configuration, enable the **Use same values for private configuration** or **Use same values for infrastructure group configuration** switch if you want to use the same values for performing the other type of Conjur configuration.
7. Click **Save**.

Editing a Conjur configuration

Procedure

1. In the Secrets Management dialog box, click **Edit** next to the Conjur configuration that you want to edit.
2. Edit the selected Conjur configuration as required. For detailed information on Conjur configuration properties, see "Adding a Conjur configuration" on the previous page.
3. Click **Save**.

Removing a Conjur configuration

Procedure

1. In the Secrets Management dialog box, click **Remove** next to the Conjur configuration that you want to remove.
2. Click **Yes** to confirm that you want to remove the selected Conjur configuration.

Configuring email delivery

You can configure HYCU to send email notifications about events and scheduled reports. For sending email notifications, HYCU supports the SMTP protocol and Microsoft 365 Graph API.

Depending on the type of email delivery that you want to configure, see one of the following sections:

- “Configuring email delivery using SMTP” below
- “Configuring email delivery for Microsoft 365” on the next page

> Accessing the Email Delivery Settings dialog box
>
> To access the Email Delivery Settings dialog box, click ⚙ **Administration**, and then select **Email Delivery Settings**.

Configuring email delivery using SMTP

Prerequisite

For using the STARTTLS or SSL/TLS security mode to secure email traffic: A valid SSL certificate must be imported to HYCU. For instructions, see “Securing SMTP connections” on page 545.

Limitation

HYCU supports only basic SMTP authentication.

Procedure

1. In the Email Delivery Settings dialog box, from the Email Delivery Method drop-down menu, select **SMTP**.
2. Provide the following information:

Required information	Description
Username	User name of the account on the SMTP server.
Password	Password of the account on the SMTP server.
Display name	Display name of the email sender.
Hostname or IP address	Host name or IP address of the SMTP server.
Port	Port number to be used (usually set to 25).
Security mode	Protocol used to secure email traffic—can be set to None, STARTTLS, or SSL/TLS.

Required information	Description
From email address	Email address from which email notifications will be sent.

3. If you want to verify the provided SMTP configuration by sending a test email, do the following:
   a. Enable the **Send test email** switch.
   b. In the Test Email Recipient field, enter the email recipient that should receive the test email.
4. Click **Save**.

After you configure email delivery using SMTP, you can continue with configuring HYCU to send email notifications. For details on how to do this, see "Setting up email notifications" on page 388.

To reconfigure your existing SMTP configuration, delete the current settings by clicking **Delete Settings**, and then redefine them.

Configuring email delivery for Microsoft 365

Prerequisite

An Azure or Azure Government service principal must be added to HYCU and it must have the Mail.Send API application permission (not delegated permission) assigned. For instructions, see "Adding an Azure service principal" on page 468 or "Adding an Azure Government service principal" on page 469.

Recommendation

The Mail.Send API application permission allows the service account to send emails as any user. It is recommended that you create an Exchange Online security group to limit sending emails only to members of the group. In this case, make sure that the email address that you enter in the From Email Address field is a member of this group.

Procedure

1. In the Email Delivery Settings dialog box, from the Email Delivery Method drop-down menu, select **Microsoft 365**.

2. From the Cloud Account drop-down menu, select the Azure or Azure Government service principal that you want to use for authentication.
3. In the Display Name field, enter the display name of the email sender.
4. In the From Email Address field, enter the email address from which email notifications will be sent.
5. If you want to verify the provided Microsoft 365 configuration by sending a test email, do the following:
   a. Enable the **Send test email** switch.
   b. In the Test Email Recipient field, enter the email recipient that should receive the test email.
6. Click **Save**.

After you configure email delivery for Microsoft 365, you can continue with configuring HYCU to send email notifications. For details on how to do this, see “Setting up email notifications” on page 388.

To reconfigure your existing configuration, delete the current settings by clicking **Delete Settings**, and then redefine them.

Upgrading HYCU

You can upgrade HYCU when a new software release version is available. HYCU can be upgraded from the HYCU repository, a URL address, or the image that you upload from your local machine. If you decide to upgrade from the HYCU repository, HYCU will use the upgrade image from the public HYCU repository. However, if your data protection environment has no internet access, you can also configure HYCU to use the upgrade image from a private local repository.

Prerequisites

- Back up or create snapshots of your HYCU backup controller and the connected HYCU instances before upgrading HYCU. You can remove these snapshots after completing the upgrade.
- Jobs that you do not want to be aborted must be finished (the upgrade process aborts all currently running jobs).
- The HYCU data disk must be larger than the HYCU system disk. For instructions on how to increase disk size, see “Increasing the size of the

HYCU virtual disks" on page 552.

- *Only if you plan to perform the upgrade from the public HYCU repository.* The HYCU backup controller must have access to the repository.
- *Only if you plan to perform the upgrade from the private local repository.* Your web server must host a location that stores the following:
  - One or more HYCU image files
  - JSON index file

    For instructions on how to create the JSON index file, see "Creating the JSON index file" on page 518.

Considerations

- Any users that were signed in to the HYCU web user interface of the HYCU backup controller that is being upgraded should perform a hard reload of the web user interface page in their web browser after the process completes.
- Upgrading removes any previously added update packages from the update folder on the HYCU backup controller.
- The EFI and BIOS boot modes are supported.
- The HYCU backup controller activities are automatically suspended during the upgrade.
- *Only if you do not have access to the public HYCU repository.* You can configure the HYCU backup controller to upgrade from the private local repository. For details on how to customize the HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.

Recommendation

Before upgrading HYCU, back up the HYCU backup controller. For instructions, see "Backing up virtual machines" on page 182.

Procedure

1. Click ⚙ **Administration**, and then select **Software Upgrade**.
2. In the Software Upgrade dialog box, on the Release tab, check the current version of HYCU. You can also check whether any newer version is available on the HYCU Support portal by clicking the **Check for new version** link.
3. Depending on how you want to perform the upgrade, do one of the following:

I want to...	Instructions
Upgrade from the public HYCU repository.	a. From the Upgrade From drop-down menu, select **HYCU repository**. b. From the Available Versions drop-down menu, select the version to which you want to upgrade HYCU. c. Select the **Download image to HYCU backup controller and perform upgrade locally** check box if your network capacity is insufficient for an uninterrupted upgrade procedure or if you experience frequent network outages.
Upgrade from a URL link.	a. From the Upgrade From drop-down menu, select **URL**. b. In the URL field, enter the full path to the upgrade image file. Example `https://myserver.com/images/hycu-5.3.0-1234.hyimg` c. Select the **Download image to HYCU backup controller and perform upgrade locally** check box if your network capacity is insufficient for an uninterrupted upgrade procedure or if you experience frequent network outages.
Upgrade from an image file that was uploaded from your local machine.	a. From the Upgrade From drop-down menu, select **Upload file**. b. In the File field, click **Choose a file** to select the upgrade image file on your machine, or drag and drop the upgrade image file into the dialog box.

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

You will be signed out of HYCU. The HYCU virtual machine will reboot automatically. After the upgrade process completes, you can sign in to the HYCU web user interface.

Only if you are using HYCU instances. After the HYCU upgrade, the HYCU backup controller automatically upgrades the connected HYCU instances. However, if for any reason, the upgrade of the HYCU instances fails, proceed as follows:

1. *For Nutanix clusters and vSphere environments:* If you want HYCU to use a different method to upgrade the HYCU instances, make sure that the following conditions are met:
   - The Nutanix cluster or the vSphere environment to which the HYCU instances are deployed is added as a source to HYCU.
   - The upgrade image is available on the source.

   Otherwise, go to the next step.
2. *For all environments:* Delete, and then recreate HYCU instances as follows:
   a. Delete the existing HYCU instances. For details, see "Deleting a HYCU instance" on page 487.
   b. Create new HYCU instances with the latest HYCU version. For details, see "HYCU instances" on page 25.

Creating the JSON index file

Prerequisite

The Python script for the automatic creation of the JSON index file must have access to the local files that are stored on the web server that hosts your private repository.

> **Note** To ensure access to the local files on your web server, you can run the script from any location, not exclusively from the HYCU backup controller.

Creating the JSON index file automatically

To create the JSON index file automatically, run the `createHycuRepo` Python script. You can find the script in the HYCU backup controller at the following location:

```
/opt/grizzly/bin/createHycuRepo.py
```

Use the following syntax:

```
python3 createHycuRepo.py <RepoRootFolder> [ -o <OutputFile> ]
```

> **Note** Without the `-o` option specified, the script will print the output to the console.

Writing the JSON index file manually

You can write the JSON index file manually by using the following example structure:

```
{
    "images": [
        {
            "version": "5.2.0-2487",
            "annotation": "HYCU from May 2025",
            "url": "example_builds/5.2.0-2487/hycu-5.2.0-
2487.hyimg",
            "size": 1732180359,
            "compatibility": {
                "min_version": "5.2.0-2334",
                "max_version":null
            },
            "type": "hyimg"
        },
        {
            "version": "5.2.0-2515",
            "annotation": "develop_2025-05-05_12:24:48_
gf9f06433754",
            "url": "example_builds/5.2.0-2515/hycu-5.2.0-
2515.hyimg",
            "size": 1732225216,
            "compatibility": {
                "min_version": "5.2.0-2334"
            },
            "type": "hyimg"
        },
       {...}
    ]
}
```

Applying HYCU updates

After you receive a HYCU update from HYCU Support, you can apply it to your current product version. An update can be applied only to an installed compatible product version. For example, an update labeled 1.2.3-4567 can be

applied to the product version 1.2.3 whereas an update labeled 1.2.4-5678 cannot.

> **Note** Each HYCU update addresses a cumulative set of issues.

Prerequisites

- *For applying an update to a HYCU backup controller:* The HYCU backup controller activities must be suspended. For instructions on how to do this, see "Setting power options" on page 509.
- Jobs that you do not want to be aborted must be finished (the update application process aborts all currently running jobs). You can check this by filtering the Jobs list by the Executing job status. For instructions, see "Filtering and sorting data" on page 406.
- *For applying an update to a HYCU instance:* The same update must be applied to the corresponding HYCU backup controller.
- *For applying an update by using the shell script:* You must know credentials of an operating system user account that has administrative user rights on the HYCU virtual machine where you plan to apply the update.

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

Considerations

- The update that you apply to the HYCU backup controller is not automatically applied to HYCU instances or HYCU Managers, if there are any in your data protection environment.
- *For applying an update to a HYCU backup controller or a HYCU Manager:* Any users that were signed in to the HYCU web user interface of the HYCU virtual machine where the update is being applied should perform a hard reload of the web user interface page in their web browser after the process completes.

Recommendation

Before applying an update to a HYCU backup controller, back up the HYCU backup controller. For instructions, see "Backing up virtual machines" on page 182.

You can apply HYCU updates:

- From the HYCU web user interface

  Use this method if you want to apply an update to a HYCU backup controller, a HYCU instance, or a HYCU Manager. For instructions, see "Applying an update by using the HYCU web user interface" below.
- By using the shell script

  Use this method if you are unable to sign in to the HYCU web user interface. For instructions, see "Applying an update by using the shell script" on page 523.

Applying an update by using the HYCU web user interface

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

- "Applying an update to a HYCU backup controller or a HYCU Manager" below
- "Applying an update to a HYCU instance" on the next page

Applying an update to a HYCU backup controller or a HYCU Manager

Procedure

1. Sign in to the HYCU web user interface.
2. Click **Administration**, and then select **Software Upgrade**.
3. In the Software Upgrade dialog box, click the **Updates** tab.
4. In the Update label column, check if the package of the preferred update has already been added to HYCU or HYCU Manager, and then do one of the following:
   - If the update label is not present, follow these steps:
     a. Click **Add**.
     b. Click **Browse**, and then browse for the update package (in the ZIP format).
     c. Click **Add Update**.
     d. Select the update label.
   - If the update label is present, select it.

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

5. Click **Apply Update**.
6. Verify that the displayed digital fingerprint matches the one that you were given by HYCU Support, and then click **Yes** to start the update process.

   You are automatically signed out of the HYCU web user interface, and can track the progress of applying the update on the web user interface sign-in page.
7. When the process completes, perform a hard reload of the HYCU web user interface page in your web browser.
8. *Only if you applied an update to a HYCU backup controller.* Do the following:
   a. Sign in to the HYCU web user interface.
   b. Resume activities of the HYCU backup controller. For instructions on how to do this, see "Setting power options" on page 509.

You can later delete any of the update packages that you do not need anymore by selecting it and clicking **Delete**.

Applying an update to a HYCU instance

Procedure

1. Sign in to the HYCU web user interface.
2. Click **Administration**, and then select **HYCU Instances**.
3. In the HYCU Instances dialog box, select the HYCU instance to which you want to apply an update, and then click **Updates**.
4. In the Update label column, check if the package of the preferred update has already been added to HYCU, and then do one of the following:
   - If the update label is not present, follow these steps:
     a. Click **Add**.
     b. Click **Browse**, and then browse for the update package (in the ZIP format).
     c. Click **Add Update**.

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

     d. Select the update label.
   - If the update label is present, select it.

   > **Tip** You can click **Info** to review the list of issues that the update resolves.

5. Click **Apply Update**.
6. Verify that the displayed digital fingerprint matches the one that you were given by HYCU Support, and then click **Yes** to start the update process.

The HYCU instance status icon in the HYCU Instances dialog box turns gray to indicate the ongoing process. You can track the progress of the process by checking the status of the corresponding job in the Jobs panel. When the update is applied, the HYCU instance status icon turns green.

You can later delete any of the update packages that you do not need anymore by selecting it, and then clicking **Delete**.

Applying an update by using the shell script

Procedure

1. Sign in to the web user interface that you are using to manage your virtualization environment, and connect to the HYCU virtual machine where you plan to apply the update.
2. Sign in to the operating system with a user account that has administrative user rights.
3. Open a command shell, and then run the following command:

```
cd /opt/grizzly/bin/
```

4. Run the following command to retrieve the list of update packages that are already added to the HYCU virtual machine:

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

5. If the label of the preferred update is not present on the list, follow these steps:
   a. Extract the contents of the update package (in the ZIP format). The package contains the main update file, installation instructions, and digital fingerprints.
   b. Use the `/usr/bin/cksum` and `/usr/bin/md5sum` commands to verify that the digital fingerprint of the main update file matches the one that you were given by HYCU Support.
   c. Copy the main update file in the archived TAR (`.tar.gz`) format to the

following directory on the HYCU virtual machine:

```
/hycudata/opt/grizzly/updates
```

> Tip Run the following command to review the list of issues that the update resolves:
>
> ```
> sudo ./HycuPatch.sh -patch_info <UpdateLabel>
> ```

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

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

7. *Only if you applied an update to a HYCU backup controller.* Do the following:
   a. Sign in to the HYCU web user interface.
   b. Resume activities of the HYCU backup controller. For instructions on how to do this, see "Setting power options" on page 509.

Configuring SSL certificates

To establish trusted and secure communication in your data protection environment, you must configure SSL certificates.

> Accessing the SSL Certificates dialog box
>
> To access the SSL Certificates dialog box, click ⚙ **Administration**, and then select **SSL Certificates**.

In the SSL Certificates dialog box that opens, you can view the information about your SSL certificate—the certificate name, the certificate common name, the certificate expiry date, the certificate key type, and the generated certificate signing request (CSR).

Consideration

After you create or import an SSL certificate, make sure to update also the HYCU network settings by specifying this certificate. For details on how to do this, see "Configuring your network" on page 496.

Recommendation

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

Depending on what you want to do, see one of the following:

I want to...	Procedure
Create a self-signed certificate.	"Creating a self-signed certificate" below
Create a certificate signing request.	"Creating a certificate signing request" on the next page
Import a custom certificate to HYCU.	"Importing a custom certificate" on page 527

Creating a self-signed certificate

Procedure

1. In the SSL Certificates dialog box, click ✙ **Generate**.
2. Select **Generate self-signed certificate**, and then click **Next**.
3. Provide the following certificate-related information:
   - Name
   - Common name
   - Organization
   - Organization unit
   - Location
   - Country
   - *Optional.* State
   - Key algorithm
   - Key size

   > ⓘ Important The maximum number of characters in each field is 64.
4. Click **Generate**.

The self-signed certificate is added to the list of SSL certificates. Keep in mind that each SSL certificate that is generated through HYCU is valid for 824 days and that you must maintain the validity of the certificate.

Creating a certificate signing request

Procedure

1. In the SSL Certificates dialog box, click **Generate**.
2. Select **Generate certificate signing request**, and then click **Next**.
3. Provide the following certificate-related information:
   - Name
   - Common name
   - Organization
   - Organization unit
   - Location
   - Country
   - *Optional.* State
   - Key algorithm
   - Key size
   - *Optional.* Subject alternative name (SAN)

     The SAN is a list of possible names of the HYCU backup controller. Each name can be one of the following:
     - FQDN (for example, `hycu-bc.hycu.local`)
     - Wildcard prefix "*." followed by a domain name (for example, `*.hycu.local`)
     - IPv4 address (for example, `10.1.100.1`)
     - IPv6 address (for example, `fe80::1234:5678:9abc:def0`)

   > **Important** The maximum number of characters in each field is 64.
4. Click **Generate**.

An SSL private key is added to the list of SSL certificates and the icon in the CSR column indicates that the certificate signing request has been created. Click **Download Generated CSR** to download the generated CSR.

After you create and download the CSR, you can send it to a certificate authority to create a certificate. The certificate that is created from the CSR by the certificate authority must be added to the SSL private key to complete the SSL key pair in HYCU. For instructions, see “Completing the SSL key pair” on the next page.

Completing the SSL key pair

Procedure

1. In the SSL Certificates dialog box, select the SSL private key, and then click **Edit**.
2. Browse for the following files:
   - Certificate: The file with the certificate that was created from the CSR.
   - *Optional.* CA certificate/chain: The file with the CA-signed certificate or trust chain certificates.
3. Click **Update**.

Importing a custom certificate

HYCU enables you to import an SSL key pair, or a CA-signed certificate or trust chain certificates.

Prerequisites

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

Considerations

- If the certificate uses a wildcard for the Common Name (CN), make sure that the Certificate Subject Alt Name field includes all possible host names or FQDNs, and their corresponding IP addresses. Otherwise, the certificate may be recognized as invalid by your web browser.
- If you are importing an SSL key pair from a PFX file, consider the following:
  - The PFX file must contain the entire trust chain to the root CA certificate.
  - The PFX must contain a single private key along with its associated certificate. If multiple keys exist in the PFX file, the import may fail or only one private key may be imported.

