

syneto

Simplify • Accelerate • Protect™

SynetoOS User Guide

Version 4.4

August 2019

COPYRIGHT NOTICE

© 2018 SYNETO SRL. All rights reserved.

™ *Syneto - Your Data • Our Business* and the Syneto logo are trademarks of SYNETO SRL.

All other trademarks in this document are the property of their respective owners.

No parts of this publication may be used for any purposes other than documentation for a Syneto product.

Syneto may have trademarks, copyright, patents, or other intellectual property rights covering the subject matter of this publication. This document, and the possession of this document, does not give you any license to these trademarks, copyrights, patents, or other intellectual properties.

The information contained in this document describes the software configuration of a Syneto HYPERSeries appliance for the specific version mentioned on the cover. This document can change or become obsolete by newer product versions without notice. Syneto shall not be liable for technical or editorial errors or omissions contained herein.



TABLE OF CONTENTS

[DEFAULT USERNAMES AND PASSWORDS](#)

- [ESXi management interface](#)
- [SynetoOS management interface](#)
- [IPMI management interface](#)

[ABOUT THIS PRODUCT](#)

- [What is it?](#)
- [Version](#)
- [Document updates](#)
- [Product licensing](#)
- [VMware vSphere integration](#)
- [Used Ports](#)

[GATHERING REQUIRED INFORMATION](#)

- [vSphere license](#)
- [IP addresses](#)
- [Finding IP addresses](#)
- [vCenter and ESXi hosts credentials](#)
- [Active Directory credentials](#)

[ACTIVATE THE WEB MANAGEMENT INTERFACE](#)

- [Access ESXi web console](#)
- [Configure basic Management VM networking](#)
- [First time quick setup](#)
- [Activate the appliance](#)
- [Access the management interface](#)
- [Configure using IPMI](#)
 - [Default credentials](#)
- [Management interface main components](#)

[INITIAL CONFIGURATION](#)

- [Configure ESXi networking](#)
- [Install ESXi license](#)
- [Configure Management VM networking](#)
- [Update Syneto HYPER to latest version](#)
- [Join vCenter](#)
- [Configure date / time](#)
- [Email server and Alert thresholds](#)

[Recommended thresholds](#)

[Define recurring integrity check intervals](#)

[THE DASHBOARD](#)

[DATASTORES](#)

[Introduction](#)

[Provisioning and managing](#)

[Provision new datastore](#)

[Provisioning new datastore when multiple pools are available](#)

[Deleting a datastore](#)

[Protecting](#)

[Monitoring](#)

[Virtual machines](#)

[Used space](#)

[Recovering data from a snapshot](#)

[Get access to a previous version of the datastore](#)

[Clean up when you are done](#)

[Reverting a clone](#)

[Malware / Ransomware use case](#)

[Registering additional virtual machines from a clone](#)

[Serving non-Syneto hosts](#)

[Defining external hosts](#)

[Mounting datastore to additional hosts](#)

[Unmounting datastore from hosts](#)

[VIRTUAL MACHINES](#)

[Introduction](#)

[Provisioning and managing](#)

[Power Management](#)

[Remote Console](#)

[Protecting](#)

[Consistency levels](#)

[Monitoring](#)

[CPU](#)

[Memory](#)

[Used space](#)

[Additional information](#)

[Analyzing performance](#)

[Dataset IOPS](#)

[Dataset bandwidth](#)

[Dataset average block size](#)

[Recovering data from a snapshot](#)

[Get access to a previous version of the virtual machine](#)

[Take a manual snapshot](#)

[Clean up when you are done](#)

[Migrating Virtual Machines](#)

[How to migrate virtual machines](#)

[SHARES](#)

[User management](#)

[Active Directory integration](#)

[Managing local users and groups](#)

[Provisioning and managing](#)

[SMB - Create share](#)

[SMB - Manage share](#)

[AFP - Create share](#)

[AFP - Manage share](#)

[Protecting](#)

[Monitoring](#)

[Type](#)

[Used space](#)

[Quota](#)

[Analyzing performance](#)

[Dataset IOPS](#)

[Dataset bandwidth](#)

[Dataset average block size](#)

[Recovering data from a snapshot](#)

[Access a previous version of your share](#)

[Windows previous versions](#)

[Clean up when you are done](#)

[Reverting a clone](#)

[VOLUMES](#)

[Introduction](#)

[Access management](#)

[iSCSI remote initiators](#)

[iSCSI local targets](#)

[Provisioning and managing](#)

[Protecting](#)

[Monitoring](#)

[Used space](#)

[Provisioning](#)

[Analyzing performance](#)

[Dataset IOPS](#)

[Dataset bandwidth](#)

[Dataset average block size](#)

[Recovering data from a snapshot](#)

[Access a previous version of your volume](#)

[WARNING: Do not change data on the clones. Clones cannot be protected, and after the Cleanup process \(see below\), all changes will be deleted.](#)

[Clean up when you are done](#)

[DRIVES AND POOLS](#)

[Introduction](#)

[Replacing a faulted drive](#)

[Replacing a faulted OS drive](#)

[Replacing a faulted cache drive](#)

[Creating a pool](#)

[Expanding a pool](#)

[Importing a disk pool](#)

[Exporting a disk pool](#)

[Protecting against silent data corruption](#)

[EXPANSION SHELVES](#)

[Attaching expansion shelves](#)

[Expanding a pool with JBOD drives](#)

[DATA PROTECTION AND REPLICATION](#)

[Introduction](#)

[Configuring replication targets](#)

[Adding a replication target](#)

[Removing a replication target](#)

[Editing a replication target](#)

[Configuring virtual machine credentials](#)

[Protecting a datastore](#)

[Configuring the snapshot schedule for a datastore](#)

[Checking the progress of a snapshot or replication of a datastore](#)

[Canceling a running replication of a datastore](#)

[Monitoring the snapshots of a datastore](#)

[Configuring the type of snapshots for the virtual machines](#)

[Exceptions: Excluding time intervals from Schedules](#)

[Protecting a share](#)

[Configuring the snapshot schedule for a share](#)

[Checking the progress of a snapshot or replication of a share](#)

[Canceling a running replication of a share](#)

[Monitoring the snapshots of a share](#)

[Exceptions: Excluding time intervals from Schedules](#)

[Protecting a volume](#)

[Configuring the snapshot schedule for a volume](#)

[Checking the progress of a snapshot or replication of a volume](#)

[Canceling a running replication of a volume](#)

[Monitoring the snapshots of a volume](#)

[Exceptions: Excluding time intervals from Schedules](#)

[Recovering files from a virtual machine](#)

[Browse file recover](#)

[Download files](#)

[Attach files to a VM](#)

[SYSTEM CONFIGURATION](#)

[Networking](#)

[Ethernet interfaces](#)

[Missing or broken interfaces](#)

[Date/time and timezones](#)

[Password change](#)

[Changing External ESXi Host Password](#)

[Alerts and notifications](#)

[Alerts Page](#)

[New Software Update](#)

[Notification boxes](#)

[Notification emails](#)

[ANALYTICS](#)

[Introduction](#)

[Low level analytics](#)

[Mid level analytics](#)

[Case 1 - I read a lot of data, but I don't see any read activity in analytics](#)

[Case 2 - I write continuously to my Syneto HYPER, but pool graphics show write IOPS in spikes instead of a continuous line](#)

[Case 3 - I want to make sure my pool can read data fast enough from the drives](#)

[Case 4 - I have several disk pools and one is very slow](#)

[High level analytics](#)

[Miscellaneous analytics](#)

[SOFTWARE UPDATES](#)

[Update without restart](#)

[Update with restart](#)

[TROUBLESHOOTING](#)

[A snapshot replication fails](#)

[Cannot enable customer support](#)

[Can't connect to Active Directory](#)

[Can't create app-consistent/live snapshots because existing VMware snapshot](#)

[After reinstall, data pool cannot be imported](#)

[Error when creating a datastore, share, or volume](#)

[Datastores are not accessible on vSphere management interface after Syneto HYPER IP change](#)

[Datastores are not accessible on vSphere management interface after ESXi host IP change](#)

[I don't know if Syneto Central is accessible](#)

[Dialog appears asking for the new local ESXi password](#)

[APPENDIX A: CREATING A CENTRAL ACCOUNT](#)

[APPENDIX B: ACCESSING CUSTOMER SUPPORT](#)

[Contacting Syneto Customer Support](#)

[Support for VMware ESXi hypervisor issues](#)

[APPENDIX C: SERENITY SERVICE](#)

[APPENDIX D: GRANULAR FILE RECOVERY](#)

DEFAULT USERNAME AND PASSWORDS

Thank you for using a Syneto product. For a quick reference, are the default usernames and passwords for the various management interfaces.