Procedure

1. In the SSL Certificates dialog box, click **Import**.
2. Depending on whether you want to import an SSL key pair, a CA-signed certificate or trust chain certificates from a file, or a CA-signed certificate or trust chain certificates from a host, select one of the following options, and then click **Next** and follow the instructions:

Option	Instructions
Import SSL keypair from PEM files	a. Enter a name for your certificate. b. Browse for the following files: • *Optional.* CA certificate/chain: The file with the CA-signed certificate or trust chain certificates. • Certificate: The file with the certificate corresponding to the private key that you are importing. • Private key: The file with the private key that is associated with the certificate that you are importing. The private key should be created with the RSA or ECDSA algorithm in the PKCS#1 or PKCS#8 format. The minimum key size for private keys created with the RSA algorithm is 2048 bits. **Note** If you use Conjur for managing your HYCU secrets, you can enable the **Retrieve values from secrets manager** switch if you want to provide the secret instead of browsing for the file. For details on managing secrets, see "Managing secrets" on page 509.
Import SSL keypair from PFX file	a. Enter a name for your certificate. b. Browse for the PFX file that contains the required SSL key pair. c. *Optional.* Enter the passphrase of the PFX file.

Option	Instructions
Import CA certificate/chain from file	a. Enter a name for your certificate. b. Browse for the file with the CA-signed certificate or trust chain certificates.
Import CA certificate/chain from host	a. Enter the server host name or IP address, and the port. The following examples show which host names and ports to use in common configuration scenarios: • If HYCU is configured to use HTTPS for WinRM connections to virtual machines, enter the host name or IP address of the virtual machine for which you want to establish an HTTPS connection, and the HTTPS port (usually 5986). • If HYCU is configured to use LDAP over SSL (LDAPS), enter the LDAPS server host name or IP address, and the LDAPS port (usually 636). • If you are using STARTTLS or SSL/TLS for SMTP connections, enter the SMTP server host name or IP address, and the port for authenticated SMTP connections (465 for the SSL/TLS security mode, and 587 or 25 for the STARTTLS security mode). b. From the Security mode drop-down menu, select the preferred security mode. c. Click **Retrieve**. d. Review the certificates and select the one that you want to import. If you select the CA-signed certificate for the import, its trust chain certificates will be trusted as well.

3. Click **Import**.

You can also change the name of any self-signed or custom certificate (click **Edit** and make the required modification) or delete the ones that you do not need anymore (click **Delete**).

Sharing telemetry data with HYCU

You can configure HYCU to collect telemetry data. This data helps HYCU to provide proactive support and improved performance to better meet your data protection environment needs.

Sharing diagnostic data through telemetry enables proactive, contextualized support for HYCU as follows:

1. Collects detailed data on your data protection environment that includes the syslog files, HYCU internal data base (PostgreSQL) logs, system activity information (sar), data distribution statistics for file share backups, HYCU license information, and other detailed information on your specific infrastructure, and then sends this data to HYCU Support.

   > **Important** HYCU does not collect any sensitive information from your data protection environment.

2. Analyzes collected data, generates internal reports, and identifies eventual problems or unfavorable trends considerably reducing issue resolution time.
3. Provides you with feedback on your HYCU environment that addresses eventual issues and instructs you on how to adjust your environment and to improve infrastructure and performance.

> **Note** You need to enable telemetry data sharing for each HYCU backup controller that you want to include in the advanced troubleshooting.

Prerequisite

You must have a valid HYCU Support user account.

Consideration

When a Pay-as-you-go license is applied to HYCU, sharing telemetry data with HYCU is enabled by default and cannot be disabled.

Accessing the Telemetry dialog box

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

Procedure

In the Telemetry dialog box, use the **Share telemetry data with HYCU, Inc.** switch to allow HYCU to collect your telemetry data, and then click **Save**.

HYCU starts collecting data and sends it to HYCU Support. Later, the telemetry diagnostic data is sent to HYCU Support once a day. You can view the collection job status in the Jobs panel.

If you later decide that you no longer want to share your telemetry data with HYCU, disable the **Share telemetry data with HYCU, Inc.** option for each configured HYCU backup controller.

> Note When the **Share telemetry data with HYCU, Inc.** option is enabled, you can send the log files to HYCU Support. For more information, see "Setting up logging" on page 493.

Removing HYCU

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

To remove HYCU, follow these steps:

1. Sign in to HYCU, and then unassign policies from all entities as follows:
   - To unassign policies from virtual machines:
     a. In the navigation pane, click **Virtual Machines**.
     b. Select all virtual machines, and then click **Set Policy**.
     c. Click **Unassign**.
     d. Click **Yes** to confirm that you want to unassign the policies from the selected virtual machines.
   - To unassign policies from applications:
     a. In the navigation pane, click **Applications**.
     b. Select all discovered applications, and then click **Set Policy**.
     c. Click **Unassign**.
     d. Click **Yes** to confirm that you want to unassign the policies from the selected applications.
   - To unassign policies from file shares:

a. In the navigation pane, click **Shares**.
b. Select all file shares, and then click **Set Policy**.
c. Click **Unassign**.
d. Click **Yes** to confirm that you want to unassign the policies from the selected file shares.

- To unassign policies from volume groups:
  a. In the navigation pane, click **Volume Groups**.
  b. Select all volume groups, and then click **Set Policy**.
  c. Click **Unassign**.
  d. Click **Yes** to confirm that you want to unassign the policies from the selected volume groups.
- To unassign policies from Buckets:
  a. In the navigation pane, click **Buckets**.
  b. Select all Buckets, and then click **Set Policy**.
  c. Click **Unassign**.
  d. Click **Yes** to confirm that you want to unassign the policies from the selected buckets.

2. *Only if you used HYCU instances.* Remove the existing HYCU instances. For instructions, see "Deleting a HYCU instance" on page 487.
3. *Only if HYCU was used for file share protection.* Remove the file server snapshots created by HYCU. To do so, on the HYCU backup controller, run the `/opt/grizzly/bin/HycuCleanup.pl` script as follows:

```
sudo perl HycuCleanup.pl -c <FileServer> -u <Username> -p
<Password> -dnfs -all
```

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

   > **Important** By running this command, you will also remove all file server snapshots whose names start with `hycu-` (case insensitive).

4. *For Nutanix clusters:* On the HYCU backup controller, run the `/opt/grizzly/bin/HycuCleanup.pl` script as follows:
   - To remove virtual machine and volume group snapshots created by HYCU:

```
sudo perl HycuCleanup.pl -c <NutanixCluster> -u <Username> -
p <Password> -dvms -all
```

```
sudo perl HycuCleanup.pl -c <NutanixCluster> -u <Username> -
p <Password> -dvgs -all
```

In these instances, *<NutanixCluster>* is the name of the Nutanix cluster in the following format: `https://<ServerName>:<Port>`.

> **Important** By running these commands, you will also remove all third-party snapshots created by using Nutanix REST API v3 whose names start with the IP address.

- To remove volume groups created by HYCU:

```
sudo perl HycuCleanup.pl -c <NutanixCluster> -u <Username> -
p <Password> -dvg -all
```

In this instance, *<NutanixCluster>* is the name of the Nutanix cluster in the following format: `https://<ServerName>:<Port>`.

> **Important** By running this command, you will also remove all volume groups created by using Nutanix REST API v3 whose names start with `HYCU-` (case insensitive).

5. Remove data from targets. To do so, on each target, delete the `bkpctrl` folder.
6. Sign in to the management console of the environment in which HYCU is deployed, and then delete the HYCU backup controller virtual machine. For instructions, see your platform documentation.

Chapter 13

Tuning your data protection environment

Administration tasks that you perform through the **Administration** menu to customize HYCU for your data protection environment are usually sufficient to successfully manage it. However, sometimes the needs of your organization require additional administration tasks to be performed for optimal performance, a higher security level, or interaction with external applications, as well as for taking advantage of a broader spectrum of HYCU options.

I want to...	Procedure
Access the HYCU backup controller virtual machine by using SSH.	“Accessing the HYCU backup controller virtual machine by using SSH” on the next page
Enable HTTPS for WinRM connections.	“Enabling HTTPS for WinRM connections” on page 538
Ensure that HYCU can use Kerberos authentication.	“Using Kerberos authentication” on page 538
Configure FIPS-compliant mode for HYCU.	“Configuring FIPS mode for HYCU” on page 540
Set up LDAPS authentication.	“Setting up LDAPS authentication” on page 541
Set up two-factor authentication.	“Setting up two-factor authentication” on page 542
Manage API keys.	“Managing API keys” on page 543
Manage FIDO authenticators.	“Managing FIDO authenticators” on page 544
Secure SMTP connections.	“Securing SMTP connections” on

I want to...	Procedure
	page 545
Set up HYCU to use multiple networks.	“Setting up HYCU to use multiple networks” on page 545
Increase the size of the HYCU virtual disks.	“Increasing the size of the HYCU virtual disks” on page 552
Assign required privileges to a vSphere user.	“Assigning privileges to a vSphere user” on page 552
Configure Prism Central user permissions.	“Configuring Prism Central user permissions” on page 556
Use the HYCU REST API to automate tasks.	“Using the HYCU REST API Explorer” on page 560
Use `hyCLI`.	“Using the command-line interface” on page 560
Use the pre and post scripts to perform necessary actions before and after the backup and the restore are performed.	“Using the pre and post scripts” on page 561

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. The only two exceptions for which you should use SSH are restarting the HYCU application server (the Grizzly server) or the entire appliance.

> ⓘ **Important** Using SSH to perform any tasks other than restarting the HYCU application server or the entire appliance is not recommended.

After you deploy the HYCU virtual appliance, you can use the following default credentials to access the HYCU backup controller virtual machine by using SSH:

User name: **hycu**

Password: **hycu/4u**

Changing the default SSH password

For security purposes, it is highly recommended that you change the default SSH password. To do so, follow these steps:

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

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

   When requested, enter the default password.
2. Change the password for the `hycu` user:

   ```
   passwd
   ```

   When requested, enter the default password again, and then enter and verify your new password.

Configuring SSH public key authentication

Adding an SSH public key to HYCU and using it to access the HYCU backup controller enables you to add an additional layer of security to your data protection environment by providing a more secure alternative to SSH password authentication. If you are using HYCU for file share or bucket protection and you configure SSH public key authentication for accessing the HYCU backup controller, you can use the same SSH public key also to access your HYCU instances. For added security, you can choose to disable SSH password authentication.

Limitation

The supported SSH key types are RSA, ECDSA, and Ed25519.

> ### Accessing the SSH Authentication dialog box
>
> To access the SSH Authentication dialog box, click ⚙ **Administration**, and then select **SSH Authentication**.

Procedure

1. In the SSH Authentication dialog box, click **Add Public Key**.
2. Enter a name for the SSH public key and the SSH public key.
3. Click **Save**.

The SSH public key is added to HYCU. For each added key, the name, creation date, and key fingerprint are displayed.

You can also delete any of the existing SSH public keys by selecting the key, and then clicking **Delete**.

If after configuring SSH public key authentication you want to disable SSH password authentication, you can do so by disabling the **Allow password authentication** switch, and then clicking **Save**.

Disabling SSH access

You can disable SSH access at any time. To do so, follow these steps:

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

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

   When requested, enter the password for the `hycu` user.
2. Shut down the SSH service:

   ```
   sudo systemctl stop sshd.service
   ```

   When requested, enter the password for the `hycu` user.
3. Disable the SSH service:

   ```
   sudo systemctl disable sshd.service
   ```

   If requested, enter the password for the `hycu` user.

After performing this procedure, your SSH connection will be disabled. To re-enable SSH, you need to connect to the HYCU backup controller virtual machine through the console of the respective source.

Managing the HYCU application server

To manage the HYCU application server, follow these steps:

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

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

   When requested, enter the password for the `hycu` user.
2. Perform the preferred operation on the HYCU application server:

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

   When requested, enter the password for the `hycu` user.

> ⓘ **Important** If you plan to restart the PostgreSQL server, make sure the HYCU application server is stopped before and started after restarting the PostgreSQL server.

Enabling HTTPS for WinRM connections

If you want to add an additional layer of security, you can configure HYCU to use HTTPS for WinRM connections to virtual machines.

Procedure

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

1. Configure WinRM for HTTPS. For details on how to do this, see Microsoft documentation.
2. *Only if WinRM is configured with a certificate that was signed by a private certificate authority or with a self-signed certificate.* Import the CA-signed certificate or trust chain certificates to HYCU. For instructions, see “Importing a custom certificate” on page 527.

Using Kerberos authentication

For stronger security, HYCU by default uses Kerberos to authenticate with the following:

- SMB targets
- Windows virtual machines

- File shares that use the SMB protocol

To ensure that HYCU can use Kerberos authentication, your environment must be correctly configured. Otherwise, the legacy NTLM authentication will be used.

The prerequisites that must be fulfilled to enable HYCU to use Kerberos authentication in different data protection scenarios are listed in the following table:

Data protection scenario	Prerequisites
Using an SMB target	Your firewall must be configured to allow communication with the Kerberos Domain Controller. For details, see "Targets" on page 30.
Protecting Windows VMs	• The user name that you use for assigning a credential group to the virtual machine must include the DNS domain name. For details, see "Enabling access to data" on page 174. • If you plan to restore individual files, your firewall must be configured to allow communication with the Kerberos Domain Controller. For details, see "Restore of individual files and application awareness" on page 34.
Protecting applications running on Windows VMs	• Your firewall must be configured to allow communication with the Kerberos Domain Controller. For details, see "Restore of individual files and application awareness" on page 34. • Your DNS server must be configured to perform reverse DNS lookups. For details, see "Enabling access to application data" on page 252.
Protecting file shares that use the SMB protocol	• Your firewall must be configured to allow communication with the Kerberos Domain Controller. For details, see "File share protection" on page 36. • Additional prerequisites apply for different file server types. For details, see "Adding a file server" on page 85.

Configuring FIPS mode for HYCU

HYCU can be configured to operate to be compliant with the Federal Information Processing Standards (FIPS) 140-2 that establish security requirements for cryptography modules (which encryption algorithms and methods for generating encryption keys can be used).

Depending on the nature of your business, you can either enable or disable FIPS mode for HYCU. To check whether FIPS mode is enabled (disabled by default), open a remote session to the HYCU backup controller, and then as the root user or by using sudo, run the following command:

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

Limitations

- *For Linux virtual machines:* Applications cannot be discovered and therefore protected.
- SMB file shares cannot be protected.

Considerations

- When FIPS mode is enabled, using Kerberos authentication is mandatory and the fallback to NTLM authentication is disabled. For details on how to ensure that HYCU can use Kerberos authentication, see “Using Kerberos authentication” on page 538.
- After you upgrade HYCU, FIPS mode will be disabled. If required, make sure to re-enable it.
- *Only if you use HYCU instances.* You must enable FIPS mode for each HYCU instance separately (independent of the HYCU backup controller).

Enabling FIPS mode for HYCU

Procedure

Open a remote session to the HYCU backup controller, and then as the root user or by using sudo, do the following:

1. Stop the HYCU application server:

   ```
   systemctl stop grizzly.service
   ```

2. Enable FIPS-compliant mode:

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

3. Reboot the HYCU backup controller:

```
reboot
```

Disabling FIPS mode for HYCU

Procedure

Open a remote session to the HYCU backup controller, and then as the root user or by using sudo, do the following:

1. Stop the HYCU application server:

```
systemctl stop grizzly.service
```

2. Disable FIPS-compliant mode:

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

3. Reboot the HYCU backup controller:

```
reboot
```

Setting up LDAPS authentication

If you want to add an extra layer of protection and ensure the confidentiality of data, you can configure HYCU to use LDAP over SSL (LDAPS) for secure user authentication. For this authentication to work, HYCU must trust the LDAPS server certificate. Depending on the type of the LDAPS server certificate, trust is established in one of the following ways:

- If the LDAPS server certificate was signed by a public certificate authority, HYCU will trust it automatically.
- If the LDAPS server certificate was signed by a private certificate authority, or if it is self-signed, you must import the CA-signed certificate or trust chain certificates to HYCU. For instructions, see “Importing a custom certificate” on page 527.

Setting up two-factor authentication

You can set up two-factor authentication to add an extra layer of security when signing in to HYCU. The following authentication methods are supported:

- Time-based one-time passwords (OTP) generated by an OTP application.
- Authenticators compliant with the FIDO protocol (FIDO authenticators), such as security keys and fingerprint reader.

When setting up two-factor authentication for HYCU, you must complete the following tasks:

Task	Instructions
1. Perform the necessary preparation steps for the selected authentication method.	• *For OTP*: Provide instructions to users and make sure that they have access to an OTP application. • *For FIDO authenticators:* Make sure that the following conditions are met: ◦ Authenticators are set up correctly. For instructions, see the authenticator documentation. ◦ The DNS is configured properly. ◦ The host name is resolved properly.
2. Create or edit a user for whom you want to enable two-factor authentication, and then add this user to a user group.	Follow the procedures described in “Creating a user” on page 445 and “Adding a user to a group” on page 450.

After you perform these tasks, users must authenticate their sign-ins by signing in to HYCU as described in “Signing in to HYCU” on page 68.

Managing API keys

API keys are needed in the following scenarios:

- If you enable two-factor authentication for using the REST API.
- If you enable API key authentication when adding a Hybrid Cloud Edition controller to HYCU Manager.

You can generate or revoke your API keys by using the API keys option.

Consideration

As a user with the Administrator role assigned, you can edit other users' information through the Self-Service panel. For details, see "Creating a user" on page 445.

> **Accessing the API Keys dialog box**
>
> To access the API Keys dialog box, click at the upper right of the screen, and then select **API Keys**.

Generating an API key

Procedure

1. In the API Keys dialog box, click **New**.
2. Enter a name for the API key.
3. *Optional.* Set the expiry date. If you do not set the expiry date, the API key does not expire.
4. Click **Generate**. The API key is displayed.
5. Write the API key down and store it safely.

   > **Important** For security reasons, the API key will never be displayed again, so make sure to write it down and keep it safe. You can copy the API key to the clipboard by clicking **Copy to Clipboard**.
   > Your API key can be used to access your data, therefore, treat it like a password.

6. Click **Finish**.

Revoking an API key

Procedure

1. In the API Keys dialog box, select the API key that you want to revoke, and then click **Revoke**.
2. Click **Revoke** to confirm that you want to revoke the API key. The API key is immediately revoked.

Managing FIDO authenticators

If the FIDO two-factor authentication method is enabled for your account, you need to set up a FIDO authenticator. You can add or revoke your FIDO authenticators by using the FIDO Authenticators option.

Adding a new FIDO authenticator

Considerations

- As a user in the Infrastructure group with the Administrator role assigned, you can edit other users' information through the Self-Service panel. For details, see “Creating a user” on page 445.
- Make sure that you use a fully qualified domain name when signing in to HYCU and that DNS is correctly configured. Otherwise, authentication may fail.

> Accessing the FIDO Authenticators dialog box
>
> To access the FIDO Authenticators dialog box, click at the upper right of the screen, and then select **FIDO Authenticators**.

Procedure

1. In the FIDO Authenticators dialog box, click **New**. The Security Setup wizard opens.
2. Follow the wizard instructions to create the FIDO authenticator. The process depends on the type of authenticator you select and the operating system version.
3. In the Name field, enter a name for the FIDO authenticator.
4. Click **Register**.

Revoking a FIDO authenticator

Procedure

1. In the FIDO Authenticators dialog box, select the authenticator that you want revoke, and then click **Revoke**.
2. Click **Revoke** to confirm that you want to revoke the FIDO authenticator. The FIDO authenticator is immediately revoked.

Securing SMTP connections

If you want to add an extra layer of protection and ensure the confidentiality of data, you can configure HYCU to use SMTP over SSL/TLS or STARTTLS for secure user authentication. For this authentication to work, HYCU must trust the SMTP server certificate. Depending on the type of the SMTP server certificate, trust is established in one of the following ways:

- If the SMTP server certificate was signed by a public certificate authority, HYCU will trust it automatically.
- If the SMTP server certificate was signed by a private certificate authority, or if it is self-signed, you must import the CA-signed certificate or trust chain certificates to HYCU. For instructions, see “Importing a custom certificate” on page 527.

Setting up HYCU to use multiple networks

You can set up HYCU to operate in a multi-network environment, allowing it to have two network adapters assigned to different VLANs or network segments. This is especially useful if you have dedicated storage used for backups in a different network than HYCU. For example:

- HYCU could be located on the 10.0.0.0/16 VLAN and a storage box could be located on the 192.168.0.0/24 VLAN.
- You need to access the HYCU web user interface from a network other than the virtual machine network. In this case, it is recommended to have a dedicated NIC for data transfer that must be on the same VLAN as the

Nutanix Controller virtual machines, in addition to the NIC for the web user access.

> **Note** *For Nutanix clusters:* While the bulk of data traffic during a backup takes place over the additional network, part of it is still done through the management network. This is because HYCU uses the Nutanix data services IP address to consume data through Nutanix Volumes, which must be in the same subnet as the management network of the CVMs.
> For details on this limitation, see Nutanix documentation.

Limitation

You cannot set up HYCU to use multiple networks in AWS GovCloud (US) Azure, or Azure Government environments.

Considerations

- *For file server and object server environments:*
  - The main network must correspond to a network segment where both the HYCU backup controller and the additional HYCU instances can see and establish a connection to each other.
  - Both virtual machines (the HYCU backup controller and one or more connected HYCU instances) must be able to connect to the file server or the object server.
  - Each network adapter must be on a different subnet.
  - *Only if the DNS servers are specified.* The DNS servers on all subnets must return the same results.
  - *For Nutanix ESXi clusters:* When upgrading HYCU, network settings on all additional network adapters will be set to the default values. Make sure to reconfigure the HYCU instance after the upgrade.

Depending on the environment in which you want to set up HYCU to use multiple networks, perform one of the following procedures:

- “Setting up HYCU to use multiple networks on a Nutanix AHV or Nutanix ESXi cluster” on the next page
- “Setting up HYCU to use multiple networks in a vSphere environment” on page 548
- “Setting up HYCU to use multiple networks in a XenServer environment” on page 549

- "Setting up HYCU to use multiple networks in an Azure Local environment" on page 550
- "Setting up HYCU to use multiple networks on a Hyper-V cluster" on page 551

Setting up HYCU to use multiple networks on a Nutanix AHV or Nutanix ESXi cluster

Procedure

1. Sign in to the Nutanix Prism web console, and then add an additional network adapter:
   a. In the menu bar, click **Home**, and then select **VM**.
   b. Click the **Table** tab to display the VM Table view, and then, from the list of virtual machines, select your HYCU virtual machine.
   c. Click **Update**, and then navigate to the Network Adapters (NIC) section.
   d. Click **Add New NIC**, and then select the required VLAN and click **Add**.
   e. Click **Save**.

   For details, see Nutanix documentation.
2. Configure the network. To do so, depending on how the VLAN is set up, select one of the following approaches:
   - VLAN has IP address (DHCP) management enabled

     Assign the IP address directly from the Nutanix Prism web console.
   - VLAN does not have IP address (DHCP) management enabled

     Depending on whether you are configuring an additional network for a HYCU instance, a HYCU backup controller, or a HYCU Manager, do one of the following:
     - *For a HYCU instance:* Configure the network manually:
       a. Open a remote session to the HYCU backup controller virtual machine:

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

       b. Open the `ifcfg-mainnetwork.template` file located at `/opt/grizzly/misc/`, and then follow the instructions provided in this template. Make sure to run the specified commands as the root user or by using sudo.

- *For a HYCU backup controller or a HYCU Manager:* Configure the network from the HYCU backup controller or HYCU Manager user interface. For details, see “Configuring your network” on page 496.

After the new network adapter is properly configured, you can add a target located on another VLAN to HYCU.

Setting up HYCU to use multiple networks in a vSphere environment

Procedure

1. Sign in to the vSphere Client, and then add an additional network adapter:
   a. Click the **VMs** tab, and then navigate to your HYCU backup controller.
   b. Right-click the HYCU backup controller, and then select **Edit Settings**.
   c. From the New device drop-down menu, select **Network**, and then click **Add**.
   d. From the New Network drop-down menu, select the required network.

      > ⓘ Important Make sure not to select a vSphere distributed switch (dvSwitch) for the virtual NIC option.

   e. Click **OK**.

   For details, see VMware documentation.
2. Depending on whether you are configuring an additional network for a HYCU instance, a HYCU backup controller, or a HYCU Manager, do one of the following:
   - *For a HYCU instance:* Configure the network manually:
     a. Open a remote session to the HYCU backup controller virtual machine:

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

     b. Open the `ifcfg-mainnetwork.template` file located at `/opt/grizzly/misc/`, and then follow the instructions provided in this template. Make sure to run the specified commands as the root user or by using sudo.
   - *For a HYCU backup controller or a HYCU Manager:* Configure the network from the HYCU backup controller or HYCU Manager user interface. For details, see “Configuring your network” on page 496.

After the new network adapter is properly configured, you can add a target located on another network to HYCU.

Setting up HYCU to use multiple networks in a XenServer environment

Procedure

1. Sign in to XenCenter, and then add an additional network adapter:
   a. Navigate to your HYCU backup controller.
   b. On the Networking tab, click **Add Interface...**.
   c. Select the required network and MAC address, and, optionally, specify a QoS limit.
   d. Click **Add**.

   For details, see XenServer documentation.
2. Depending on whether you are configuring an additional network for a HYCU backup controller, a HYCU instance, or a HYCU Manager, do one of the following:
   - *For a HYCU backup controller or a HYCU Manager:* Configure the network from the HYCU backup controller or HYCU Manager user interface. For details, see "Configuring your network" on page 496.
   - *For a HYCU instance:* Configure the network manually:
     a. Open a remote session to the HYCU backup controller virtual machine:

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

     b. Open the `ifcfg-mainnetwork.template` file located at `/opt/grizzly/misc/`, and then follow the instructions provided in this template. Make sure to run the specified commands as the root user or by using sudo.

After the new network adapter is properly configured, you can add a target located on another network to HYCU.

Setting up HYCU to use multiple networks in an Azure Local environment

Procedure

1. Sign in to Windows Admin Center.
2. Add an additional network adapter:
   a. Go to **Virtual machines**, and then click the HYCU virtual machine.
   b. Click **Settings**, and then go to **Networks**.
   c. Click **Add network adapter**.
   d. From the Isolation mode drop-down menu, select **VLAN**, and then define the VLAN identifier.
   e. Click **Save network settings**.
3. Configure the network. To do so, depending on how the VLAN is set up, select one of the following approaches:
   - VLAN has IP address (DHCP) management enabled

     Assign the IP address directly from Windows Admin Center.
   - VLAN does not have IP address (DHCP) management enabled

     Depending on whether you are configuring an additional network for a HYCU instance, a HYCU backup controller, or a HYCU Manager, do one of the following:
     - *For a HYCU instance:* Configure the network manually:
       a. Open a remote session to the HYCU backup controller virtual machine:

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

       b. Open the `ifcfg-mainnetwork.template` file located at `/opt/grizzly/misc/`, and then follow the instructions provided in this template. Make sure to run the specified commands as the root user or by using sudo.
     - *For a HYCU backup controller or a HYCU Manager:* Configure the network from the HYCU backup controller or HYCU Manager user interface. For details, see “Configuring your network” on page 496.

After the new network adapter is properly configured, you can add a target located on another network to HYCU.

Setting up HYCU to use multiple networks on a Hyper-V cluster

Procedure

1. Sign in to Windows Admin Center.
2. Add an additional network adapter:
   a. Go to **Virtual machines**, and then click the HYCU virtual machine.
   b. Click **Settings**, and then go to **Networks**.
   c. Click **Add network adapter**.
   d. From the Isolation mode drop-down menu, select **VLAN**, and then define the VLAN identifier.
   e. Click **Save network settings**.
3. Configure the network. To do so, depending on how the VLAN is set up, select one of the following approaches:
   - VLAN has IP address (DHCP) management enabled

     Assign the IP address directly from Windows Admin Center.
   - VLAN does not have IP address (DHCP) management enabled

     Depending on whether you are configuring an additional network for a HYCU instance, a HYCU backup controller, or a HYCU Manager, do one of the following:
     - *For a HYCU instance:* Configure the network manually:
       a. Open a remote session to the HYCU backup controller virtual machine:

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

       b. Open the `ifcfg-mainnetwork.template` file located at `/opt/grizzly/misc/`, and then follow the instructions provided in this template. Make sure to run the specified commands as the root user or by using sudo.
     - *For a HYCU backup controller or a HYCU Manager:* Configure the network from the HYCU backup controller or HYCU Manager user interface. For details, see “Configuring your network” on page 496.

After the new network adapter is properly configured, you can add a target located on another network to HYCU.

Increasing the size of the HYCU virtual disks

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

Procedure

1. Sign in to the management console of the environment in which HYCU is deployed.
2. Shut down the HYCU backup controller.
3. Increase the size of the HYCU disk as required.
4. Turn on the HYCU backup controller.

For instructions on how to perform these steps, see your platform documentation.

Assigning privileges to a vSphere user

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

Procedure

1. Sign in to the vSphere Client as an administrator.
2. Click **Administration** > **Roles**.
3. Add a new role, and then type its name (for example, **HYCU**).
4. Depending on your data protection environment, select the required privileges for the role:

Nutanix ESXi cluster

Privilege category	Backup privileges	Restore privileges	Upgrade and HYCU instance creation privileges
Cryptographic operations	Direct Access	Direct Access	Not applicable
Datastore	Browse datastore	Browse datastore	• Allocate space • Low level file operations
Network	Not applicable	Not applicable	Assign network
vApp	Not applicable	Not applicable	Import
Virtual Machine > Configuration	Not applicable	Not applicable	• Add existing disk • Add new disk • Change Settings • Remove disk
Virtual Machine > Interaction	Not applicable	Not applicable	Power On
Virtual Machine > Inventory	Not applicable	Not applicable	• Create from existing • Remove
Virtual Machine > Provisioning	Not applicable	Not applicable	Clone virtual machine
vSphere Tagging	Assign or Unassign vSphere Tag	Assign or Unassign vSphere Tag	Not applicable

vSphere environment

Privilege category	Backup privileges	Restore privileges	Upgrade privileges
Cryptographic operations	Direct Access	Not applicable	Not applicable
Datastore	• Browse datastore • Low level file operations	• Allocate space • Low level file operations	• Allocate space • Low level file operations
Global	• Disable methods • Enable methods	Not applicable	Not applicable
Host > Local operations	Not applicable	• Create virtual machine • Delete virtual machine • Reconfigure virtual machine	Not applicable
Network	Not applicable	• Assign network • Configure	Assign network
Resource	Not applicable	Assign virtual machine to resource pool	Not applicable
vApp	Not applicable	Add virtual machine	Import
Virtual Machine > Configuration	• Toggle disk change tracking • Change Settings	All privileges	• Add existing disk • Add new disk • Add or remove device • Change Settings

Privilege category	Backup privileges	Restore privileges	Upgrade privileges
			• Remove disk • Rename
Virtual Machine > Interaction	Power On	• Answer question • Connect devices • Power Off • Power On	Power On
Virtual Machine > Inventory	Not applicable	• Create new • Register • Remove • Unregister	• Create from existing • Remove
Virtual Machine > Provisioning	• Allow read-only disk access • Allow virtual machine download • *For backing up a template:* Mark as template • *For backing up a template:* Mark as virtual machine	Allow disk access	Clone virtual machine
Virtual Machine > Snapshot management	• Create snapshot • Remove snapshot	Revert to snapshot	Not applicable
vSphere Tagging	Assign or Unassign vSphere Tag	Assign or Unassign vSphere Tag	Not applicable

5. Click **Create** to create the role.
6. Click **Administration** > **Global Permissions**.
7. Add a new permission for the role that you have created.

For details, see VMware documentation.

Configuring Prism Central user permissions

If you want to perform certain data protection tasks related to virtual machines running on a Nutanix AHV or a Nutanix ESXi cluster, the following prerequisites must be fulfilled:

- The cluster that hosts the virtual machines must be registered with Prism Central.
- The role that is assigned to your Prism Central user must have access to the required operations.

The following tables list the data protection tasks for which this is applicable and the associated required operations for each type of Nutanix cluster.

Nutanix AHV

Data protection task	Required operations
Setting up automatic policy assignment For details on setting up automatic policy assignment, see “Setting up automatic policy assignment” on page 151.	• Category > View Value Category
Setting up automatic virtual machine assignment by using tags For details on setting up automatic virtual machine assignment by using tags, see “Setting up automatic virtual machine assignment by using tags” on page 453.	
Backing up virtual machines from their replicas in remote office/branch office (ROBO) environments if Nutanix DR is	• AHV VM > View Virtual Machine

Data protection task	Required operations
enabled in Prism Central	
Restoring or cloning a virtual machine to an overlay network For details on restoring virtual machines, see “Restoring virtual machines” on page 184.	• AHV VM > View Virtual Machine • Subnet > View Subnet
Restoring or cloning a secondary IP address	• AHV VM > View Virtual Machine • AHV VM > Update Virtual Machine NIC List
Restoring or cloning VM categories	• AHV VM > View Virtual Machine • Category > View Value Category • AHV VM > Update Virtual Machine Categories
Restoring or cloning VM owners	• AHV VM > View Virtual Machine • AHV VM > Update Virtual Machine Owner
Restoring or cloning VM projects	• AHV VM > View Virtual Machine • AHV VM > Update Virtual Machine Project
Restoring and powering on a virtual machine	• AHV VM > View Virtual Machine • AHV VM > Delete Virtual Machine • Task > View Task
Validating the virtual machine backup and automatically deleting the cloned virtual machine For details on validating the virtual machine backup, see “Validating the virtual machine backup” on page 228 and “Setting up a validation policy” on page 425.	

Nutanix ESXi

Data protection task	Required operations
Backing up virtual machines from their replicas in remote office/branch office (ROBO) environments if Nutanix DR is enabled in Prism Central	• AHV VM > View Virtual Machine
Restoring or cloning a virtual machine to an overlay network For details on restoring virtual machines, see "Restoring virtual machines" on page 184.	• AHV VM > View Virtual Machine • Subnet > View Subnet

For details on how to create roles, grant them access to operations, and assign them to users in Prism Central, see Nutanix documentation.

Setting permissions in AWS GovCloud (US)

The following list shows a minimum set of permissions that are required for protecting virtual machines and applications running on virtual machines in an AWS GovCloud (US) data protection environment:

Service	Required permissions
EC2	ec2:AllocateAddress ec2:AssignIpv6Addresses ec2:AssociateAddress ec2:AttachNetworkInterface ec2:AttachVolume ec2:CopySnapshot ec2:CreateNetworkInterface ec2:CreateSnapshot ec2:CreateSnapshots ec2:CreateSubnet ec2:CreateTags

Service	Required permissions
	ec2:CreateVolume ec2:CreateVpc ec2:DeleteSnapshot ec2:DeleteVolume ec2:DeregisterImage ec2:DescribeAddresses ec2:DescribeAvailabilityZones ec2:DescribeImages ec2:DescribeInstanceAttribute ec2:DescribeInstances ec2:DescribeInstanceStatus ec2:DescribeInstanceTypeOfferings ec2:DescribeInstanceTypes ec2:DescribeKeyPairs ec2:DescribeMacHosts ec2:DescribeNetworkAcls ec2:DescribeNetworkInterfaceAttribute ec2:DescribeNetworkInterfaces ec2:DescribeRegions ec2:DescribeSecurityGroups ec2:DescribeSnapshotAttribute ec2:DescribeSnapshots ec2:DescribeSubnets ec2:DescribeVolumes ec2:DescribeVpcAttribute ec2:DescribeVpcEndpoint ec2:DescribeVpcEndpointServices ec2:DescribeVpcs ec2:DescribeVpnConnections ec2:DetachNetworkInterface ec2:DetachVolume ec2:GetConsoleOutput ec2:GetInstanceUefiData ec2:ImportImage ec2:ImportSnapshot ec2:ImportVolume ec2:ModifyAddressAttribute

Service	Required permissions
	ec2:ModifyNetworkInterfaceAttribute ec2:RegisterImage ec2:RunInstances ec2:StartInstances ec2:StopInstances ec2:TerminateInstances
Elastic Block Store (EBS)	All READ and WRITE permissions
IAM	iam:PassRole
Systems Manager	ssm:GetParameters

Using the HYCU REST API Explorer

HYCU provides a REST API that can be used by external applications to interact with the HYCU backup controller, retrieve information from it, and automate tasks. All functionality exposed through the HYCU user interface is also available through the HYCU REST API. You can use the HYCU REST API Explorer to interact with the API and view the expected input and output formats for each endpoint.