ESXi management interface

Username: root

Password: Syneto.Default1

SynetoOS management interface

Username: admin

Password: admin

IPMI management interface

Username: ADMIN

Password: ADMIN

ABOUT THIS PRODUCT

What is it?

Syneto HYPER is a modern hyperconverged product specifically designed for the needs of the small and medium business. It features built-in virtualization (using industry leading hypervisor ESXi), NAS file sharing (SMB for Windows networks and AFP for Mac networks) and an embedded backup and disaster recovery service that protects both virtual machines and datastores. The operating system is called **SynetoOS** and it runs as a controller VM on top of the hypervisor.

Version

This document is for SynetoOS version **4.4.2**.

Document updates

As we release new software versions and learn more about the way our customers use HYPER and SynetoOS, we update and expand the documentation with new chapters and new information.

To download the latest version of the product documentation, we highly recommend that you periodically check the Central¹ account associated with your HYPER machine.

Product licensing

HYPER does not require the explicit installation of Syneto license files. The product is automatically activated based on your support subscription by our Syneto Central server at the first time you log in. See [Activate using a Central account](#) for more details.

NOTE: You will need to procure and install a VMware vSphere license. **vSphere Essentials** is the minimum required license. Check with your Syneto representative, as Syneto can supply the required VMware vSphere licenses on request.

VMware vSphere integration

Syneto products integrate both with standalone ESXi servers and within VMware vCenter deployments.

Each Syneto HYPER comes with an ESXi local hypervisor and it can serve data to any other VMware ESXi from your existing network (standalone or in or vCenter clusters).

From now on, we will refer to the ESXi or vCenter as "vSphere". For example "**vSphere management interface**" will refer to either ESXi or vCenter configuration interface. When a distinction is necessary, the document will refer to the particular product.

¹ The official Syneto support portal: <https://central.syneto.eu>.

Used Ports

In order to allow remote management, VMWare integration, and snapshot replication, SynetoOS uses a set of TCP ports. The list of services exposed over the network, and their role can be found in the table below.

SynetoOS

TCP Port	Service	Description
443	HTTPS / Web Server	Web Management Interface.
80	HTTP / Web Server	Allows accessing the management with an http:// prefix. Redirects immediately to https:// for security reasons.
22	SSH	Secure Shell service allows you to connect to the CLI management interface. It is also used for secure snapshot replication.
111, 968, 2049, 4045	NFS	Used for sharing datastores to ESXi hosts.
137,138	SMB	Used for SMB sharing.
548	AFP	Used for AFP sharing.
2003, 2004, 7002	carbon_cache	Used for Analytics.
5353 (UDP)	mdnsd	DNS resolver. Used to resolve host names.
9000-9099	mbuffer	One port opened temporarily for each non-encrypted snapshot replication. Port is closed after a snapshot is received. These ports cannot be NATed, or you have to NAT all of them one-to-one on you gateway.
11211	memcached	Used by the Web Management Interface to cache various data in the memory.

ESXi Ports

See official VMware documentation:

<https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.security.doc/GUID-171B99EA-15B3-4CC5-8B9A-577D8336FAA0.html>

IPMI Ports

TCP Ports: 80, 443, 5901, 5900, 5120, 5123 and UDP Ports: 623

GATHERING REQUIRED INFORMATION

Before starting to use Syneto HYPER, there are several pieces of information you need to acquire:

- A valid serial number for VMware vSphere
- Networking settings for the hypervisor and controller VM
- Credentials for accessing existing ESXi hosts or vCenter servers
 - If you want to host on HYPER datastores for existing hosts
- Active Directory credentials
 - For creating SMB shares in a Windows network

vSphere license

HYPER machines come preinstalled with a temporary ESXi license (valid 60 days) and you need to install a final VMware license for continued use of the product. While the temporary license is active, all product features are enabled. During this time frame, you should contact the person administering your VMware licenses, and ask for a **vSphere essentials** license key.

NOTE: The format of the license is: XXXXX-XXXXX-XXXXX-XXXXX-XXXXX.

More information about how to use the vSphere license can be found in the chapter [Install ESXi license](#).

IP addresses

To function properly, HYPER needs three IPs in the management network. These can be self-assigned via DHCP, or they can be statically configured. If you don't have a DHCP server on your network, ask your network administrator for the following IP addresses:

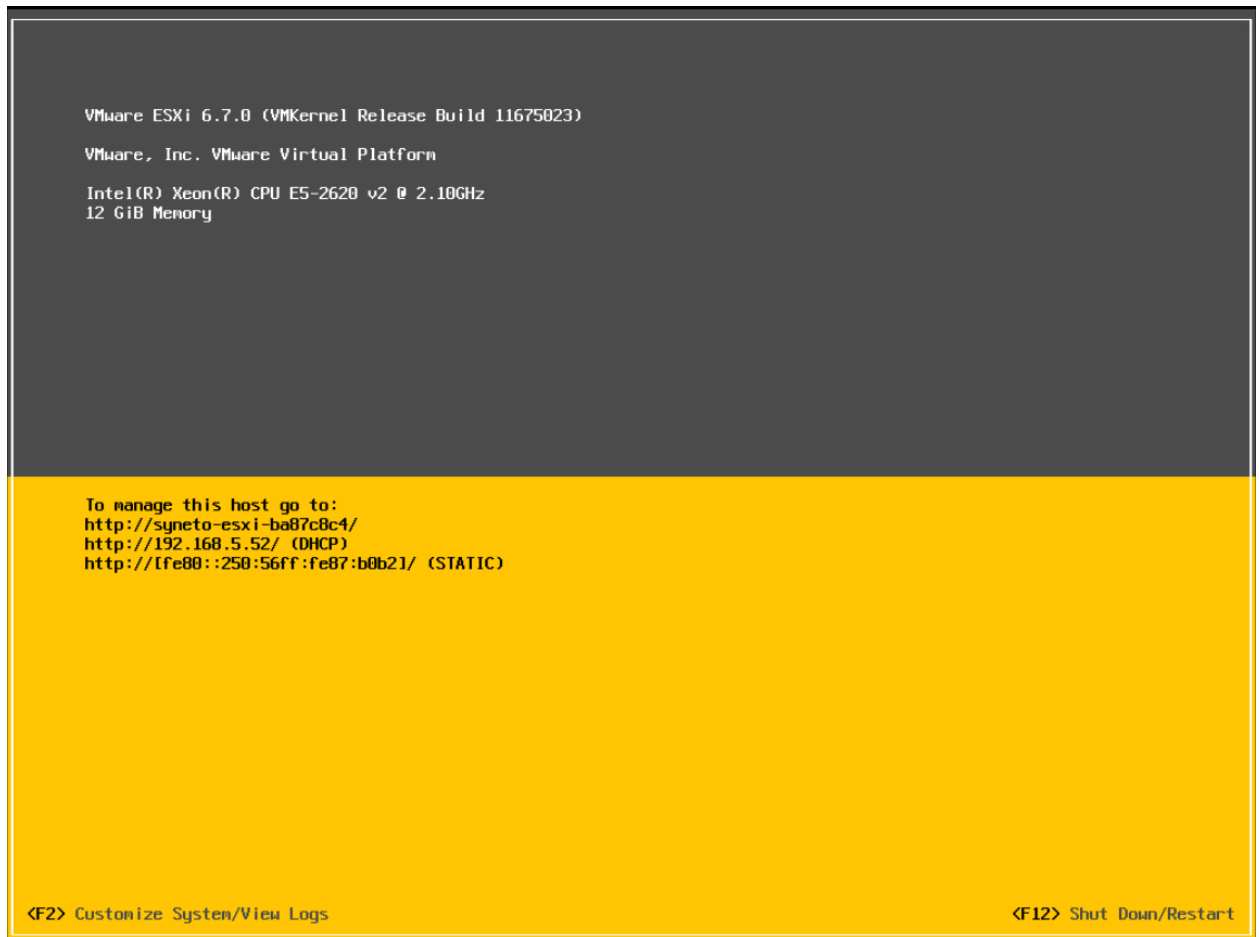
- Allocate three IPs in the **management network**, for:
 - IPMI (machine management)
 - ESXi (hypervisor management)
 - SynetoOS (controller VM)
- **DNS server** IP address - you can configure between one and three DNS servers in the management network.
- **Gateway** IP address - you will need this to access the UI and the Central Portal