To access the HYCU REST API Explorer, follow these steps:

1. Click ⓘ **Help** at the upper right of the screen, and then select **REST API Explorer**. The HYCU REST API Explorer opens.
2. In the list of functionality groups, you can expand the preferred group by clicking **List Operations**. A list of API endpoints is displayed.
3. Click any of the endpoints to show the description, the parameters, and the output format. You can fill in the fields, and then click **Try it out!** to call an API and get output data.

Using the command-line interface

> ⓘ Important The hyCLI functionality is deprecated and will be removed from HYCU in the next release.

You can manage your data protection environment also by using the HYCU command-line user interface (hyCLI). hyCLI provides the functionality comparable to the HYCU web user interface and enables you to implement scripts for automating certain tasks.

To use hyCLI, contact HYCU Support.

Using the pre and post scripts

If you want to use the pre/post scripts to perform necessary actions before and after the backup and the restore are performed, these scripts should return an exit code of 0 for success and any other value for failure. In the latter case, the data protection operation is also affected as follows:

- An exit code is greater than 0: The status of the job (and the backup in the case of the backup operation) will be set to Warning and the job will continue.
- An exit code is less than 0: The status of the job (and the backup in the case of the backup operation) will be set to Failed.

During the execution of the scripts, the following environment variables are exported:

Environment variable	Description
`HYCU_BKPCTRL_URL`	HYCU backup controller URL
`HYCU_BKPCTRL_UUID`	HYCU backup controller UUID
`HYCU_VM_UUID`	Virtual machine UUID
`HYCU_BACKUP_UUID`	Backup UUID
`HYCU_JOB_UUID`	Job UUID
`HYCU_TARGET_UUID`	Target UUID
`HYCU_VM_NAME`	Virtual machine name[a]
`HYCU_TARGET_NAME`	Target name[a]
`HYCU_TARGET_PATH`	Path to the data on the target
`HYCU_SUCCESS`	*Available only for post scripts.* Success of the data protection operation.

Environment variable	Description
HYCU_PREEXEC_RETURN_CODE	*Available only for post scripts.* Exit code of the pre script.

[a] If the name contains the space character or any of the following characters: " ' , ; & % € ( ) < > { } | ^ ` ˇ , these characters are replaced with an underscore before the export.

For details on how to specify pre and post scripts, see the following sections:

- “Specifying pre/post-backup and pre/post-snapshot scripts” on page 179
- “Restoring individual files” on page 243

Chapter 14

Monitoring data protection environments

HYCU Manager is designed to provide you with the visibility you need to proactively monitor all your data protection environments, allowing you to view their overall status from a single console. With HYCU Manager, data protection information received from all registered HYCU controllers is consolidated in one place and you can easily access this information for the on-premises (HYCU) and the following cloud data protection environments:

- HYCU Data Protection as a Service for Azure (HYCU for Azure)
- HYCU for Microsoft 365 and Google Workspace

Before you can start monitoring your data protection environments, you must complete the following tasks:

Task	Instructions
1. Deploy the HYCU virtual appliance in the HYCU Manager mode.	“Deploying the HYCU virtual appliance” on page 23
2. Add a HYCU controller to HYCU Manager.	“Adding a HYCU controller” on the next page

After you add HYCU controllers to HYCU Manager, you can perform various tasks to effectively monitor your data protection environments. Depending on what you want to do in HYCU Manager, see one of the following sections:

I want to...	Instructions
Get an at-a-glance overview of the data collected from all the data protection environments for which I am responsible.	“Using the HYCU Manager console ” on page 569
View data protection information that	“Viewing entity data” on page 571

I want to...	Instructions
is received from added Hybrid Cloud Edition controllers.	
View all events that occurred in HYCU Manager and all events that are associated with added Hybrid Cloud Edition controllers.	“Viewing events” on page 573
Obtain reports on different aspects of the monitored data protection environments.	“Using HYCU Manager reports” on page 574
Perform various administration tasks related to HYCU Manager.	“Performing administration tasks” on page 577

> Tip You can set up the appearance of your HYCU Manager console to fit your preferences. For details, see “Setting up the appearance of your HYCU web user interface” on page 438.

Managing HYCU controllers

You can use the HYCU Controllers panel to add, edit, and delete the HYCU controllers, as well as to view the information about each of them.

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

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

Adding a HYCU controller

I want to monitor...	HYCU controller to add	Instructions
On-premises (HYCU) data protection environment	Hybrid Cloud Edition controller	“Adding a Hybrid Cloud Edition controller” on the

I want to monitor…	HYCU controller to add	Instructions
		next page
HYCU for Azure data protection environment	Azure controller	"Adding an Azure controller" below
HYCU for Microsoft 365 and Google Workspace data protection environment	Microsoft 365 and Google Workspace controller	"Adding a Microsoft 365 and Google Workspace controller" on the next page

Adding a Hybrid Cloud Edition controller

Procedure

1. In the HYCU Controllers panel, click ✙ **Add**.
2. Select **Hybrid Cloud Edition controller**, and then click **Next**.
3. Enter the name of the HYCU backup controller.
4. Enter the URL of the HYCU backup controller.
5. Depending on the type of authentication you want to use, do one of the following:
   - *Basic authentication:* Make sure the **Use API key authentication** switch is disabled, and then enter the user name and password of an infrastructure group administrator.
   - *API key authentication:* Enable the **Use API key authentication** switch, and then enter your API key. For details on how to generate and revoke an API key, see "Managing API keys" on page 543.
6. Click **Add**.

After you add a Hybrid Cloud Edition controller, you can view data protection information received from this HYCU controller for each related virtual machine, application, file share, server, volume group, and bucket. For details, see "Viewing entity data" on page 571.

Adding an Azure controller

Prerequisites

- You must own a premium tier Platform license. For details, see "Licensing" on page 488.

- You must have an active subscription for HYCU for Azure. For details, see HYCU for Azure documentation.
- A cloud account must be added to HYCU. For instructions, see "Adding an Azure service principal" on page 468.

Consideration

If the required cloud account is not added to HYCU, the option for adding the Azure controller is grayed out in the HYCU Manager console.

Procedure

1. In the HYCU Controllers panel, click **Add**. The Add Controller dialog box opens.
2. Select **Azure controller**, and then click **Next**. The Add Controller > Azure Controller dialog box opens.
3. Select the HYCU for Azure protection sets that you want to monitor. You can also search for a protection set by entering its name in the Search field.

   > **Tip** You can see all Azure resource groups that are included in each available protection set by clicking >.

4. Click **Add**.

After you add an Azure controller, you can select it and click **Details** to see its protection sets and all included resource groups.

Adding a Microsoft 365 and Google Workspace controller

Prerequisite

You must have an active subscription for HYCU for Microsoft 365 and Google Workspace. For details, see the *HYCU for Microsoft 365 and Google Workspace Quick Start Guide*.

Procedure

1. In the HYCU Controllers panel, click **Add**. The Add Controller dialog box opens.
2. Select **Microsoft 365 and Google Workspace controller**, and then click **Next**. The Add Controller > Microsoft 365 and Google Workspace Controller dialog box opens.
3. Enter the name of the Microsoft 365 and Google Workspace controller.

4. Enter the URL of your HYCU for Microsoft 365 and Google Workspace web user interface.
5. Enter the access and reseller tokens that you received when you subscribed to HYCU for Microsoft 365 and Google Workspace.
6. Click **Add**.

Viewing information about HYCU controllers

In the HYCU Controllers panel, you can view specific information about each added HYCU controller. However, keep in mind that not all information might be applicable to your data protection scenario.

> Tip To narrow down the list of displayed HYCU controllers, you can use the filtering options described in "Filtering and sorting data" on page 406.

Considerations

- If you click the name of the HYCU controller, you are directed to the relevant web user interface.
- *For the Hybrid Cloud Edition controller*: If you use the same user account to sign in to both HYCU Manager and HYCU, after you click the HYCU controller in HYCU Manager, you are automatically directed and signed in to the related HYCU web user interface. Keep in mind that single sign-on is available only if the version of both HYCU Manager and the HYCU backup controller is 5.3.0 or later.
- *For the Hybrid Cloud Edition controller and the Azure controller:* You can click any of the icons representing the information about the HYCU controller and you are automatically directed to the specific panel listing all the corresponding items. If your HYCU controller is a Hybrid Cloud Edition controller, these items are also filtered according to your selection. For example, if you click an icon representing the percentage of the protected virtual machines, you are directed to the Virtual Machines panel listing only all the protected virtual machines.

HYCU controller information	Description
Name	Name of the HYCU controller. A Hybrid Cloud Edition controller is represented by the

HYCU controller information	Description
	icon and the name of the HYCU backup controller. If you use HYCU Manager to monitor also the cloud data protection environments, you can view cloud controllers. A cloud controller is represented by: • *HYCU for Azure*: The icon and the name of the Azure service principal and the HYCU for Azure protection set. • *HYCU for Microsoft 365 and Google Workspace*: The icon and the name of the Microsoft 365 and Google Workspace controller.
Version	HYCU software release version on the HYCU backup controller.
Status	Status of the HYCU controller (Active, Unavailable, Suspended, or Undefined).
Backups	Exact number and percentage of successful and failed backups (also per user groups).
DR-ready VMs	Number of DR-ready virtual machines and servers. A virtual machine or a server is DR-ready if all backups in the current backup chain are stored on one of the cloud targets.
VM protection	Exact number and percentage of protected and unprotected virtual machines (also per user groups).
App protection	Exact number and percentage of protected and unprotected applications (also per user groups).
Share protection	Exact number and percentage of protected and unprotected file shares (also per user groups).
Bucket protection	Exact number and percentage of protected and unprotected buckets (also per user groups).
Policy compliance	Percentage of compliant and non-compliant policies.
Target utilization	Percentage of used storage space on targets.

You can export data that you view in the HYCU Controllers panel to a file in JSON or CSV format. For details on how to do this, see “Exporting the contents of the panel” on page 416.

Editing a HYCU controller

Limitation

You cannot edit the Azure controller.

Procedure

1. In the HYCU Controllers panel, select the HYCU controller that you want to edit, and then click **Edit**.
2. Edit the selected HYCU controller as required.
3. Click **Save**.

Removing a HYCU controller

Procedure

1. In the HYCU Controllers panel, select the HYCU controller that you want to remove from HYCU Manager, and then click **Remove**.
2. Click **Remove** to confirm that you want to remove the selected HYCU controller.

Using the HYCU Manager console

The HYCU Manager console provides you with an at-a-glance overview of the data collected from all the data protection environments for which you are responsible.

> Accessing the Console panel
>
> To access the Console panel, in the navigation pane, click **Console**.

Within each widget in the HYCU Manager console, you can find information related to your data protection environments. However, keep in mind that not all widgets might be applicable to your data protection scenario.

Console widget	Description
Virtual Machines	Shows the percentage of compliant and non-compliant virtual machines, the percentage of protected and unprotected virtual machines, the percentage of R-Shield compliant and non-compliant virtual machines, and the percentage of protected virtual machines that have the DR-ready status in your data protection environments. A virtual machine is considered: • Compliant: If the time since the last successful backup is lower than its RPO and the estimated time to recover is lower than its RTO. • Protected: If it has at least one valid backup available and does not have the Exclude policy assigned. • R-Shield compliant: If the changes in backup data size are below the threshold defined in the R-Shield policy, and no malware or ransomware is detected. • DR-ready: If all backups in the current backup chain are stored on one of the cloud targets.
Applications	Shows the percentage of compliant and non-compliant applications, the percentage of protected and unprotected applications, and the percentage of R-Shield compliant and non-compliant applications in your data protection environments. An application is considered: • Compliant: If the time since the last successful backup is lower than its RPO and the estimated time to recover is lower than its RTO. • Protected: If it has at least one valid backup available and does not have the Exclude policy assigned. • R-Shield compliant: If the changes in backup data size are below the threshold defined in the R-Shield policy, and no malware or ransomware is detected.
HYCU Controllers	Shows the number of available and unavailable HYCU controllers in your data protection environments.
Backups	Shows the percentage of successful and unsuccessful backups in your data protection environments for the last 7 days, the number of successful backups, and the

Console widget	Description
	number of virtual machines with the DR-ready status. You can safely ignore the DR-ready label if you do not plan to use the SpinUp functionality.
Shares	Shows the percentage of protected and unprotected file shares and the percentage of R-Shield compliant and non-compliant file shares in your data protection environments. A file share is considered: • Protected: If it has at least one valid backup available and does not have the Exclude policy assigned. • R-Shield compliant: If the changes in backup data size are below the threshold defined in the R-Shield policy, and no malware or ransomware is detected.
Targets	Shows the list of all targets in your data protection environments, and the information on how much space is used and available for storing data on each target and on all targets in your data protection environments combined.
Policies	Shows the number of entities in your data protection environments, the number of entities that have no policy assigned, and the number of entities that are compliant and non-compliant with the RPO and RTO set in their assigned policy. The number of compliant and non-compliant entities for specific policies is also shown.
M365 / G Workspace	Shows the overview of protected users, Microsoft 365 SharePoint sites, Groups and Teams, and Google Workspace Shared Drive files in your data protection environments. For users, the total number of protected emails, files, contacts, calendar items, and tasks is also displayed.

Viewing entity data

If you have one or more Hybrid Cloud Edition controllers added to HYCU Manager, you can view data protection information received from these HYCU

controllers for each related virtual machine, application, file share, server, volume group, and bucket.

Depending on what kind of data protection information you want to view, access one of the following panels:

- Accessing the Applications panel

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

- Accessing the Virtual Machines panel

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

- Accessing the Shares panel

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

- Accessing the Volume Groups panel

  To access the Volume Groups panel, in the navigation pane, click **Volume Groups**.

- Accessing the Buckets panel

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

> **Note** HYCU Manager performs the automatic synchronization of entities every five minutes. However, you can at any time update the list of entities also manually by clicking **Refresh** in the corresponding panel.

Within each panel, besides viewing data protection information related to your entities, you can also do the following:

I want to...	Instructions
Navigate directly from HYCU Manager to the web user interface of the HYCU backup controller containing the entity whose restore point I select. **Note** By navigating to	To navigate to the web user interface of the HYCU backup controller containing the preferred restore point, follow these steps: 1. In the Applications, Virtual Machines, Shares, Volume Groups, or Buckets panel, click the entity to whose restore point you want to navigate. The Detail view appears at the bottom of the screen.

I want to...	Instructions
the specific web user interface, you will be able not only to perform restore operations related to the selected restore point, but also all other operations that are available in your data protection environment.	2. In the Detail view, select the preferred restore point. 3. Click ➪ **Navigate to Restore Point**. You are automatically directed to the corresponding HYCU web user interface. For instructions on how to sign in to HYCU, see "Signing in to HYCU" on page 68.
Apply different types of filters to entities, including filtering them by the HYCU controller to which they belong.	To filter entities by the HYCU controller, follow these steps: 1. In the selected panel, click ▽ **Filters**. 2. In the side panel that opens, from the HYCU Controllers drop-down menu, select the preferred Hybrid Cloud Edition controller. 3. Click **Apply Filters**. For more details on filtering, see "Filtering and sorting data" on page 406.
Export data to a file in JSON or CSV format.	For details on how to do this, see "Exporting the contents of the panel" on page 416.

Viewing events

You can use the Events panel to view the following:

- All events that occurred in HYCU Manager.
- All events that are associated with Hybrid Cloud Edition controllers that you added to HYCU Manager.

You can check details about the selected event, list events that match the specified filter, configure HYCU to send notifications when events occur, export the contents of the panel to a file in JSON or CSV format, and enable the purge of events.

Accessing the Events panel

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

I want to...	Procedure
View events and check details about the selected event.	"Managing HYCU events" on page 387
Apply filters to events. **Note** You can filter events by where they originated—in HYCU Manager or in connection with any of the added Hybrid Cloud Edition controllers.	"Filtering and sorting data" on page 406
Configure HYCU to send notifications when events occur.	"Configuring event notifications" on page 388
Export event data.	"Exporting the contents of the panel" on page 416
Enable the purge of events.	"Enabling the purge of events and jobs" on page 391

Using HYCU Manager reports

The HYCU Manager reports allow you to analyze data that is specific to one or more of your monitored HYCU controllers, providing you with detailed information about particular data protection areas. By using the HYCU Manager reports, you are always up-to-date with the changes in your data protection environments, which allows you to identify and prevent potential issues.

After you create a report as described in "Creating reports" on the next page, you can continue as follows:

- Schedule reports. For details, see "Scheduling reports" on the next page.

  Note When scheduling the reports, you can also choose to send them by email.

- Generate reports. For details, see "Generating reports" on page 577.

- Download reports. For details, see “Downloading and viewing reports” on page 577.

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

Creating reports

Procedure

1. In the Reports panel, click **New**.
2. Enter a name for your report, and, optionally, its description.
3. From the HYCU Controllers drop-down menu, select one or more HYCU controllers for which you want to create a report.

   If you want to create the report for all HYCU controllers in HYCU Manager, select **Select all**.
4. From the Parameters drop-down menu, select one or more parameters that you want to include in the report. For example, if you select Applications, the report will include all relevant application data, enabling you to view the complete list of applications for the selected HYCU controllers.

   If you want to include all available parameters in the report, select **Select all**.
5. In the Time Range field, specify the time range for the report. You can select one of the predefined time ranges, or select **Custom**, and then specify the start date and hour and the end date and hour for the report.
6. In the Number of Report Versions field, specify how many versions of the report you want to keep.
7. Click **Save**.

You can later do the following:

- Edit any of the existing reports by clicking **Edit** and making the required modifications.
- Delete any of the existing reports by clicking **Delete**.

Scheduling reports

You can use scheduling to generate reports automatically at a particular time each day, week, or month. You can download these reports or schedule them to

be delivered by email.

Prerequisite

For sending reports by email: An SMTP server must be configured. For details, see “Configuring email delivery” on page 512.

Procedure

1. In the Reports panel, from the list of available reports, select the one that you want to be generated on a regular basis, and then click **Scheduler**.
2. In the Schedule Date field, specify the date and time when you want the report generation to begin.
3. From the Interval drop-down menu, select how often you want the reports to be generated (daily, weekly, or monthly).
4. *Only if you want to schedule the automatic delivery of the reports to email recipients.* Do the following:
   a. In the Email Address field, enter one or more email recipients that should receive the reports. If you are entering more than one email address, make sure to press the spacebar after entering each one.
   b. In the Email Text field, enter the custom information to be included in the report email. For example, you can specify the purpose of the report, list further actions to be done by the recipient, or add contact information for additional report details.
5. Click **Schedule**.

> Tip The reports that are scheduled are marked by in the Scheduled column of the Reports panel.

You can later do the following:

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

Generating reports

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

Procedure

1. In the Reports panel, from the list of available reports, select the one that you want to generate.
2. In the Detail view that appears at the bottom of the screen, click **Generate**.

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

You can delete the saved report versions that you do not need anymore. To do so, select the report version that you want to delete, and then click **Delete**.

Downloading and viewing reports

Procedure

1. In the Reports panel, from the list of available reports, select the report that you want to download and view.
2. In the Detail view that appears at the bottom of the screen, select the preferred report version.
3. Click **Download**.

The report is automatically downloaded to your default download location (in the XLSX format) and you can view it.

Performing administration tasks

After you deploy the HYCU virtual appliance in HYCU Manager mode, you can perform various administration tasks through the **Administration** menu.

> **Note** The procedures for administering HYCU deployed in the HYCU Manager mode are the same as for HYCU deployed in the HYCU Backup Controller mode. Therefore, in most cases, you can follow the same instructions.

> Keep in mind that a varied set of administration tasks is available depending on the selected deployment mode.

I want to...	Procedure
Add an Azure service principal to be able to monitor my HYCU for Azure data protection environment.	“Adding an Azure service principal” on page 468
Integrate HYCU Manager with identity providers.	“Integrating HYCU with identity providers” on page 471
Configure log file settings to troubleshoot problems if HYCU does not perform as expected.	“Setting up logging” on page 493
Change network settings.	“Changing network settings” on page 496
Securely store, access, and manage my credentials (secrets) by employing the Conjur secrets management solution.	“Managing secrets” on page 509
Configure an SMTP server.	“Configuring email delivery” on page 512
Upgrade HYCU Manager to a new available version. **Important** Before upgrading, make sure you have added the source where your HYCU Manager virtual machine resides as described in “Adding sources” on page 73.	“Upgrading HYCU” on page 515
Set power options for the HYCU controllers so that their activities are suspended or resumed.	“Setting power options for the HYCU controllers” on page 583
Upgrade HYCU controllers from HYCU Manager.	“Upgrading HYCU controllers” on page 584
Access the HYCU Manager virtual machine by using SSH.	“Accessing the HYCU backup controller virtual machine by

I want to...	Procedure
	using SSH" on page 535
Configure the SSL certificate.	"Configuring SSL certificates" on page 524
Manage HYCU Manager users.	"Managing users" below

In addition, you can use the HYCU REST API Explorer. For details, see "Using the HYCU REST API Explorer" on page 560.

Managing users

You can use the Manage Users dialog box to give the specified users access to HYCU Manager. Managing users includes creating, editing, deleting, and activating or deactivating users.

> Accessing the User Management dialog box
>
> To access the User Management dialog box, from the ⚙ **Administration** menu, select **User Management**.

Creating a new user

Consideration

Only if you plan to add an LDAP group. If the LDAP group that you add contains a user with the at sign (@) in their name, such a user will not be able to sign in to HYCU Manager.

Procedure

1. In the User Management dialog box, click ✚ **New**.
2. Depending on what kind of user you are adding, enter one of the following:
   - *For a HYCU user, an AD user, an OIDC user, an OIDC group, or an LDAP user:* User name

     > ⓘ **Important** Keep in mind the following:
     > - *For an AD user:* When entering a name, make sure it complies with the SAM account name limitations—name length may not exceed 20 characters and contain any of the following characters: "/ \ [ ] : ; | = , + * ? < >. In addition, HYCU does not allow the at sign (@) in the name.

> If your environment requires it, these limitations can be overridden by editing the `ad.username.filter.regex` configuration setting. However, this is not supported and could cause authentication issues. For details on how to customize HYCU configuration settings, see "Customizing HYCU configuration settings" on page 628.
>
> - *For an LDAP user:* HYCU does not allow the at sign (@) in the name.

- *For an AD group or an LDAP group:* Common name

3. From the Authentication Type drop-down menu, select one of the following authentication types, and then follow the instructions:

Authentication type	Instructions
HYCU	a. From the Language drop-down menu, select the preferred language for the user. b. In the Name field, enter a display name for the user. c. *Optional.* In the Email field, enter the email address of the user. d. In the Password field, enter the user password. **Note** The minimum password length is six characters.
OIDC User	a. From the Language drop-down menu, select the preferred language for the user. b. From the Identity Provider drop-down menu, select the identity provider. c. In the Identity Provider User ID field, enter the ID of the identity provider user. **Note** Depending on your identity provider, the user ID corresponds to the following: • *Active Directory Federation Services:* Object GUID • *Google:* User email address • *Keycloak*: User ID

Authentication type	Instructions
	• *Microsoft:* Object ID • *Okta:* Part of the URL when you navigate to the user's profile For details, see the respective identity provider documentation.
OIDC Group	a. From the Language drop-down menu, select the preferred language for the group. b. From the Identity Provider drop-down menu, select the identity provider. c. In the Identity Provider Group ID field, enter the ID of the identity provider group. Note Depending on your identity provider, the group ID corresponds to the following: • *Active Directory Federation Services:* Object GUID • *Keycloak*: Group ID • *Microsoft:* Group Object IDs • *Okta:* Group name For details, see the respective identity provider documentation.
AD User	a. From the Language drop-down menu, select the preferred language for the user. b. From the Identity Provider drop-down menu, select the Active Directory the AD user belongs to.
AD Group	a. From the Language drop-down menu, select the preferred language for the user. b. From the Identity Provider drop-down menu, select the Active Directory the AD group belongs to.
LDAP User	a. From the Language drop-down menu, select the

Authentication type	Instructions
	preferred language for the user. b. From the Identity Provider drop-down menu, select the identity provider that the LDAP user belongs to.
LDAP Group	a. From the Language drop-down menu, select the preferred language for the group. b. From the Identity Provider drop-down menu, select the identity provider that the LDAP group belongs to.

4. *Only if you are adding a HYCU user, an AD user, an AD group, an LDAP user, or an LDAP group.* Use the **Two-factor authentication** switch if you want to enable two-factor authentication for the user, and then select one of the following two-factor authentication methods:
   - **Time-based one-time password**

     This option enables the use of a time-based one-time password (OTP) generated by an OTP application. The user needs to set up an OTP during the first sign-in after two-factor authentication is enabled.
   - **FIDO**

     This option enables the use of an authenticator complying with FIDO protocols (FIDO authenticator). The user needs to register a FIDO authenticator. For details, see "Managing FIDO authenticators" on page 544.
5. *Only if you enabled two-factor authentication.* To prevent the user from disabling two-factor authentication, make sure the **User cannot disable two-factor authentication** check box is selected. If you clear the check box, the user can disable two-factor authentication. An infrastructure group administrator can disable two-factor authentication even if this option is enabled.

   > **Note** If a user disables two-factor authentication, the administrator is notified with a security warning.
6. Click **Save**.

The user is added to the list of all users.

You can later do the following:

- Edit any of the existing HYCU or identity provider users by clicking **Edit** and making the required modifications. Keep in mind that the built-in user, AD users, and AD groups cannot be edited.
- Enable or disable specific users from signing in to HYCU. For details, see "Performing administration tasks" on page 577.
- Delete any of the existing users by clicking **Delete**. Keep in mind that the built-in user cannot be deleted.

Setting power options for the HYCU controllers

You can set power options for the HYCU controllers so that their activities are suspended or resumed.

Limitation

You can set the power options only for Hybrid Cloud Edition controllers.

Procedure

1. Select one or more HYCU controllers for which you want to set the preferred power option.
2. Click **Power Options**.
3. Depending on what you want to do, select one of the following power options:

Power option	Description
Suspend all	Pauses all HYCU controller activities. If you want the HYCU controller activities to automatically resume after a specified amount of time, in the Auto Resume After field, specify the number of hours (1–168) to pass before the activities are resumed. All currently running jobs are allowed to complete normally. All jobs that are in the queue will start when the HYCU controller is resumed. While activities are paused, you cannot start any new backup jobs. > **Tip** After you suspend all HYCU controller activities, the status icon is shown if some jobs are still running on the HYCU controller. You can hover over the icon to see the exact number of currently

Power option	Description
	running jobs.
Suspend cleanup	Pauses the cleanup of targets and, if enabled, the purge of events and jobs. The snapshot cleanup is not affected.
Resume	Allows HYCU controller activities to continue.

4. Click **Save**.

Upgrading HYCU controllers

HYCU Manager enables you to upgrade the HYCU controllers when a new software release version is available, or to apply an update to the current product version of the HYCU controllers.

Prerequisites

- The version of both HYCU Manager and the HYCU backup controller must be 5.2.0 or later.
- The activities on the HYCU controllers that you want to upgrade must be suspended and no jobs may be running. For instructions, see “Setting power options for the HYCU controllers” on the previous page.
- Back up or create snapshots of your HYCU controllers before upgrading them. You can remove these snapshots after completing the upgrade.

Limitation

You can upgrade only Hybrid Cloud Edition controllers.

Procedure

1. Select one or more HYCU controllers that you want to upgrade.
2. Click **Upgrade**.
3. Depending on whether you want to upgrade the HYCU controllers to a new software release version or to apply a HYCU update to the current product version, select one of the following options:
   - **Upgrade to a Newer Version**:
     a. Select how you want to perform the upgrade, and then follow these steps:

- *From the public HYCU repository:*
  a. From the Upgrade From drop-down menu, select **HYCU repository**.
  b. From the Available Versions drop-down menu, select the version to which you want to upgrade HYCU controllers.
  c. Select the **Download image to HYCU backup controller and perform upgrade locally** check box if your network capacity is insufficient for an uninterrupted upgrade procedure or if you experience frequent network outages.
- *From a URL link:*
  a. From the Upgrade From drop-down menu, select **URL**.
  b. In the URL field, enter the full path to the upgrade image file.

  > Example `https://myserver.com/images/hycu-5.3.0-1234.hyimg`

  c. Select the **Download image to HYCU backup controller and perform upgrade locally** check box if your network capacity is insufficient for an uninterrupted upgrade procedure or if you experience frequent network outages.
- *From an image file from your local machine:*
  a. From the Upgrade From drop-down menu, select **Upload file**.
  b. Under Update File, click **Choose a file** to browse for the upgrade image file on your machine, or drag and drop the upgrade image file into the dialog box.

b. Click **Software Upgrade**.
c. Click **Yes** to confirm that you want to upgrade the HYCU controllers.

- **Apply Update**:
  a. Under Update File, click **Choose a file** to browse for the update package (in the ZIP format), or drag and drop the update package into the dialog box.
  b. Click **Update**.

Chapter 15

Protecting data across on-premises and cloud environments

The SpinUp functionality ensures business continuity of your data protection environment across different infrastructures. You can ensure data resilience by migrating virtual machines across the on-premises and cloud (AWS, Google Cloud, Azure, or Azure Government) infrastructures. In the event of a disaster in your on-premises environment, the SpinUp functionality provides disaster recovery of data to cloud.

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

- “Protecting data across on-premises and AWS environments” below
- “Protecting data across on-premises and Google Cloud environments” on page 599
- “Protecting data across on-premises and Azure environments” on page 609
- “Protecting data across on-premises and Azure Government environments” on page 619

Protecting data across on-premises and AWS environments

You can use the SpinUp functionality to migrate protected data across the on-premises and Amazon Web Services (AWS) environments. In the event of a disaster, it provides disaster recovery of data to AWS.

Depending on what you want to do, see one of the following:

I want to...	Instructions
Migrate protected data across the on-	“Migrating virtual machines across

I want to...	Instructions
premises and AWS environments.	different environments" on the next page
Perform disaster recovery of data to AWS.	"Performing disaster recovery of data to AWS" on page 597

Prerequisites

- You must have an active subscription for HYCU R-Cloud. For instructions, see HYCU R-Cloud documentation.
- An AWS user account must be added to HYCU. For instructions, see "Adding an AWS user account" on page 464.
- A HYCU account must be added to HYCU. For instructions, see "Adding a HYCU account" on page 470.

Migrating virtual machines across different environments

You can migrate protected data across the on-premises and AWS environments as follows:

- "Migrating data to cloud" below
- "Migrating data from cloud" on page 595

Migrating data to cloud

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

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

Prerequisites

- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected.
- In the HYCU R-Cloud web user interface, make sure that the AWS user account is granted the Administrator role in the Subscription context.

- Your AWS user account must have the required permissions for migrating data to AWS. For a list of these permissions, see "Setting permissions in AWS" on page 593.
- *For Windows virtual machines:* In AWS, the required permissions must be specified in your IAM policy for VM Import/Export. Make sure to use `hycu-tmp` instead of `mys3bucket`, `disk-image-file-bucket`, and `export-bucket` in the sample policies. For instruction on how to specify required permissions for VM Import/Export, see AWS documentation.
- You must own a premium tier Platform license. For details, see "Licensing" on page 488.

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for migrating data.
- *For Nutanix clusters:* You cannot migrate volume groups.
- *For vSphere environments:*
  - You cannot migrate virtual machine templates.
  - Migrating data from snapshots is not supported.
- *For Windows virtual machines:* If the virtual machine type uses the Intel 82599 VF interface for enhanced networking, the AWS import image mode is the only supported adaptation mode.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for migrating data.
- *For Linux virtual machines:* If the Prepare OS adaptation overlay task completes with a warning due to missing kernel modules, make sure to select a virtual machine type that does not require them.

Depending on whether you want to migrate virtual machine or application data to cloud, access one of the following panels:

- #### Accessing the Virtual Machines panel

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

- Accessing the Applications panel

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

Procedure

1. In the Virtual Machines or Applications panel, select the entity that you want to migrate.
2. In the Detail view that appears at the bottom of the screen, select the virtual machine or application restore point that you want to use for the migration.

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

3. Click **SpinUp to Cloud**.
4. Select **SpinUp VM to AWS**, and then click **Next**.
5. From the HYCU Account drop-down menu, select the HYCU account.
6. From the AWS User Account drop-down menu, select the AWS user account.

   > **Note** By default, the AWS account to which the selected AWS user account belongs and to which the entity will be migrated is displayed.

7. From the AWS Account ID, select the AWS account ID.
8. From the Region drop-down menu, select the AWS region to which you want to migrate the entity.
9. From the Availability Zone drop-down menu, select the preferred availability zone within the selected AWS region.
10. Click **Next**.
11. From the SpinUp From drop-down menu, select which tier you want to use for the migration. Your restore point can contain one or more tiers among which you can select:
    - **Automatic**: Ensures the fastest migration of data to cloud.
    - **Backup**
    - **Copy**
    - **Archive**
    - **Snapshot**
12. In the New VM Name field, enter a name for the migrated virtual machine instance.

> **Important** Make sure the migrated virtual machine instance name is unique.

13. In the vCPU Threads field, enter the number of CPUs for the migrated virtual machine instance multiplied by the number of cores per CPU and the number of threads per core.
14. In the Memory field, set the amount of memory (in GiB) for the migrated virtual machine instance. The default value is the amount of memory in GiB of the original virtual machine.
15. From the Virtual Machine Type drop-down menu, select the machine type for the migrated virtual machine instance.

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

16. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the migrated virtual machine.
17. Under Network Interfaces, do the following:
    a. Click **Add Network Interface**.
    b. From the Subnets drop-down menu, select the subnet.
    c. From the Security Groups drop-down menu, select the AWS security group.
    d. Under Public Address Type, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	The network interface does not use a public IP address.
Auto-assign	The network interface uses an automatically allocated public IP address. **Note** Auto-assign will not work if the Auto-assign public IPv4 address on a subnet option is set to `No` or if more than one network interface is specified.

Elastic IP (Reserved)	The network interface uses an elastic public IP address that was reserved in Amazon EC2 in advance.
Elastic IP (New)	The network interface uses a new elastic public IP address. **Note** Allocation of the IP address in Amazon EC2 is performed at the very beginning of the migration. If the allocation fails, the migration task is terminated without being logged.

e. Under Private Address Type, select the private IP address for the network interface. You can select between the following options:

Option	Description
Auto-assign	The network interface uses an automatically allocated private IP address.
Custom	The network interface uses a private IP address that you define. **Important** Using this option might result in IP address conflicts.

f. Click **Add**.

You can also edit any of the existing network interfaces (click **Edit** and make the required modifications) or delete the ones that you do not need anymore (click **Delete**). Keep in mind that you cannot migrate the virtual machine without a network interface.

18. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
19. *Only if the platform readiness check was not performed for the selected restore point.* Use the **Adapt OS for migration** switch to apply the configuration changes that are required for the migration to cloud.

    If you enable this option for Linux virtual machines, HYCU applies the following configuration changes:

- Includes the required kernel modules into `initramfs`.
- Enables the serial console on serial port 1.
- Changes the network configuration to use the DHCP.

20. *Only if you enabled the Adapt OS for migration option for Windows virtual machines:* Under Adaptation mode, select one of the following modes:

Adaptation mode	Select this mode if you want to...
Use HYCU mode	Enable HYCU to apply the configuration changes that are required for the migration to cloud. This is the recommended mode in which adapting the operating system for the migration takes less time to complete compared to the AWS import image mode. If you enable this option, HYCU applies the following configuration changes: • Installs the required drivers. • Cleans up the existing devices to facilitate the discovery of new devices. • Enables the EMS console on serial port 1. • Resets the TCP/IP and Winsock stacks to use the DHCP.
Use AWS import image mode	Use the AWS import image to apply the configuration changes that are required for the migration to cloud. For details on the AWS import image, see AWS documentation. This mode allows you to replace the existing OS license.

21. *Only if you selected the AWS import image mode.* Under Operating System License, select one of the following options:

OS license option	Select this option if you want to...
Keep existing license	Keep the existing OS license on the migrated virtual machine instance. **Important** Make sure that the existing license is applicable also in

OS license option	Select this option if you want to...
	AWS.
Available only for the Windows Server OS. **Replace existing license with AWS license**	Replace the existing OS license with an AWS license on the migrated virtual machine instance.