NOTE: HYPER uses an internal virtual network (172.16.254.1/24) for communication between the controller VM and the hypervisor. Changing this network is not possible.

Finding IP addresses

In networks managed using DHCP, the administrator might want to find the management IP addresses allocated by the DHCP server. As stated above, there are 3 IP Addresses that are required.

The ESXi management IP addresses are shown on ESXi's DCUI splash screen, accessible using **Alt+F2** on the physical video console:



The SynetoOS (and IPMI management - if configured) IP addresses are shown on the ESXi service console shell, accessible using **Alt+F1** on the physical video console:

```
VMware ESXi 6.7.0 build-11675023 http://www.vmware.com
Copyright (c) 2007-2017 VMware, Inc.

SynetoOS http://www.syneto.eu
Copyright (c) 2007-2018 Syneto SrL.

SynetoOS management: https://192.168.5.152/
```

vCenter and ESXi hosts credentials

HYPER can provision datastores for an unlimited number of existing ESXi hosts: both standalone and managed by vCenter. To provide storage for ESXi/vCenter, you will need for the following informations:

- IP address / hostname
- Username
- Password

When adding a vCenter, you will be able to mount the datastores on all the ESXi hosts on that vCenter.

For more information about connecting to vCenter, can be found in the [Join vCenter](#) chapter.

Active Directory credentials

You need Active Directory credentials for provisioning SMB shares in a network managed by Active Directory.

You will need to get from your system administrator the following information:

- Domain name

NOTE: The domain name is case sensitive.

NOTE: The format is: `domain.mycompany`.

- Domain controller (a server that responds to security authentication requests within the Windows Server domain)

NOTE: The format is: `dc.dn.mycompany.com`. Alternatively, an IP address can be provided.

- Username - for a user with sufficient privileges to add computers to the domain
- Password

More information about when and how to configure Active Directory integration can be found in the [Shares](#) chapter.

ACTIVATE THE WEB MANAGEMENT INTERFACE

Administration of the appliance is performed via one of the web management interfaces. There are three interfaces:

- **IPMI:** through this interface you can connect to the ESXi console, the same way as if you connected a monitor and keyboard to the machine. You can change the ESXi password or network configuration. In case of a reinstall, you can attach the SynetoOS ISO and start the installation from your desk.
- **ESXi:** from this interface, you can create new virtual machines and upload ISOs. You can power off, power on or suspend a virtual machine. You can also mount iSCSI volumes.
- **Management VM:** the SynetoOS interface. You can monitor and provision datastores for virtual machines, AFP and SMB shares and iSCSI volumes. You can also backup these entities and replicate them to be able to recover in case of a disaster.

This chapter describes the activation procedure for each of the interfaces. The detailed sequence of steps is the following:

1. Get physical access to HYPER using monitor and keyboard
2. Activate the ESXi web console
3. Configure basic networking of the management VM
4. Activate machine with online Syneto Central service
5. Get access to the HYPER Management VM web console.
6. Activate IPMI remote system management interface

Access ESXi web console

After your Syneto HYPER is connected to the physical infrastructure, power it on and attach a monitor and a keyboard. The monitor will display the ESXi console.

Press **F2** and enter the default ESXi username and password (username: `root`; password: `Syneto.Default1`). A menu will open, where you can modify different options, set a new password, take troubleshooting actions and more.

At this point you can configure static ESXi networking (recommended).

```
VMware ESXi 6.7.0 (VMKernel Release Build 11675023)
VMware, Inc. VMware Virtual Platform
Intel(R) Xeon(R) CPU E5-2620 v2 @ 2.10GHz
12 GiB Memory

To manage this host go to:
http://syneto-esxi-ba87c8c4/
http://192.168.5.52/ (DHCP)
http://[fe80::250:56ff:fe87:b0b21]/ (STATIC)

<F2> Customize System/View Logs                                <F12> Shut Down/Restart
```


On the ESXi console, you can see the IP address to the ESXi UI. Insert the address in a browser (of your choice) and log in to the UI with the default credentials (default username: `root`; password: `Syneto.Default1`).

You now have access to the ESXi web console.

Note: For more information about ESXi configuration, please review the [Configure ESXi networking](#) chapter.

Configure basic Management VM networking

After accessing the ESXi web console, go to *Virtual Machines* → *SynetoOS*. Open its console and log in with the default credentials (username: `admin`; password: `admin`). If DHCP has been enabled on your network, you will see that there are already two IP addresses assigned to the SynetoOS virtual machine. If HYPER was connected to a network without a DHCP server, only the internal network will be configured (`172.16.254.2`).



If you do not have DHCP, or want to override the default values for IP, DNS and gateway, run the following commands:

```

net ip edit vmxnet3s0 address <IPv4_address>/<netmask>
net dns add <IPv4_address>
net gateway set <IPv4_address>
  
```

To review the current settings for IP, gateway and DNS, run the following commands:

```

net ip show
net gateway show
net dns show
  
```

Take a note of the IP for `vmxnet3s0` as we'll be using it later on.

Finally, let's configure the domain name:

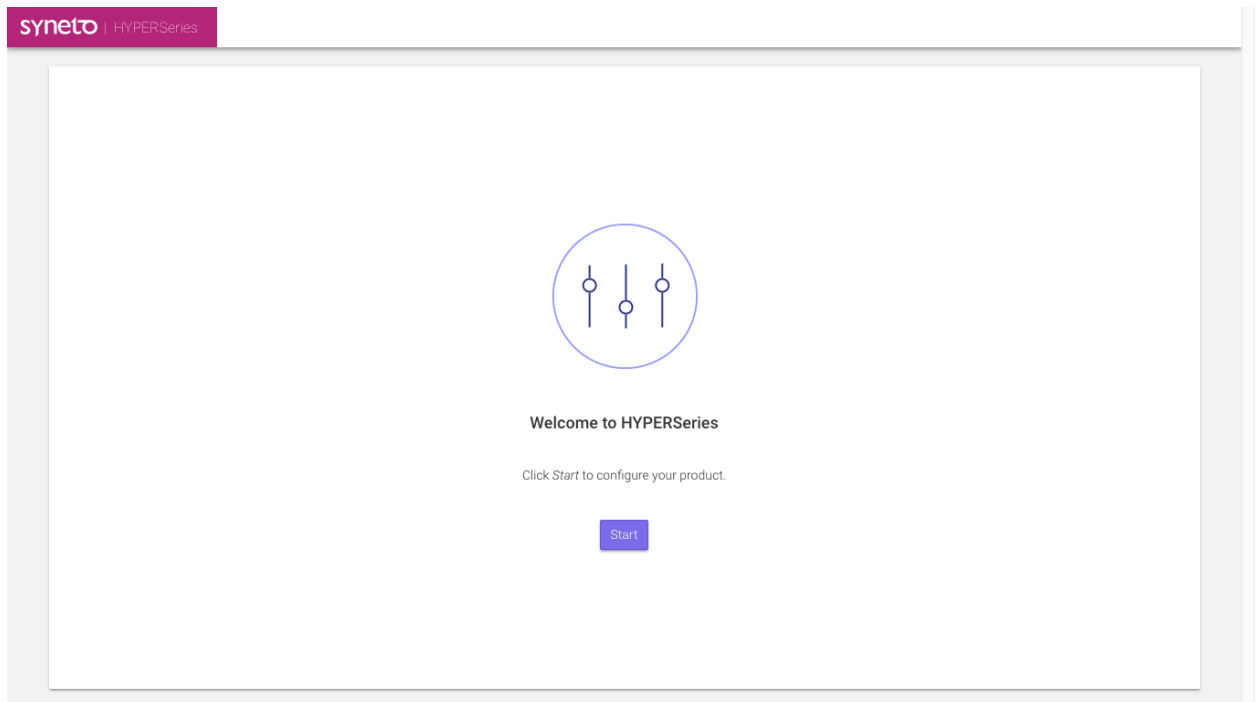
```
net dns setdomain <domain>
```

At this point, we can access the management VM web interface at `https://<IP_address>`. The IP is the one configured for `vmxnet3s0`. At the login screen, enter the default credentials (user:admin; password: admin).

First time quick setup

Steps: Start

Click **Start** to configure your product.



Step 1: Network

Configure networking settings for the Management VM and the Hypervisor.

The screenshot shows the SynetoOS configuration interface for network settings. The interface is titled "Network" and includes a sidebar with a list of configuration steps: 1. Network (selected), 2. Date and time, 3. Maintenance service, 4. Passwords, 5. Active Directory, 6. Notifications, and 7. Updates. The main content area is titled "Management VM (Syneto OS)" and contains the following fields:

- Hostname:** syneto-os-2ff8aaef
- Domain name:** dev.syneto.net
- DNS 1:** 192.168.1.15
- DNS 2:** #####
- DNS 3:** #####
- Default gateway:** 192.168.1.1
- Use a proxy?

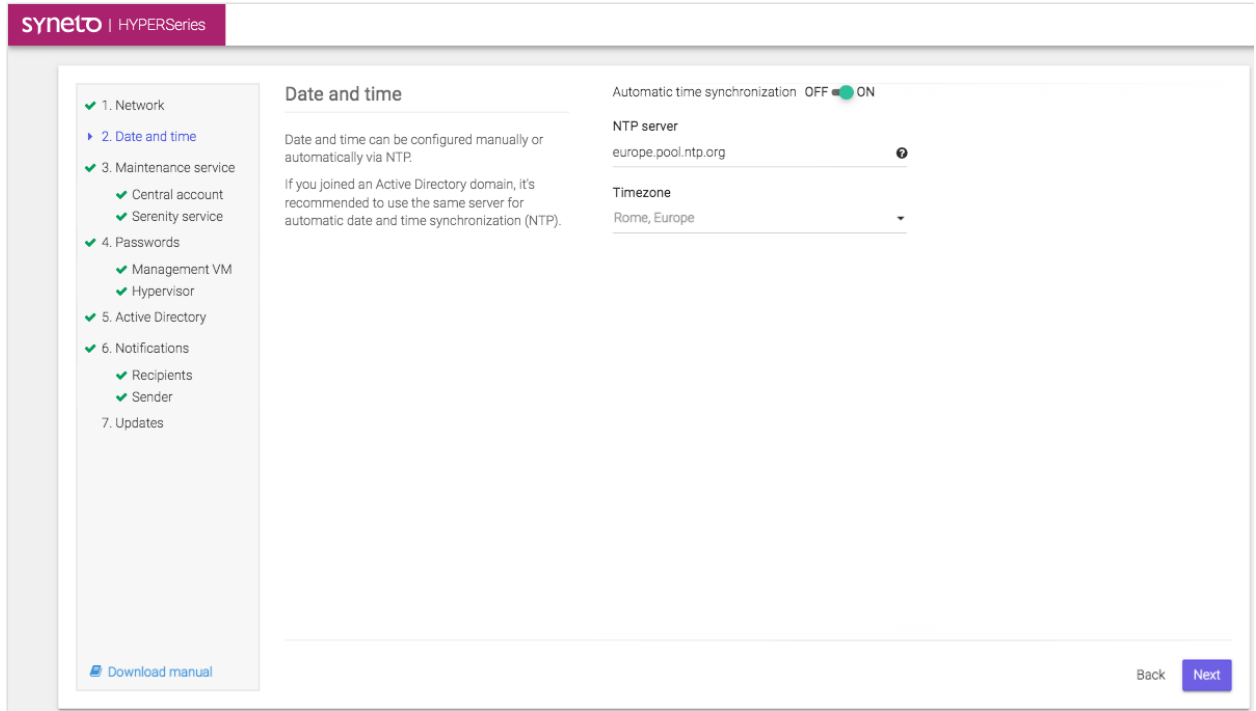
A yellow warning box is present in the center of the main content area, stating: "WARNING: A working Internet connection is mandatory for successful product activation. Check with your network administrator for proxy or firewall restrictions." At the bottom right of the interface, there are "Back" and "Next" buttons.

NOTE: A working Internet connection is mandatory for successful product activation. Check with your network administrator for proxy or firewall restrictions.

Step 2: Date and time

Date and time can be configured manually or automatically via NTP.

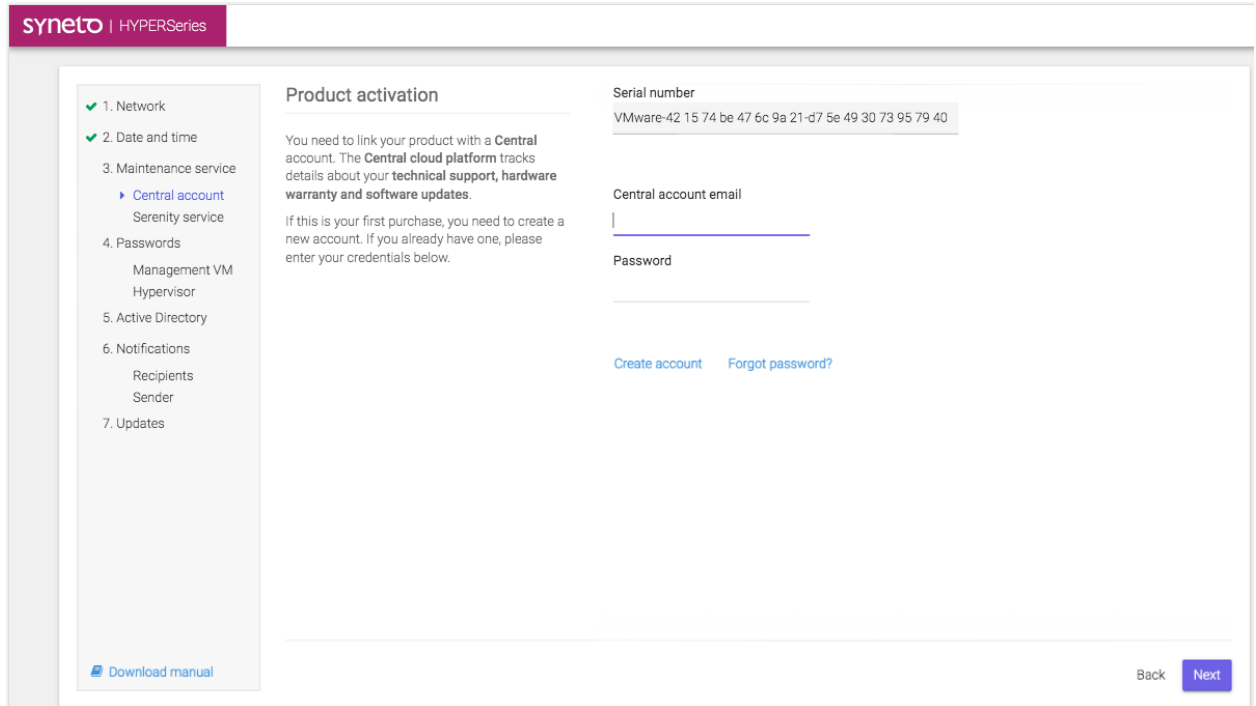
If you joined an Active Directory domain, it's recommended to use the same server for automatic date and time synchronization (NTP).