22. Click **SpinUp**.

The Migration to cloud job starts. When it finishes successfully, you can check the migrated virtual machine instance in the Instances panel in HYCU R-Cloud. For details, see HYCU R-Cloud documentation.

After migrating data to cloud

- *For Windows virtual machines:* If you decided to keep the existing OS license on the migrated virtual machine instance, reactivate the Windows license.
- Enable the protection of the migrated virtual machine instances by using HYCU R-Cloud. For details, see HYCU R-Cloud documentation.

Setting permissions in AWS

The following table shows the minimum set of permissions that are required for migrating data to AWS:

Service	Required permissions
EC2	ec2:AllocateAddress ec2:AssociateAddress ec2:AttachNetworkInterface ec2:AttachVolume ec2:CopySnapshot ec2:CreateNetworkInterface ec2:CreateSecurityGroup ec2:CreateTags ec2:CreateVolume ec2:DeleteSnapshot ec2:DeleteVolume ec2:DeregisterImage ec2:DescribeAddresses ec2:DescribeAvailabilityZones

Service	Required permissions
	ec2:DescribeImages ec2:DescribeInstanceAttribute ec2:DescribeInstances ec2:DescribeInstanceTypeOfferings ec2:DescribeInstanceTypes ec2:DescribeKeyPairs ec2:DescribeMacHosts ec2:DescribeNetworkAcls ec2:DescribeNetworkInterfaceAttribute ec2:DescribeNetworkInterfaces ec2:DescribeRegions ec2:DescribeSecurityGroups ec2:DescribeSnapshotAttribute ec2:DescribeSnapshots ec2:DescribeSubnets ec2:DescribeVolumes ec2:DescribeVpcAttribute ec2:DescribeVpcEndpoints ec2:DescribeVpcEndpointServices ec2:DescribeVpcs ec2:DescribeVpnConnections ec2:DetachNetworkInterface ec2:DetachVolume ec2:GetConsoleOutput ec2:GetInstanceUefiData ec2:ModifyAddressAttribute ec2:ModifyImageAttribute ec2:ModifyInstanceAttribute ec2:ModifyNetworkInterfaceAttribute ec2:RegisterImage ec2:RunInstances
Elastic Block Store (EBS)	ebs:CompleteSnapshot ebs:PutSnapshotBlock ebs:StartSnapshot

Migrating data from cloud

You can migrate virtual machine instances from cloud by using the SpinUp functionality.

Limitation

You cannot migrate virtual machine instances from cloud to an Azure Local environment.

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

Procedure

1. In the Virtual Machines panel, click **SpinUp VM from Cloud**.
2. Select **SpinUp VM from AWS**, and then click **Next**.
3. From the AWS user account drop-down menu, select the AWS user account.

   > **Note** By default, the AWS account to which the selected AWS user account belongs and from which the virtual machine instance will be migrated is displayed.

4. From the HYCU account drop-down menu, select the HYCU account.
5. From the Virtual machine drop-down menu, select the virtual machine instance that you want to migrate.
6. From the Checkpoint drop-down menu, select the checkpoint (restore point) from which you want to migrate virtual machine instance data.

   > **Important** Before you migrate a virtual machine that you want to migrate from the cloud, must be backed up.

7. Click **Next**. The VM Settings dialog box opens.
8. From the Storage container drop-down menu, select where you want to migrate the virtual machine instance.
9. In the New VM name field, enter a name for the migrated virtual machine.
10. *Only if the virtual machine instance that you are migrating was created in the on-premises environment, migrated to cloud, and now you are migrating it back to the on-premises environment.* If you want the migrated virtual machine to have the same virtual machine settings as it had in the on-premises environment, enable the **Keep original on-premises settings** option, and

then continue with step 13.

Otherwise, leave the Keep original on-premises settings option disabled and continue with the next step.

11. Specify the following values for the migrated virtual machine:

Value	Description
vCPU(s)	Number of virtual CPUs.
Cores per vCPU	Number of cores per virtual CPU.
Memory	Amount of memory (in GiB).

> **Note** The default values are the ones that the virtual machine had in the environment in which it was created, either in the on-premises or cloud one.

12. Under Network adapters, depending on your data protection needs, do one of the following:
    - Add one or more network adapters:
        a. Click **Add Network Adapter**. The New Network Adapter dialog box opens.
        b. From the Network drop-down menu, select the network.
        c. Click **Save**.
    - Edit any of the existing network adapters to connect the virtual machine to a different network. To do so, select a network adapter, click **Edit**, and make the required modification.
    - Delete any of the existing network adapters by selecting it, and then clicking **Delete**. If you delete all the existing network adapters, your virtual machine will be migrated without network connectivity.
13. Use the **Power virtual machine on** switch if you want to turn the migrated virtual machine on after the migration.
14. Click **SpinUp**.

The Migration from cloud job starts. When it finishes successfully, you can view the migrated virtual machine in the Virtual Machines panel.

After migrating data from cloud

- *For virtual machines on a Nutanix AHV cluster:* Make sure that the latest version of NGT is installed on the virtual machine. For instructions, see

Nutanix documentation.

- *For virtual machines on a Nutanix ESXi cluster:* Make sure that the latest versions of VMware Tools and NGT are installed on the virtual machine. For instructions, see Nutanix and VMware documentation.
- *For virtual machines in a vSphere environment:* Make sure that the latest version of VMware Tools is installed on the virtual machine. For instructions, see VMware documentation.
- *For Windows virtual machines:* Reactivate the Windows licenses.
- *For Linux virtual machines:* If a virtual machine on a Nutanix ESXi cluster or in a vSphere environment does not boot, change the controller type from SCSI to SATA, and then install the necessary SCSI drivers to switch back to SCSI.
- *Only if you migrated virtual machines without network connectivity.* Make sure to configure the network settings on the virtual machine.
- Enable protection of the migrated data. For details on how to do this, see "Protecting virtual machines" on page 154 and "Protecting applications" on page 252.

Performing disaster recovery of data to AWS

You can perform disaster recovery of data from the on-premises environment to AWS in the event of a disaster.

Prerequisites

- You must have the HYCU virtual appliance image for AWS. To obtain the image, on AWS Marketplace, browse the AMI Catalog for the following Amazon Machine Image (AMI):

Image name	The image name is represented in the following format: `hycu-<Version>-<Revision>` For example: `hycu-5.3.0-3634`
Owner	The owner is represented by the following AWS account ID: 787223699828

  For instructions, see AWS documentation.
- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected and must have

the DR-ready status. For more information, see “SpinUp specifics” on page 170.

Considerations

- When the HYCU backup controller is deployed in AWS, changing network settings is prevented in HYCU.
- Make sure the imported target is in the region to which you plan to migrate your virtual machines. This ensures the disaster recovery process is as fast and as cost-effective as possible.
- After you deploy the HYCU backup controller and use it to perform disaster recovery, you can keep the HYCU backup controller to stay prepared for disaster recovery in the future.

Procedure

1. Deploy a HYCU backup controller. To do so, select the HYCU virtual appliance image in the AWS AMI Catalog, and then click **Launch Instance with AMI**. For instructions, see AWS documentation.
2. In AWS, create a new firewall rule to allow ingress network traffic through TCP port 8443 from the entire subnetwork to which the HYCU backup controller belongs. For instructions, see AWS documentation.
3. Sign in to the HYCU web user interface by specifying the following URL:

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

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

   > **Important** The credentials you provided in AWS during virtual machine instance creation cannot be used to sign in to HYCU and perform disaster recovery of data to AWS. For details on what credentials you can use to sign in to HYCU or to access the HYCU backup controller by using SSH, see “Signing in to HYCU” on page 68 or “Accessing the HYCU backup controller virtual machine by using SSH” on page 535.

4. Add an AWS user account. For instructions, see “Adding an AWS user account” on page 464.
5. Import the Amazon S3 target on which your backup data is stored to HYCU:
   a. In the Targets panel, click **Import**. The Import Target dialog box opens.

b. From the Type drop-down menu, select **Amazon S3 / S3 Compatible**.
c. In the Service endpoint field, enter the service endpoint URL.
d. In the Bucket name field, enter the Amazon S3 bucket name as it was specified in the original target configuration.
e. In the Access key ID field, enter the access key ID of your AWS user account.
f. In the Secret access key, enter the secret access key of your AWS user account.
g. Enable the **Path style access** switch if you want HYCU to use a path-style URL (`https://s3.amazonaws.com/<BucketName>`) to access the bucket. HYCU by default uses a virtual-hosted–style URL (`https://<BucketName>.s3.amazonaws.com`).
h. Click **Next**. The Import Backup Catalog dialog box opens.
i. Select the HYCU backup controller whose backup data you want to import, and then click **Next**.
j. In the Multiple Targets dialog box, one or more targets that store backup data are displayed. If any additional targets are found, select them one by one and specify the values so that they match the original target configuration. For each target, click **Validate** to check the configuration.
k. After you validate all the targets, click **Import**.

6. Migrate your virtual machines or applications to cloud. For instructions, see "Migrating data to cloud" on page 587.

Protecting data across on-premises and Google Cloud environments

You can use the SpinUp functionality to migrate protected data across the on-premises and Google Cloud environments. In the event of a disaster, it provides disaster recovery of data to Google Cloud.

Depending on what you want to do, see one of the following:

I want to...	Instructions
Migrate protected data across the on-premises and Google Cloud environments.	"Migrating virtual machines across different environments" on the next page

I want to...	Instructions
Perform disaster recovery of data to Google Cloud.	“Performing disaster recovery of data to Google Cloud” on page 607

Prerequisites

- You must have an active subscription for HYCU R-Cloud. For instructions, see HYCU R-Cloud documentation.
- A Google Cloud service account must be added to HYCU. For instructions, see “Adding a Google Cloud service account” on page 466.
- A HYCU account must be added to HYCU. For instructions, see “Adding a HYCU account” on page 470.

Migrating virtual machines across different environments

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

- “Migrating data to cloud” below
- “Migrating data from cloud” on page 604

Migrating data to cloud

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

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

Prerequisites

- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected.
- You must own a premium tier Platform license. For details, see “Licensing” on page 488.

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for migrating data.
- *For Nutanix clusters:* You cannot migrate volume groups.
- *For vSphere environments:*
  - You cannot migrate virtual machine templates.
  - Migrating data from snapshots is not supported.
- You cannot select a virtual machine type that uses the Hyperdisk Balanced disk type for the migrated virtual machine instance.

Consideration

If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for migrating data.

Depending on whether you want to migrate virtual machine or application data to cloud, access one of the following panels:

- #### Accessing the Virtual Machines panel

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

- #### Accessing the Applications panel

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

Procedure

1. In the Virtual Machines or Applications panel, select the entity that you want to migrate.
2. In the Detail view that appears at the bottom of the screen, select the virtual machine or application restore point that you want to use for the migration.

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

3. Click **SpinUp to Cloud**.

4. Select **SpinUp VM to Google Cloud**, and then click **Next**.
5. From the HYCU Account drop-down menu, select the HYCU account.
6. From the Cloud Account drop-down menu, select the Google Cloud service account to which the project where you want to migrate the virtual machine is linked.
7. From the Project, Target Region, and Target Zone drop-down menus, select the required values.
8. Click **Next**.
9. From the SpinUp From drop-down menu, select which tier you want to use for the migration. Your restore point can contain one or more tiers among which you can select:
   - **Automatic**: Ensures the fastest migration of data to cloud.
   - **Backup**
   - **Copy**
   - **Archive**
   - **Snapshot**
10. In the New VM Name field, enter a name for the migrated virtual machine instance.

    > ⓘ Important Make sure the migrated virtual machine instance name is unique.

11. In the vCPU Cores field, enter the number of virtual CPUs for the migrated virtual machine multiplied by the number of cores per virtual CPU.
12. In the Memory field, set the amount of memory (in GiB) for the migrated virtual machine instance. The default value is the amount of memory in GiB of the original virtual machine.
13. From the Virtual Machine Type drop-down menu, select the machine type for the migrated virtual machine instance.

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

14. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the migrated virtual machine.

15. Under Network Interfaces, the default network interface is displayed and you can check to which network it is assigned (based on the selected project and region). If required, you can also modify network settings.

    Modifying network settings

    Depending on your data protection needs, you can leave the default network interface or do one of the following:

    - Add a new network interface:
      a. Click **Add Network Interface**.

         Note The maximum number of network interfaces that you can add depends on the selected virtual machine type.

      b. From the Target Networks drop-down menu, select a network to which you want to add the migrated virtual machine instance. You can choose among the networks configured in the selected project and other networks that your cloud account has access to.
      c. Select the external address type for the network interface and, if required, the name of the preferred external IP address resource. For details, see HYCU R-Cloud documentation.
      d. Select the internal address type for the network interface and, if required, depending on the address type, do one of the following:
         - In the Internal Address field, enter the preferred IP address.
         - From the Internal Address drop-down menu, select the name of the preferred internal IP address resource.

         For details, see HYCU R-Cloud documentation.
      e. Click **Add**.
    - Select another network for the existing network interface by selecting it, clicking **Edit** and making the required modifications.
    - Delete the existing network interface by selecting it, and then clicking **Delete**.

16. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
17. *Only if the platform readiness check was not performed for the selected restore point.* Use the **Adapt OS for migration** switch to apply the configuration

changes that are required for the migration to cloud.

If you enable this option, HYCU applies the following configuration changes to the virtual machines:

- *Windows:*
  - Installs the required drivers.
  - Cleans up the existing devices to facilitate the discovery of new devices.
  - Enables the EMS console on serial port 1.
  - Resets the TCP/IP and Winsock stacks to use the DHCP.
- *Linux:*
  - Includes the required kernel modules into `initramfs`.
  - Enables the serial console on serial port 1.
  - Changes the network configuration to use the DHCP.

18. Click **SpinUp**.

The Migration to cloud job starts. When it finishes successfully, you can check the migrated virtual machine instance in the Instances panel in HYCU R-Cloud. For details, see HYCU R-Cloud documentation.

After migrating data to cloud

- Install the Google Compute Engine guest environment on the virtual machine.
- *For Windows virtual machines:* Reactivate the Windows licenses.
- Enable protection of the migrated virtual machines by using HYCU R-Cloud. For details, see HYCU R-Cloud documentation.

Migrating data from cloud

You can migrate virtual machine instances from cloud by using the SpinUp functionality.

Limitation

You cannot migrate virtual machine instances from cloud to an Azure Local environment.

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

Procedure

1. In the Virtual Machines panel, click **SpinUp VM from Cloud**.
2. Select **SpinUp VM from Google Cloud**, and then click **Next**.
3. From the Cloud account drop-down menu, select the Google Cloud service account to which the project containing the virtual machine instance that you want to migrate is linked.
4. From the HYCU account drop-down menu, select the HYCU account.
5. From the Project drop-down menu, select the Google Cloud project to which the virtual machine instance that you want to migrate belongs.
6. From the Virtual machine drop-down menu, select the virtual machine instance that you want to migrate.
7. From the Checkpoint drop-down menu, select the checkpoint (restore point) from which you want to migrate virtual machine instance data.
8. Click **Next**. The VM Settings dialog box opens.
9. From the Storage container drop-down menu, select where you want to migrate the virtual machine instance.
10. In the New VM name field, enter a name for the migrated virtual machine.
11. *Only if the virtual machine instance that you are migrating was created in the on-premises environment, migrated to cloud, and now you are migrating it back to the on-premises environment.* If you want the migrated virtual machine to have the same virtual machine settings as it had in the on-premises environment, enable the **Keep original on-premises settings** option, and then continue with step 13.

    Otherwise, leave the Keep original on-premises settings option disabled and continue with the next step.
12. Specify the following values for the migrated virtual machine:
    - The number of virtual CPUs.
    - The number of cores per virtual CPU.
    - The amount of memory (in GiB).

    > **Note** The default values are the ones that the virtual machine had in the environment in which it was created, either in the on-premises or cloud one.

13. Under Network adapters, depending on your data protection needs, do one of the following:

- Add one or more network adapters:
  a. Click **Add network adapter**. The New Network Adapter dialog box opens.
  b. From the Networks drop-down menu, select the network for the virtual adapter.
  c. Click **Save**.
- Edit any of the existing network adapters to connect the virtual machine to a different network. To do so, select a network adapter, click **Edit**, and make the required modification.
- Delete any of the existing network adapters by selecting it, and then clicking **Delete**. If you delete all the existing network adapters, your virtual machine will be migrated without network connectivity.

14. Use the **Power virtual machine on** switch if you want to turn the migrated virtual machine on after the migration.
15. Click **SpinUp**.

The Migration from cloud job starts. When it finishes successfully, you can view the migrated virtual machine in the Virtual Machines panel.

After migrating data from cloud

- Remove the Google Compute Engine guest environment from the virtual machine.
- *For virtual machines on a Nutanix AHV cluster:* Make sure that the latest version of NGT is installed on the virtual machine. For details on how to do this, see Nutanix documentation.
- *For virtual machines on a Nutanix ESXi cluster:* Make sure that the latest versions of VMware Tools and NGT are installed on the virtual machine. For details on how to do this, see Nutanix and VMware documentation.
- *For virtual machines in a vSphere environment:* Make sure that the latest version of VMware Tools is installed on the virtual machine. For details on how to do this, see VMware documentation.
- *For Linux virtual machines:* If a virtual machine on a Nutanix ESXi cluster or in a vSphere environment does not boot, change the controller type from SCSI to SATA, and then install the necessary SCSI drivers to switch back to SCSI.
- *For Windows virtual machines:* Reactivate the Windows licenses.

- *Only if you migrated virtual machines without network connectivity.* Make sure to configure the network settings on the virtual machine.
- Enable protection of the migrated data. For details on how to do this, see "Protecting virtual machines" on page 154 and "Protecting applications" on page 252.

Performing disaster recovery of data to Google Cloud

You can perform disaster recovery of data from the on-premises environment to Google Cloud in the event of a disaster.

Prerequisites

- You must have a Google Account with the following permissions:
  - To access Google Cloud Storage buckets in the Google Cloud project where you want to deploy your new HYCU backup controller.
  - To deploy Google Compute Engine virtual machine instances to the Google Cloud project where you want to deploy your new HYCU backup controller.
  - To set up a firewall rule in the Google Cloud network where you plan to deploy your new HYCU backup controller.
- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected and must have the DR-ready status. For more information, see "SpinUp specifics" on page 170.

Considerations

- When the HYCU backup controller is deployed in Google Cloud, changing network settings is prevented in HYCU.
- Make sure the imported target is in the region to which you plan to migrate your virtual machines. This ensures the disaster recovery process is as fast and as cost-effective as possible.
- After you deploy the HYCU backup controller and use it to perform disaster recovery, you can keep the HYCU backup controller to stay prepared for disaster recovery in the future.

Procedure

1. Deploy a HYCU backup controller by using the HYCU R-Cloud web user interface. For details on how to do this, see HYCU R-Cloud documentation.
2. In Google Cloud, in the VPC network pane, in the Firewall rules context, create a new firewall rule to allow ingress network traffic through the TCP port 8443 from the entire subnetwork which the HYCU backup controller belongs to. For details on how to do this, see Google Cloud documentation.
3. Sign in to the HYCU web user interface by specifying the following URL:

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

   In this instance, `<IPAddress>` is the external IP address of the newly deployed HYCU backup controller.
4. Add a Google Cloud service account with permissions to access the Google Cloud Storage buckets where backup data of the protected virtual machines is stored. For details on how to do this, see “Adding a Google Cloud service account” on page 466.
5. Import the Google Cloud target with your backup data:
   a. In the Targets panel, click **Import**. The Import Target dialog box opens.
   b. In the Bucket Name field, enter the name as it was specified in the original target configuration.
   c. From the Cloud Account drop-down list, select an imported Google Cloud service account, and then click **Next**.
   d. Click the target name to confirm your selection, and then click **Next**.
   e. In the Multiple Targets dialog box, one or more targets that store backup data are displayed. If any additional targets are found, select them one by one and specify the values so that they match the original target configuration. For each target, click **Validate** to check the configuration.
   f. After you validated all the targets required for your restore, click **Import**.
6. Migrate your virtual machines or applications to cloud. For instructions, see “Migrating data to cloud” on page 600.

Protecting data across on-premises and Azure environments

You can use the SpinUp functionality to migrate protected data across the on-premises and Azure environments. In the event of a disaster in the on-premises environment, it provides disaster recovery of data to Azure.

Prerequisites

- You must have an active subscription for HYCU R-Cloud or HYCU for Azure. For details, see HYCU R-Cloud or HYCU for Azure documentation.
- An Azure service principal must be added to HYCU. For instructions, see "Adding an Azure service principal" on page 468.
- *Only if your cloud data is protected with HYCU R-Cloud.* A HYCU account must be added to HYCU. For instructions, see "Adding a HYCU account" on page 470.
- A storage account that is dedicated exclusively to migration operations must be created in Azure. This storage account must be in the same region and resource group as the virtual machine that you plan to migrate, must have public network access enabled, and its type must be Standard general-purpose v2 or Premium block blobs.

Depending on what you want to do, see one of the following:

I want to...	Instructions
Migrate protected data across the on-premises and Azure environments.	"Migrating virtual machines across different environments" below
Perform disaster recovery of data to Azure.	"Performing disaster recovery of data to Azure" on page 617

Migrating virtual machines across different environments

You can migrate protected data across the on-premises and Azure environments:

- "Migrating data to cloud" on the next page
- "Migrating data from cloud" on page 615

Migrating data to cloud

You can migrate virtual machines, servers, and applications running on virtual machines and servers to Azure by using the SpinUp functionality. Keep in mind that when you migrate an application, the whole virtual machine or server on which this application is running is migrated to cloud.

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

Prerequisites

- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected.
- You must own a premium tier Platform license. For details, see "Licensing" on page 488.

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for migrating data.
- *For Nutanix clusters:* You cannot migrate volume groups.
- *For vSphere environments:*
  - You cannot migrate virtual machine templates.
  - Migrating data from snapshots is not supported.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for migrating data.
- After you migrate data to cloud, an Azure temporary disk is automatically assigned to the migrated virtual machine. This disk is not a managed disk and it is used only for short-term data storage.
- *For virtual machines with secure boot enabled:* Because Azure does not currently support the secure boot feature for virtual machines, after you migrate such a virtual machine to cloud, secure boot cannot be enabled for it.

Depending on whether you want to migrate virtual machine or application data to cloud, access one of the following panels:

- Accessing the Virtual Machines panel

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

- Accessing the Applications panel

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

Procedure

1. In the Virtual Machines or Applications panel, select the entity that you want to migrate.
2. In the Detail view that appears at the bottom of the screen, select the virtual machine or application restore point that you want to use for the migration.

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

3. Click **SpinUp to Cloud**.
4. Select **SpinUp VM to Azure**, and then click **Next**.
5. From the HYCU Account drop-down menu, select the HYCU account.

   > Important If you use HYCU for Azure to protect your data, select **No HYCU account**.

6. From the Service Principal drop-down menu, select the service principal that has access to the required resources.
7. From the Subscription drop-down menu, select the appropriate subscription for the migrated virtual machine.
8. From the Resource Group drop-down menu, select the resource group for the migrated virtual machine.
9. From the Location drop-down menu, select the geographic region for the migrated virtual machine.
10. From the Availability Zone drop-down menu, select the zone for the migrated virtual machine.

    > Note The selected geographic region and the size of the virtual machine determine to which zones you can migrate data. If you do not want to migrate data to any zone, select **None**.

11. Click **Next**.
12. From the SpinUp From drop-down menu, select which tier you want to use for the migration. Your restore point can contain one or more tiers among which you can select:
    - **Automatic**: Ensures the fastest migration of data to cloud.
    - **Backup**
    - **Copy**
    - **Archive**
    - **Snapshot**
13. In the New VM Name field, enter a name for the migrated virtual machine.
14. In the vCPU Cores field, enter the number of virtual CPUs to be assigned to the migrated virtual machine multiplied by the number of cores per virtual CPU.
15. In the Memory field, enter the amount of memory (in GiB) to be assigned to the migrated virtual machine.
16. From the Virtual Machine Type drop-down menu, select the virtual machine type.

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

17. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the migrated virtual machine.
18. Under Network Interfaces, you can view the network interface that will be added to the migrated virtual machine. By default, this is the first network interface from the subscription that you selected for the migrated virtual machine. If required, you can also modify network settings.

    ### Modifying network settings

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

    > Note When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The

> maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add Network Interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow these steps:
  a. *Only if you are adding a network interface.* From the Virtual Network drop-down menu, select the network for the network interface.

     > **Note** The list of available networks includes only the ones within the region you selected for the migrated virtual machine.

  b. Select the subnet to which the network interface should be assigned.
  c. Under Public IP Address Type, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the migrated virtual machine.
Static	A static IP address will be assigned to the network interface on the migrated virtual machine.
Existing	A preferred public IP address resource that you have created in Azure Government will be assigned to the network interface on the migrated virtual machine.

  d. Under Private IP Address Type, select the private IP address for the network interface. You can select between the following options:

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the migrated virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the migrated virtual machine.

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

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

19. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
20. *Only if the platform readiness check was not performed for the selected restore point.* Use the **Adapt OS for migration** switch to apply the configuration changes that are required for the migration to cloud.

    If you enable this option, HYCU applies the following configuration changes to the virtual machines:
    - *Windows:*
      - Installs the required drivers.
      - Cleans up the existing devices to facilitate the discovery of new devices.
      - Enables the EMS console on serial port 1.
      - Resets the TCP/IP and Winsock stacks to use the DHCP.
    - *Linux:*
      - Includes the required kernel modules into `initramfs`.
      - Enables the serial console on serial port 1.
      - Changes the network configuration to use the DHCP.
21. Click **SpinUp**.

The Migration to cloud job starts. When it finishes successfully, you can view the migrated virtual machine in the Instances panel in HYCU R-Cloud or in the Virtual Machines panel in HYCU for Azure. For details, see HYCU R-Cloud or HYCU for Azure documentation.

After migrating data to cloud

- *For Windows virtual machines:* Reactivate the Windows licenses.
- *For Linux virtual machines:* Install the Linux Integration Services for Hyper-V and Azure on the virtual machine. For details, see Microsoft documentation.
- Enable protection of the migrated virtual machines by using HYCU R-Cloud or HYCU for Azure. For details on how to do this, see HYCU R-Cloud or HYCU for Azure documentation.

Migrating data from cloud

You can migrate virtual machines from Azure by using the SpinUp functionality.

Limitations

- Migrating virtual machines with unmanaged disks is not supported.
- Migrating virtual machines with PremiumV2_LRS disks attached is not supported.
- You cannot migrate virtual machine instances from cloud to an Azure Local environment.
- *For Nutanix clusters:* You can migrate Azure Generation 2 virtual machines only to clusters that support UEFI virtual machines.

Consideration

After you migrate data from cloud, the migrated virtual machine does not contain the temporary disk that was automatically assigned to it in Azure.

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

Procedure

1. In the Virtual Machines panel, click **SpinUp VM from Cloud**.
2. Select **SpinUp VM from Azure**, and then click **Next**.
3. From the HYCU Account drop-down menu, select the HYCU account.

   > **Important** If you use HYCU for Azure to protect your data, select **No HYCU account**.

4. From the Service Principal drop-down menu, select the service principal that has access to the required resources.
5. From the Subscription drop-down menu, select the HYCU R-Cloud or HYCU for Azure subscription to which the virtual machine that you want to migrate belongs.
6. From the Resource Group drop-down menu, select the resource group to which the virtual machine that you want to migrate belongs.
7. From the Virtual Machine drop-down menu, select the virtual machine that you want to migrate.

8. From the Checkpoint drop-down menu, select the checkpoint (restore point) from which you want to migrate virtual machine data.
9. From the Storage Account drop-down menu, select the storage account that is dedicated exclusively to migration operations.
10. Click **Next**. The VM Settings dialog box opens.
11. From the Storage Container drop-down menu, select where you want to migrate the virtual machine.
12. In the New VM Name field, enter a name for the migrated virtual machine.
13. *Only if the virtual machine that you are migrating was created in the on-premises environment, migrated to cloud, and now you are migrating it back to the on-premises environment.* If you want the virtual machine to have the same virtual machine settings as it had in the on-premises environment, enable the **Keep original on-premises settings** option, and then continue with step 15.

    Otherwise, leave the Keep original on-premises settings option disabled and continue with the next step.
14. Specify the following values for the migrated virtual machine:
    - The number of virtual CPUs.
    - The number of cores to be assigned to each virtual CPU.
    - The amount of memory (in GiB).

    > Note The default values are the ones that the virtual machine had in the environment in which it was created, either in the on-premises or cloud one.

15. Under Network Adapters, depending on your data protection needs, do one of the following:
    - Add one or more network adapters:
      a. Click **Add Network Adapter**. The Network dialog box opens.
      b. From the Network drop-down menu, select the virtual network for the network adapter.
      c. Click **Add**.
    - Edit any of the existing network adapters to connect the virtual machine to a different network. To do so, select a network adapter, click **Edit**, and make the required modification.

- Delete any of the existing network adapters by selecting it, and then clicking **Delete**. If you delete all the existing network adapters, your virtual machine will be migrated without network connectivity.

16. Use the **Power virtual machine on** switch if you want to turn the migrated virtual machine on after the migration.
17. Click **SpinUp**.

The Migration from cloud job starts. When it finishes successfully, you can view the migrated virtual machine in the Virtual Machines panel.

After migrating data from cloud

- *For virtual machines on a Nutanix AHV cluster:* Make sure that the latest version of NGT is installed on the virtual machine. For details, see Nutanix documentation.
- *For virtual machines on a Nutanix ESXi cluster:* Make sure that the latest versions of VMware Tools and NGT are installed on the virtual machine. For details, see Nutanix and VMware documentation.
- *For virtual machines in a vSphere environment:* Make sure that the latest version of VMware Tools is installed on the virtual machine. For details, see VMware documentation.
- *For Windows virtual machines:* Reactivate the Windows licenses.
- *For Linux virtual machines:* If a virtual machine on a Nutanix ESXi cluster or in a vSphere environment does not boot, change the disk controller from SCSI to IDE, and then install the latest version of VMware Tools on the virtual machine. You can later set the disk controller back to SCSI.
- *Only if you migrated virtual machines without network connectivity.* Make sure to configure the network settings on the virtual machine.
- Enable protection of the migrated data. For details, see "Protecting virtual machines" on page 154 and "Protecting applications" on page 252.

Performing disaster recovery of data to Azure

You can perform disaster recovery of data from the on-premises environment to Azure in the event of a disaster.

Prerequisite

The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected and must have the

DR-ready status. For more information, see “SpinUp specifics” on page 170.

Considerations

- Make sure the imported target is in the region to which you plan to migrate your virtual machines. This ensures the disaster recovery process is as fast and as cost-effective as possible.
- After you deploy the HYCU backup controller and use it to perform disaster recovery, you can keep the HYCU backup controller to stay prepared for disaster recovery in the future. For instructions on how to upgrade the HYCU backup controller when a new software release version is available, contact HYCU Support.

Procedure

1. Deploy a HYCU backup controller by using the HYCU R-Cloud or HYCU for Azure web user interface. For details on how to do this, see HYCU R-Cloud or HYCU for Azure documentation.
2. In Azure, create a new firewall rule to allow ingress network traffic on TCP port 8443 from the entire subnetwork to which the HYCU backup controller belongs. For details, see Azure documentation.
3. Sign in to the HYCU web user interface by specifying the following URL:

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

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

   > **Important** The credentials you provided in Azure during virtual machine creation cannot be used to sign in to HYCU and perform disaster recovery of data to Azure. For details on what credentials you can use to sign in to HYCU or to access the HYCU backup controller by using SSH, see “Signing in to HYCU” on page 68 or “Accessing the HYCU backup controller virtual machine by using SSH” on page 535.

4. Import the Azure target on which your backup data is stored to HYCU:
   a. In the Targets panel, click **Import**. The Import Target dialog box opens.
   b. From the Type drop-down menu, select **AZURE**.
   c. In the Storage account name field, enter the Azure storage account name as it was specified in the original target configuration.

d. In the Secret access key field, enter the secret access key for your Azure account.
e. In the Storage container name, enter the name of the storage container that is associated with the target and where the backup data is stored.
f. Click **Next**. The Import Backup Catalog dialog box opens.
g. Select the HYCU backup controller whose backup data you want to import, and then click **Next**.
h. In the Multiple Targets dialog box, do one of the following:
   - *If backup data is stored on one target:*

     Click **Import**.
   - *If backup data is stored on more than one target:*
     - Select each target one by one and specify the values so that they match the original target configuration.
     - For each target, click **Validate** to check the configuration.
     - Click **Import**.

5. Migrate your virtual machines or applications to cloud. For instructions, see "Migrating data to cloud" on page 610.

Protecting data across on-premises and Azure Government environments

You can use the SpinUp functionality to migrate protected data from your on-premises environment to Azure Government. In the event of a disaster in the on-premises environment, it provides disaster recovery of data to Azure Government.

Prerequisite

An Azure Government service principal must be added to HYCU. For instructions, see "Adding an Azure Government service principal" on page 469.

Depending on what you want to do, see one of the following:

I want to...	Instructions
Migrate protected data from the on-premises environment to Azure Government.	"Migrating virtual machines to cloud" below
Perform disaster recovery of data to Azure Government.	"Performing disaster recovery of data to Azure Government" on page 625

Migrating virtual machines to cloud

You can migrate virtual machines, servers, and applications running on virtual machines and servers to Azure Government by using the SpinUp functionality. Keep in mind that when you migrate an application, the whole virtual machine or the server on which this application is running is migrated to cloud.

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

Prerequisites

- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected.
- You must own a premium tier Platform license. For details, see "Licensing" on page 488.

Limitations

- If a restore point contains only a Snapshot tier, you cannot use it for migrating data.
- *For Nutanix clusters:* You cannot migrate volume groups.
- *For vSphere environments:*
  - You cannot migrate virtual machine templates.
  - Migrating data from snapshots is not supported.

Considerations

- If the restore point that you select contains a tier with an incomplete backup chain (due to one or more backups, copies of backup data, or data archives missing or being stored on a deactivated target), you cannot use this tier for migrating data.

- After you migrate data to cloud, an Azure temporary disk is automatically assigned to the migrated virtual machine. This disk is not a managed disk and it is used only for short-term data storage.
- *For virtual machines with secure boot enabled:* Because Azure does not currently support the secure boot feature for virtual machines, after you migrate such a virtual machine to cloud, secure boot cannot be enabled for it.

Depending on whether you want to migrate virtual machine or application data to cloud, access one of the following panels:

- Accessing the Virtual Machines panel

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

- Accessing the Applications panel

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

Procedure

1. In the Virtual Machines or Applications panel, select the entity that you want to migrate.
2. In the Detail view that appears at the bottom of the screen, select the virtual machine or application restore point that you want to use for the migration.

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

3. Click **SpinUp to Cloud**.
4. Select **SpinUp VM to Azure Government**, and then click **Next**.
5. From the Service Principal drop-down menu, select the service principal that has access to the required resources.
6. From the Subscription drop-down menu, select the appropriate subscription for the migrated virtual machine.
7. From the Resource Group drop-down menu, select the resource group for the migrated virtual machine.
8. From the Location drop-down menu, select the geographic region for the migrated virtual machine.

9. From the Availability Zone drop-down menu, select the zone for the migrated virtual machine.

   > Note The selected geographic region and the size of the virtual machine determine to which zones you can migrate data. If you do not want to migrate data to any zone, select **None**.

10. Click **Next**.
11. From the SpinUp From drop-down menu, select which tier you want to use for the migration. Your restore point can contain one or more tiers among which you can select:
    - **Automatic**: Ensures the fastest migration of data to cloud.
    - **Backup**
    - **Copy**
    - **Archive**
    - **Snapshot**
12. In the New VM Name field, enter a name for the migrated virtual machine.
13. In the vCPU Cores field, enter the number of virtual CPUs to be assigned to the migrated virtual machine multiplied by the number of cores per virtual CPU.
14. In the Memory field, enter the amount of memory (in GiB) to be assigned to the migrated virtual machine.
15. From the Virtual Machine Type drop-down menu, select the virtual machine type.