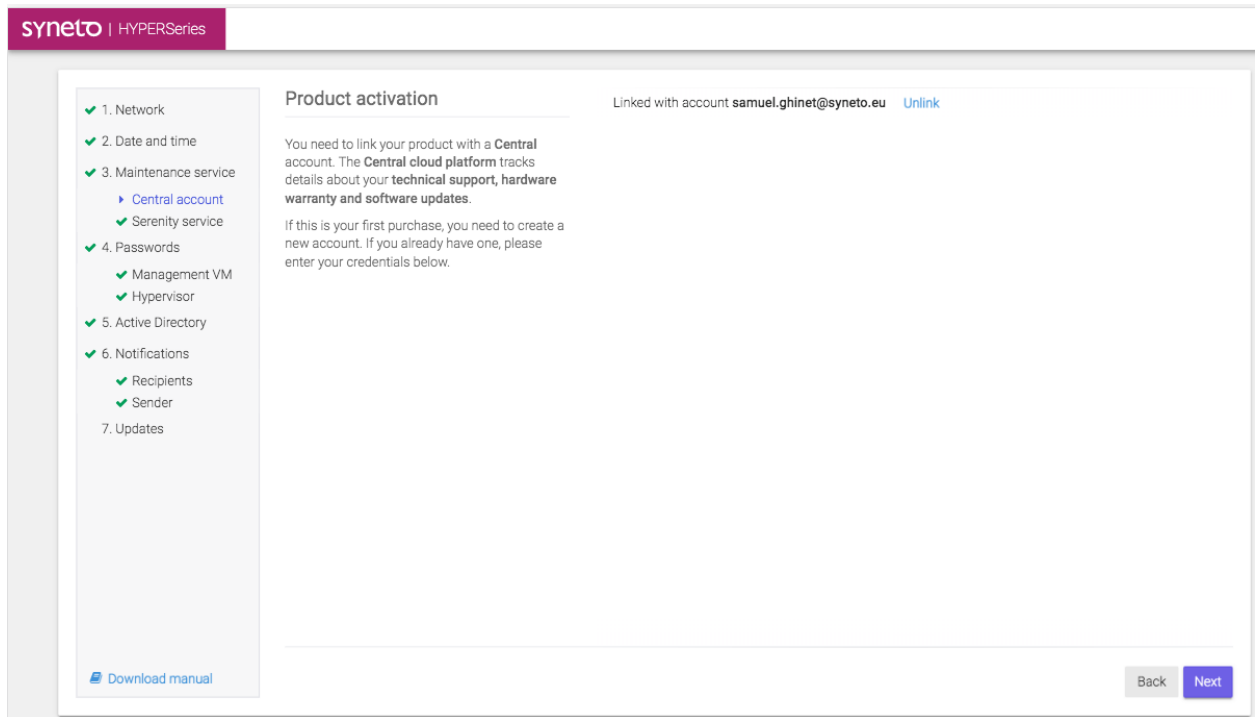
Step 3: Maintenance service → Central account

You need to link your product with a Central account. The Central cloud platform tracks details about your technical support, hardware warranty and software updates.

If this is your first purchase, you need to create a new account. If you already have one, please enter your credentials below.



Linked your account.



Step 3: Maintenance service → Serenity service

Serenity frees your inbox of the machine related issues and reports them directly to Syneto Technical Support for professional investigation.

With Serenity your machine issues are solved without you even knowing.

Serenity: **Inactive**

The Serenity service is not active on this machine. To activate Serenity you should contact your Syneto supplier.

syneto | HYPERSeries

- ✓ 1. Network
- ✓ 2. Date and time
- 3. Maintenance service
 - ✓ Central account
 - Serenity service
- 4. Passwords
 - Management VM
 - Hypervisor
- 5. Active Directory
- 6. Notifications
 - Recipients
 - Sender
- 7. Updates

[Download manual](#)

Serenity service

Serenity frees your inbox from issues and reports generated by this machine. Instead, issues are reported directly to Syneto technical support for detailed investigation.

When Serenity is enabled, Syneto engineers have direct access to investigate any detected issues.

For all confirmed incidents, you will be kept in the loop via the help desk service.

Support plan

Type	Demo	Refresh
Serenity	Inactive ⓘ	
Duration	3 years (2019/08/22 - 2022/08/21)	
Details	Temporary license for demo purposes	

Serenity

The Serenity service is not active. To activate Serenity please contact your Syneto supplier.

Disabled Enabled

Back [Next](#)

You can temporarily stop sending data to the Central servers by disabling the Serenity service. Later on, the Serenity service can be enabled or disabled from the Help page.

Serenity: Active & Disable

syneto | HYPERSeries

- ✓ 1. Network
- ✓ 2. Date and time
- ✓ 3. Maintenance service
 - ✓ Central account
 - Serenity service
- ✓ 4. Passwords
 - ✓ Management VM
 - ✓ Hypervisor
- ✓ 5. Active Directory
- ✓ 6. Notifications
 - ✓ Recipients
 - ✓ Sender
- 7. Updates

[Download manual](#)

Serenity service

Serenity frees your inbox from issues and reports generated by this machine. Instead, issues are reported directly to Syneto technical support for detailed investigation.

When Serenity is enabled, Syneto engineers have direct access to investigate any detected issues.

For all confirmed incidents, you will be kept in the loop via the help desk service.

Support plan

Type	Developers Use
Serenity	Active
Duration	3 years (2019/08/22 - 2022/08/21)
Details	Developers Use

Serenity

You can temporarily stop sending data to the Central servers by disabling the Serenity service. Later on, the service can be re-enabled from the Help page.

Disabled Enabled

Back
Next

Serenity: Active & Enabled

syneto | HYPERSeries

- ✓ 1. Network
- ✓ 2. Date and time
- ✓ 3. Maintenance service
 - ✓ Central account
 - Serenity service
- ✓ 4. Passwords
 - ✓ Management VM
 - ✓ Hypervisor
- ✓ 5. Active Directory
- ✓ 6. Notifications
 - ✓ Recipients
 - ✓ Sender
- 7. Updates

[Download manual](#)

Serenity service

Serenity frees your inbox from issues and reports generated by this machine. Instead, issues are reported directly to Syneto technical support for detailed investigation.

When Serenity is enabled, Syneto engineers have direct access to investigate any detected issues.

For all confirmed incidents, you will be kept in the loop via the help desk service.

Support plan

Type	Developers Use
Serenity	Active
Duration	3 years (2019/08/22 - 2022/08/21)
Details	Developers Use

Serenity

You can temporarily stop sending data to the Central servers by disabling the Serenity service. Later on, the service can be re-enabled from the Help page.

Disabled Enabled

Back
Next

Step 4: Password → Management VM (Syneto OS)


For security reasons we urge you to change the default password for Management VM (Syneto OS).

The screenshot shows the SynetoOS configuration interface. At the top left, there is a purple header with the text "syneto | HYPERSeries". On the left side, there is a vertical navigation menu with the following items: "1. Network", "2. Date and time", "3. Maintenance service" (with sub-items "Central account" and "Serenity service"), "4. Passwords" (with a sub-item "Management VM Hypervisor" highlighted in blue), "5. Active Directory", "6. Notifications" (with sub-items "Recipients" and "Sender"), and "7. Updates". At the bottom of the menu is a link "Download manual". The main content area is titled "Passwords" and contains the text: "For security reasons we urge you to change the default passwords." Below this text is a yellow warning box with the text: "WARNING: Once the default passwords are changed, it is important to remember them. In case they are forgotten there is no way to reset them." To the right of the warning box, there is a section titled "Management VM (Syneto OS)" with two input fields: "New password" and "Re-type new password". At the bottom right of the main content area, there are "Back" and "Next" buttons.