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

16. *Only if virtual disks have been excluded from the backup (manually or automatically):* Use the **Create excluded disks as blank** switch if you want blank disks of the same size and configuration as the excluded ones to be created and attached to the migrated virtual machine.
17. Under Network Interfaces, you can view the network interface that will be added to the migrated virtual machine. By default, this is the first network interface from the subscription that you selected for the migrated virtual machine. If required, you can also modify network settings.

Modifying network settings

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

> **Note** When adding a network interface, keep in mind that you can only add network interfaces that are attached to the same network. The maximum number of network interfaces that you can add depends on the selected virtual machine type.

Depending on how you want to modify network settings, do one of the following:

- Click **Add Network Interface** to add a network interface or click **Edit** next to the network interface that you want to edit, and then follow these steps:
  a. *Only if you are adding a network interface.* From the Virtual Network drop-down menu, select the network for the network interface.

     > **Note** The list of available networks includes only the ones within the region you selected for the migrated virtual machine.

  b. Select the subnet to which the network interface should be assigned.
  c. Under Public IP Address Type, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	No public IP address will be assigned to the network interface on the migrated virtual machine.
Static	A static IP address will be assigned to the network interface on the migrated virtual machine.
Existing	A preferred public IP address resource that you have created in Azure Government will be assigned to the network interface on the migrated virtual machine.

  d. Under Private IP Address Type, select the private IP address for the

network interface. You can select between the following options:

Option	Description
Dynamic	A dynamic IP address will be assigned to the network interface on the migrated virtual machine.
Static	The static IP address that you specify will be assigned to the network interface on the migrated virtual machine.

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

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

18. *Only if the virtual machine operating system has not been discovered yet.* Select the virtual machine operating system:
    - **Linux**
    - **Windows**
19. *Only if the platform readiness check was not performed for the selected restore point.* Use the **Adapt OS for migration** switch to apply the configuration changes that are required for the migration to cloud.

    If you enable this option, HYCU applies the following configuration changes to the virtual machines:

    - *Windows:*
      - Installs the required drivers.
      - Cleans up the existing devices to facilitate the discovery of new devices.
      - Enables the EMS console on serial port 1.
      - Resets the TCP/IP and Winsock stacks to use the DHCP.
    - *Linux:*
      - Includes the required kernel modules into `initramfs`.
      - Enables the serial console on serial port 1.
      - Changes the network configuration to use the DHCP.
20. Click **SpinUp**.

The Migration to cloud job starts.

After migrating data to cloud

- *For Windows virtual machines:* Reactivate the Windows licenses.
- *For Linux virtual machines:* Install the Linux Integration Services for Hyper-V and Azure on the virtual machine. For details, see Microsoft documentation.

Performing disaster recovery of data to Azure Government

You can perform disaster recovery of data from the on-premises environment to Azure Government in the event of a disaster.

Prerequisites

- The virtual machines that you want to migrate and the virtual machines with the applications that you want to migrate must be protected and must have the DR-ready status. For more information, see “SpinUp specifics” on page 170.
- You must have the HYCU virtual appliance image for Azure Government. To obtain the image and further instructions, contact HYCU Support.

Considerations

- When the HYCU backup controller is deployed in Azure Government, changing network settings is prevented in HYCU.
- Make sure the imported target is in the region to which you plan to migrate your virtual machines. This ensures the disaster recovery process is as fast and as cost-effective as possible.

Procedure

1. Deploy a HYCU backup controller:
   a. In Azure Government, create a managed image from the HYCU virtual appliance image.
   b. Create a virtual machine from the managed image. Make sure the virtual machine is configured with a public IP address and an additional disk of 128 GiB in size.

   For details, see Azure documentation.
2. In Azure Government, create a new firewall rule to allow ingress network traffic on TCP port 8443 from the entire subnetwork to which the HYCU backup controller belongs. For details, see Azure documentation.

3. Sign in to the HYCU web user interface by specifying the following URL:

   ```
   https://<IPAddress>:8443
   ```

   In this instance, *<IPAddress>* is the external IP address of the newly deployed HYCU backup controller.

   > ⓘ **Important** The credentials you provided in Azure Government during virtual machine creation cannot be used to sign in to HYCU and perform disaster recovery of data to Azure Government. For details on what credentials you can use to sign in to HYCU or to access the HYCU backup controller by using SSH, see "Signing in to HYCU" on page 68 or "Accessing the HYCU backup controller virtual machine by using SSH" on page 535.

4. Import the Azure Government target on which your backup data is stored to HYCU:
   a. In the Targets panel, click **Import**. The Import Target dialog box opens.
   b. From the Type drop-down menu, select **AZURE Government**.
   c. In the Storage account name field, enter the Azure Government storage account name as it was specified in the original target configuration.
   d. In the Secret access key field, enter the secret access key for your Azure Government account.
   e. In the Storage container name, enter the name of the storage container that is associated with the target and where the backup data is stored.
   f. Click **Next**. The Import Backup Catalog dialog box opens.
   g. Select the HYCU backup controller whose backup data you want to import, and then click **Next**.
   h. In the Multiple Targets dialog box, do one of the following:
      - *If backup data is stored on one target:*

        Click **Import**.
      - *If backup data is stored on more than one target:*
        - Select each target one by one and specify the values so that they match the original target configuration.
        - For each target, click **Validate** to check the configuration.
        - Click **Import**.

5. Migrate your virtual machines or applications to cloud. For instructions, see “Migrating virtual machines to cloud” on page 620.

Appendix A

Customizing HYCU configuration settings

You can find all HYCU configuration settings in the `config.properties.template` file in the `/opt/grizzly` folder on your HYCU backup controller. This file contains a list of all available configuration settings and their default values. If you want to adjust any of these configuration settings to meet your specific data protection environment needs and provide optimal performance, create a new `config.properties` file in the `/hycudata/opt/grizzly` folder, and then specify the preferred configuration settings and their new values.

> **Note** When you upgrade HYCU, the `config.properties` file will be kept. However, you may want to check the updated `config.properties.template` file for new configuration settings that you can use with the new HYCU version.

Depending on which configuration settings you want to customize, see one of the following sections:

- "Snapshot settings" on page 630
- "Utilization threshold settings" on page 630
- "SQL Server application settings" on page 631
- "Settings for aborting jobs" on page 631
- "File server settings" on page 631
- "Object server settings" on page 633
- "Data rehydration settings" on page 633
- "Disaster recovery settings" on page 634
- "User management settings" on page 634
- "Upgrade settings" on page 635

Procedure

1. Open a remote session to the HYCU backup controller virtual machine:

   ```
   ssh hycu@<HYCUBackupControllerIPAddress>
   ```

   When requested, enter the password for the `hycu` user.

   For detailed information about accessing the HYCU backup controller virtual machine by using SSH, see "Accessing the HYCU backup controller virtual machine by using SSH" on page 535.

2. Access and open the `config.properties` file by using one of the following text editors:

   - Vim:

     ```
     sudo vi /hycudata/opt/grizzly/config.properties
     ```

   - Nano:

     ```
     sudo nano /hycudata/opt/grizzly/config.properties
     ```

3. Edit any of the existing configuration settings as required.
4. Save and exit the `config.properties` file.

Changes to the configuration settings are applied based on their ReloadClass annotation in the `config.properties.template` file:

Annotation	Description
Job	The changes are applied when a new job is started.
Mount	The changes are applied in the following scenarios: • When any new target is added to HYCU. • When an existing NFS, SMB, Nutanix, iSCSI, or tape target is activated again after being deactivated with the Detach storage option enabled.
Operation	The changes are applied when a new operation that does not create a job is executed (for example, when using the HYCU web user interface, REST API, SSH, or WinRM).
Service	The changes are applied when the HYCU application server (the Grizzly server) is restarted.

If a configuration setting has no annotation, it is recommended to restart the HYCU application server (the Grizzly server). To do so, run the following command:

```
sudo service grizzly restart
```

Snapshot settings

You can use the following settings to configure the snapshot retention threshold at which an event is triggered:

Settings/descriptions
`max.snapshots.per.vm` If the number of snapshots that are retained per virtual machine exceeds the specified value, a warning event is triggered. The default value is 24.
`max.snapshots.per.cluster` If the number of snapshots that are retained per Nutanix cluster exceeds the specified value, a warning event is triggered. The default value is 2400.

Utilization threshold settings

You can use the following settings to configure the system and data disks as well as target utilization thresholds:

Settings/descriptions
`controller.disk.full.warning.threshold.fraction` If the HYCU backup controller utilization of the system or data disk exceeds the specified value, an event is triggered. The default value is 0.90.
`target.utilization.threshold.red.fraction` If the HYCU backup controller utilization of the target exceeds the specified value, its health status indicator becomes red. The default value is 0.95.
`target.utilization.threshold.yellow.fraction` If the HYCU backup controller utilization of the target exceeds the specified value, its health status indicator becomes yellow. The default value is 0.90.

For detailed information about the health status of the target, see "Viewing target information" on page 417.

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 `true`). If you want to disable it, make sure to set the value for this setting to `false`.

Settings for aborting jobs

You can use the following settings to configure when a job that has the Executing status will be aborted automatically:

Settings/descriptions
`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.

File server settings

You can use the following settings to configure file share backups:

Settings/descriptions
`afs.reindex.interval.count` Number of incremental file share backups after which a full reindex is performed, which increases the responsiveness of the file restore process. The

Settings/descriptions
default value is 5.

`afs.partial.success.threshold.count`
and
`afs.partial.success.max.fail.fraction`
 	
Maximum number and fraction of failed file backups up to which the backup status of the corresponding file share is marked as Completed with errors (and not as Failed). The default values are the following:
• *For the maximum number of failed file backups* (`afs.partial.success.threshold.count`): 10000
• *For the maximum fraction of failed file backups* (`afs.partial.success.max.fail.fraction`): 0.01
 	
Important Both of these values must be exceeded for the backup status of the file share to be marked as Failed.

`afs.instance.afs.cluster.priority` HYCU uses an internal algorithm to distribute the load among multiple HYCU instances. It prioritizes the HYCU instances that are running on the same Nutanix cluster as the file server and the HYCU instances that are running on the same Nutanix cluster as the HYCU backup controller. It also takes into account the number of jobs that are already running on each HYCU instance. Raising the value of this setting gives higher priority to the HYCU instances that are running on the same Nutanix cluster as the file server.
`afs.instance.bc.cluster.priority` HYCU uses an internal algorithm to distribute the load among multiple HYCU instances. It prioritizes the HYCU instances that are running on the same Nutanix cluster as the file server and the HYCU instances that are running on the same Nutanix cluster as the HYCU backup controller. It also takes into account the number of jobs that are already running on each HYCU instance. Raising the value of this setting gives higher priority to the HYCU instances that are running on the same Nutanix cluster as the HYCU backup controller.

Object server settings

Settings/descriptions
`afs.reindex.interval.count` Number of incremental bucket backups after which a full reindex is performed, which increases the responsiveness of the object restore process. The default value is 5.

`afs.partial.success.threshold.count`
and
`afs.partial.success.max.fail.fraction`
 	
Maximum number and fraction of failed object backups up to which the backup status of the corresponding bucket is marked as Completed with errors (and not as Failed). The default values are the following:
• *For the maximum number of failed object backups* (`afs.partial.success.threshold.count`): 10000
• *For the maximum fraction of failed object backups* (`afs.partial.success.max.fail.fraction`): 0.01
 	
ⓘ Important Both of these values must be exceeded for the backup status of the bucket to be marked as Failed.

Data rehydration settings

You can use the following settings to configure HYCU to perform data rehydration:

Settings/descriptions
`target.azure.blob.rehydration.enable` HYCU is preconfigured to perform data rehydration before performing the restore if backup data or a copy of backup data is stored in the Azure archive access tier. During a rehydration task, the data is moved from the archive access tier to the hot access tier from which HYCU can restore data. HYCU does not move data back to the archive access tier afterward. The default value is `true`.
`target.azure.blob.rehydration.threads` Number of blobs that can be rehydrated in parallel. The default value is 20.

Disaster recovery settings

You can use the following settings to enable additional scenarios for disaster recovery or adjust automatic target synchronization:

Settings/descriptions
`clone.enabled.for.hycu.dr` HYCU is preconfigured to prevent creating clones of the HYCU backup controller (the virtual machine itself or its virtual disks). > ⚠ Caution Do not activate a clone of the HYCU backup controller while the original HYCU backup controller is still active. If such activation happens, data loss may occur. All currently running backups fail and their status is set to Error. The corresponding restore points are then automatically removed by the HYCU cleaning process. If set to `true`, cloning of the HYCU backup controller is enabled and the respective restore options become available in the HYCU web user interface.
`synchronize.target.catalog.interval.minutes` When the recovery HYCU backup controller is in recovery mode, automatic target synchronization is by default performed every 60 minutes. Setting the value to 0 disables automatic target synchronization.

User management settings

You can use the following setting to completely prevent deleting protected data when changing ownership of virtual machines, file shares, and buckets:

Setting/description
`force.keep.backups.on.owner.change` If set to `true` (the default value is `false`), data protected by specific owners is never deleted—even if the option to delete such data is specified when changing ownership of virtual machines, file shares, and buckets in any of the HYCU interfaces.

Upgrade settings

You can use the following setting to define which repository should be used to perform the HYCU upgrade:

Setting/description
`upgrade.ab.hycu.repo` By default, HYCU is upgraded by using the images that are stored in the official HYCU public repository. If you want HYCU to use your private local repository, set the value to `http(s)://<RepositoryUrl>/`, where `<RepositoryUrl>` is the URL of the root folder of your private local repository.

Appendix B

After restoring a virtual machine to a different source

A virtual machine can be restored to a different source by using the Clone VM restore option as described in "Cloning a virtual machine" on page 205. However, depending on your virtual machine original environment and target environment, you might have to perform some additional steps after the restore:

VM original environment	VM target environment	Additional steps
Nutanix ESXi, vSphere, XenServer, Azure Local, Hyper-V, AWS GovCloud (US), Azure, Azure Government, or servers	Nutanix AHV	See "After restoring a virtual machine to a Nutanix AHV cluster" on the next page.
vSphere	Nutanix ESXi	See "After restoring a virtual machine to a Nutanix ESXi cluster" on page 638.
Nutanix AHV, Nutanix ESXi, XenServer, Azure Local, Hyper-V, AWS GovCloud (US), Azure, Azure Government, or servers	vSphere	• *Only if restoring a virtual machine with more than one disk.* After the restore, additional disks will be offline. Make sure to bring them back online. • *Only if restoring a Windows virtual machine from AWS GovCloud (US).* Make sure that the latest version of VMware Tools is installed on the virtual machine. For

VM original environment	VM target environment	Additional steps
		instructions, see VMware documentation. • *Only if the restored virtual machine has more than one disk.* Check the hard drive boot order of the restored virtual machine. If it differs from the one on the original virtual machine, change the boot order in BIOS.

After restoring a virtual machine to a Nutanix AHV cluster

Considerations

- *Only if restoring a virtual machine with more than one disk from a vSphere environment to a Nutanix AHV cluster.* After the restore, additional disks will be offline. Make sure to bring them back online.
- If you have not followed the recommendations described in "Preparing for the restore to a different source" on page 163, your virtual machine will not boot after the restore, and you must perform the following additional steps:
  1. Make sure that the restored virtual machine is turned off.
  2. As the administrator or the root user, sign in to the Nutanix AHV cluster by using SSH.
  3. List the virtual machine details:

     ```
     acli vm.get <VMName>
     ```

  4. Take a note of the current bus and index values in the `disk_list` section.
  5. Clone the existing disk to a new disk on the compatible bus:

     ```
     acli vm.disk_create <VMName> bus=<BusType>
     clone_from_vmdisk=vm:<VMName>:<CurrentBus>.<CurrentIndex>
     ```

In this instance, <*VMName*> is the name of the restored virtual machine, <*BusType*> is `scsi`, `ide`, or `sata`, <*CurrentBus*> is the bus value from the `disk_list` section, and <*CurrentIndex*> is the index value from the `disk_list` section.

If the original virtual machine has the SATA or SCSI disks, clone them to the SATA disks. For example:

```
acli vm.disk_create test-vm bus=sata
clone_from_vmdisk=vm:test-vm:scsi.0
```

If the original virtual machine has the IDE disks, clone them to the IDE disks. For example:

```
acli vm.disk_create test-vm bus=ide
clone_from_vmdisk=vm:test-vm:ide.0
```

After you perform the previous procedure for all the disks, follow these steps:

1. Sign in to the Nutanix Prism web console.
2. In the menu bar, click **Home**, and then select **VM**.
3. Click the **Table** tab to display the VM Table view.
4. From the list of virtual machines, select the restored virtual machine, and click **Update**.
5. Delete the source disks, and then select the boot disk and click **Save**.
6. Click **Power on** to turn on the restored virtual machine.
7. Install the Nutanix Guest Tools software bundle of the latest version on the virtual machine.
8. *Recommended for virtual machines that had the SCSI disks.* Clone the controller back to the SCSI controller.

For details on how to update a virtual machine on a Nutanix cluster, see Nutanix documentation.

After restoring a virtual machine to a Nutanix ESXi cluster

If after restoring a virtual machine from a vSphere environment to a Nutanix ESXi cluster the virtual machine does not start, you must perform additional

steps.

Steps

- If the type of controller on the restored virtual machine is not the same as it was on the original virtual machine, do the following:
  1. Sign in to the vSphere Client.
  2. Click the **VMs** tab, and then right-click the restored virtual machine and select **Edit Settings**.
  3. On the Virtual Hardware tab, modify the controller settings so that they match the ones on the original virtual machine.
- If the virtual machine uses UEFI firmware, you may need to select the boot file manually. In this case, do the following:
  1. Sign in to the vSphere Client.
  2. Access the EFI Boot Manager menu, and then do the following:
     a. Select the **Enter setup** option.
     b. Enter the boot maintenance manager by selecting **Boot option maintenance menu**.
     c. Use the **Boot from a File** option to browse for a boot file.
     d. Find a device whose name contains the GPT string that represents the boot partition, and then press **Enter** to open it.
     e. Navigate to the EFI boot file that you can find at the following location:
        - Windows: `\EFI\Microsoft\Boot\bootmgrfw.efi`
        - Linux: `/EFI/<OSName>/grubx64.efi`
     f. Press **Enter** to resume booting.

Provide feedback

For any suggestions and comments regarding this product or its documentation, send us an e-mail to:

info@hycu.com

We will be glad to hear from you!

www.hycu.com

HYCU®