NOTE: Once the default passwords are changed, it is important to remember them. In case they are forgotten there is no way to reset them.

Step 4: Password → Hypervisor (ESXi)

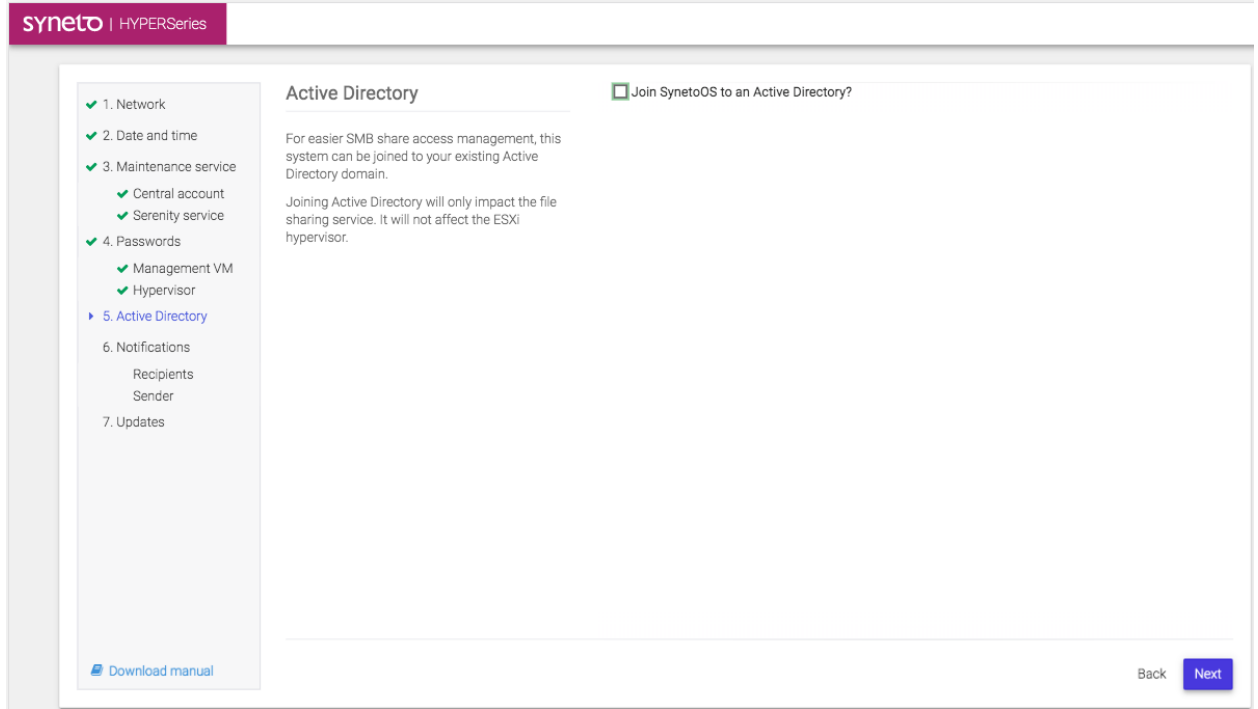
For security reasons we urge you to change the default password for Hypervisor (ESXi).



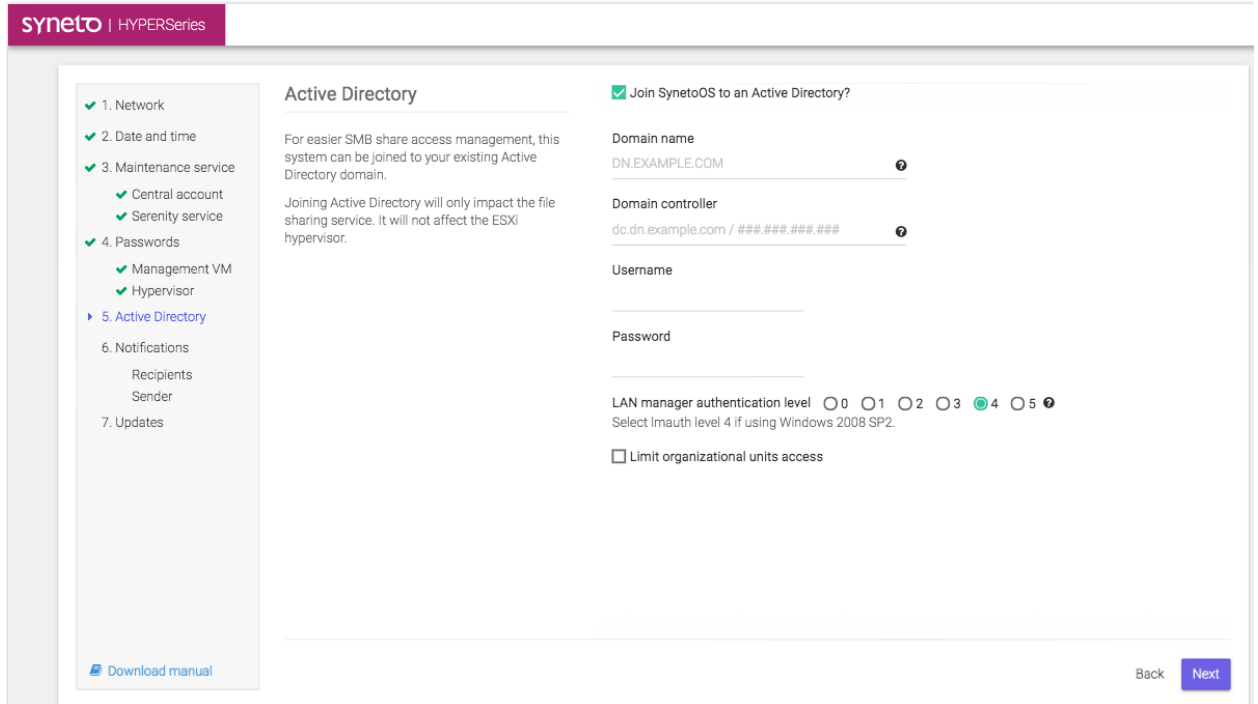
NOTE: Once the default passwords are changed, it is important to remember them. In case they are forgotten there is no way to reset them.

Step 5: Active Directory

For easier SMB share access management, this system can be joined to your existing Active Directory domain.



Check → *Join to an Active Directory*

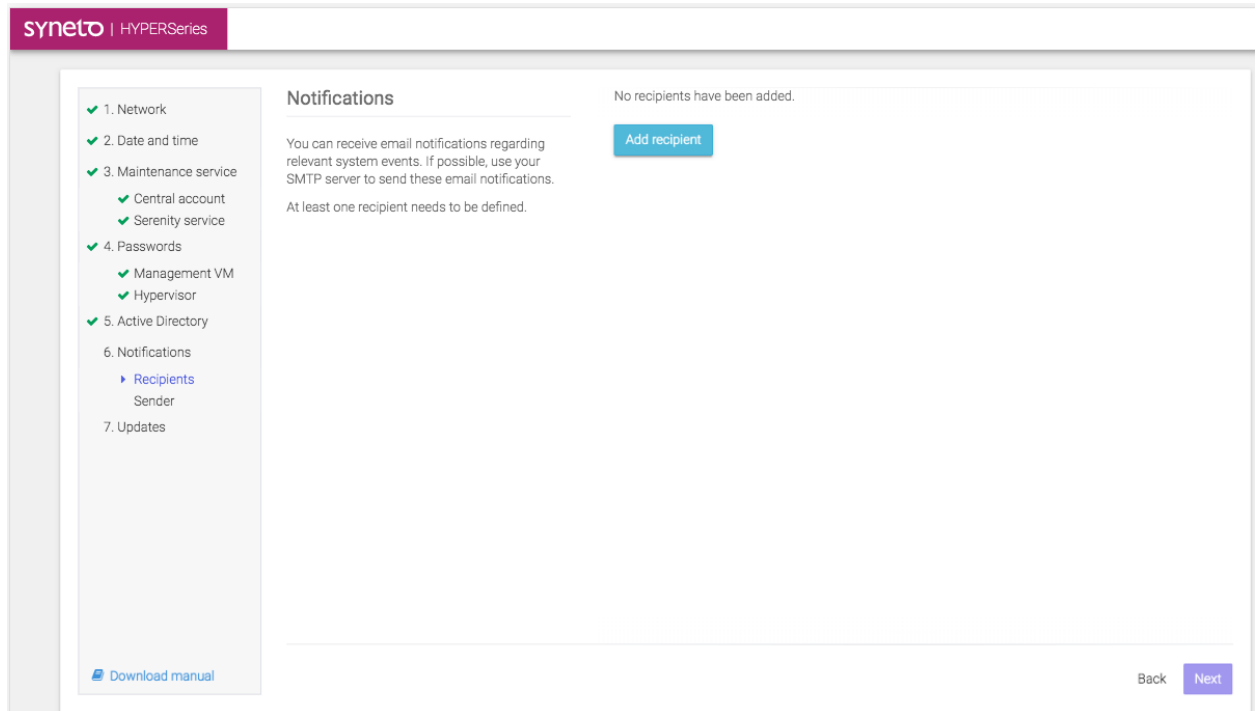


NOTE: Since SynetoOS 4.4.2, *Integrity check* can no longer be configured from the *Quick setup*. Instead, it will automatically be configured to run every 4 weeks.

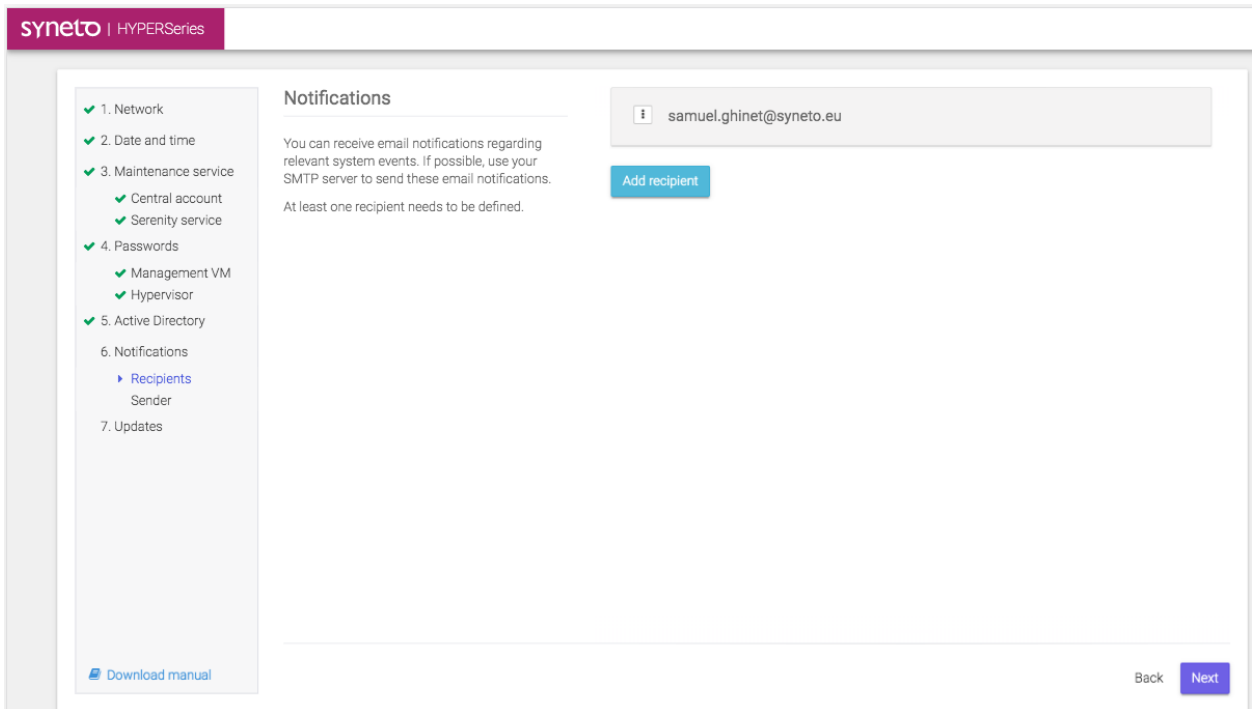
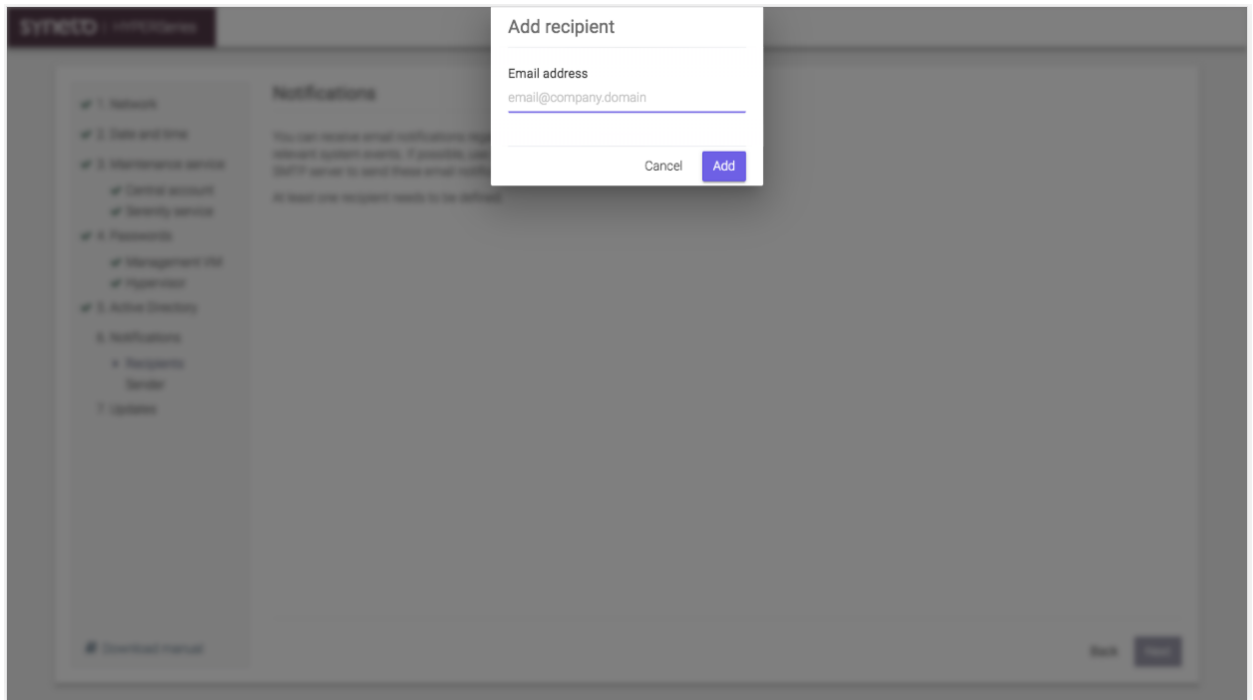
Step 6: Notification → Recipient

You can receive email notifications regarding relevant system events. If possible, use your SMTP server to send these email notifications.

At least one recipient needs to be defined.



Add recipient → Insert your email



Step 6: Notification → Sender

Configure an SMTP server to send these email notifications.

The screenshot shows the SynetoOS configuration interface for the 'Notifications' step. On the left, a navigation menu lists steps 1 through 7, with 'Notifications' selected. The main content area is titled 'Notifications' and contains instructions: 'You can receive email notifications regarding relevant system events. If possible, use your SMTP server to send these email notifications. At least one recipient needs to be defined.' To the right, the 'Outgoing email server' is set to 'SMTP'. The 'Server' field is 'smtp.gmail.com' and the 'Port' is '465'. Two checkboxes are checked: 'Use secure connection - set the proper port number for secure connection' and 'Server requires authentication'. The 'Username' field is 'samuel.ghinet@syneto.eu' and the 'Password' field is masked with dots. The 'Sender email' field is 'no-reply@syneto-os-2ff8aaef'. A 'Send test email' button is located below the password field. At the bottom right, there are 'Back' and 'Next' buttons. A 'Download manual' link is at the bottom left.

syneto | HYPERSeries

Notifications

You can receive email notifications regarding relevant system events. If possible, use your SMTP server to send these email notifications.
At least one recipient needs to be defined.

Outgoing email server Local service SMTP

Server	Port
smtp.gmail.com	465

Use secure connection - set the proper port number for secure connection

Server requires authentication

Username
samuel.ghinet@syneto.eu

Password
.....

Sender email
no-reply@syneto-os-2ff8aaef

[Send test email](#)

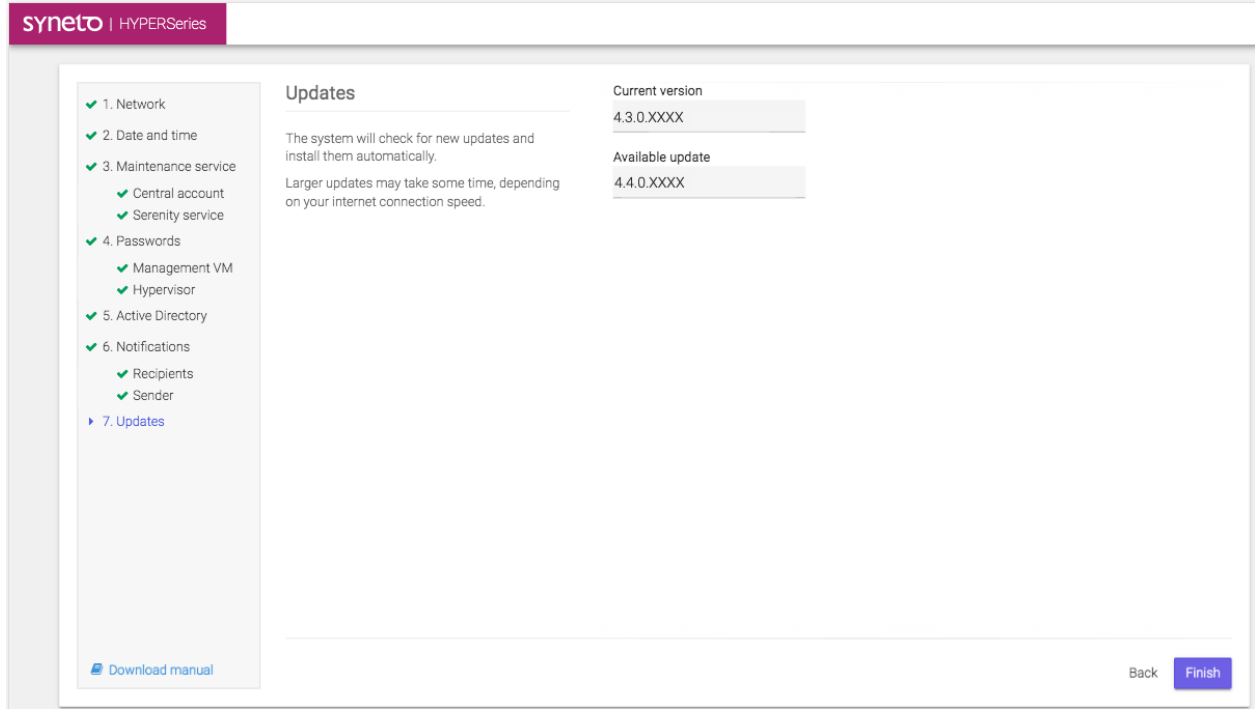
[Download manual](#)

Back [Next](#)

Step 7: Updates

The system will check for new updates and install them automatically.

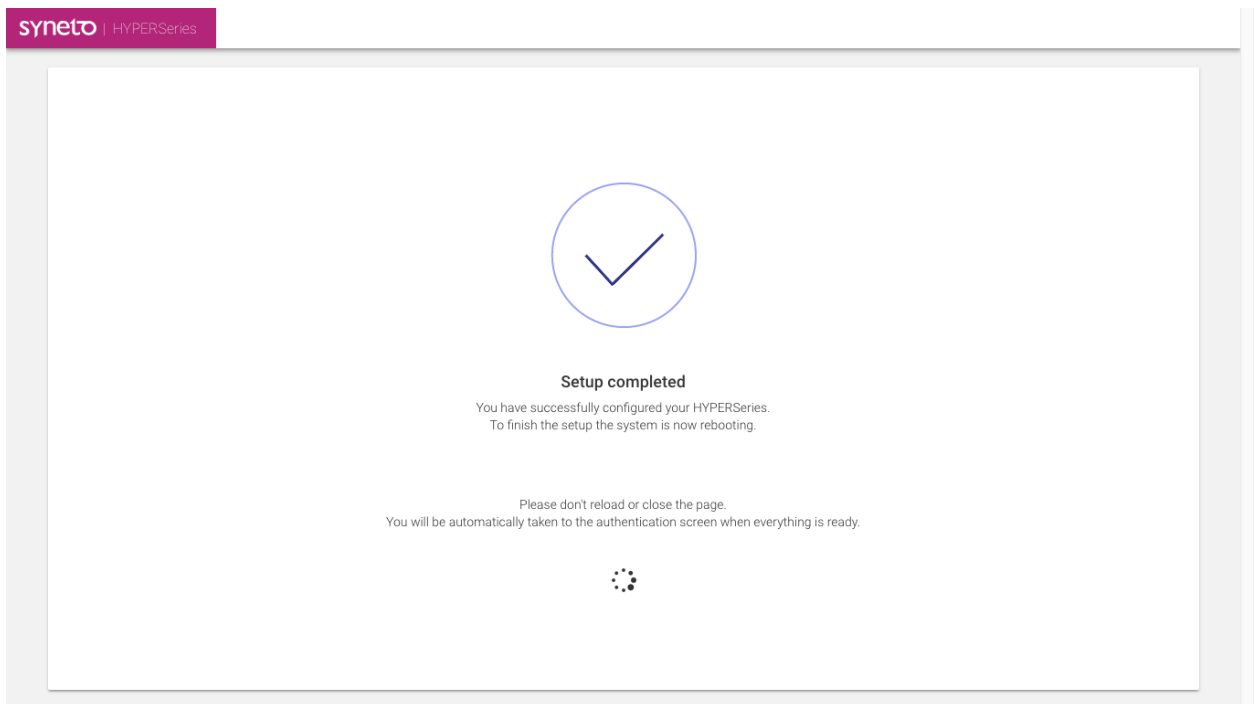
Larger updates may take some time, depending on your internet connection speed.



Step: Finish

Please don't reload or close the page.


You will be automatically taken to the authentication screen when everything is ready.



Activate the appliance

Every product has an attached support service, which entitles it to software updates and remote customer support. This support service is managed using an online portal called Syneto Central.

When you try to log in for the first time on the management VM, a Central account activation screen will appear. This chapter shows you how to activate your product.



syneto | HYPERSeries

Welcome to HYPERSeries

To begin, please link your product with a **Central** account. The **Central cloud platform** tracks details about your **technical support, hardware warranty and software updates**.

If this is your first purchase, you need to create a new account. If you already have one, please enter your credentials below.

Serial number
SYN00000122

Central account email

Password

[Forgot password?](#) [Create account](#) [Activate](#)

[Need help?](#)

NOTE: Log in to the management UI is not possible without registering the product to Syneto Central first.

Upon receiving your Syneto HYPER, it will already be registered with Central, but in order to work on it you must create a Central account. Please read [Appendix A](#) to learn how to create a Central account.

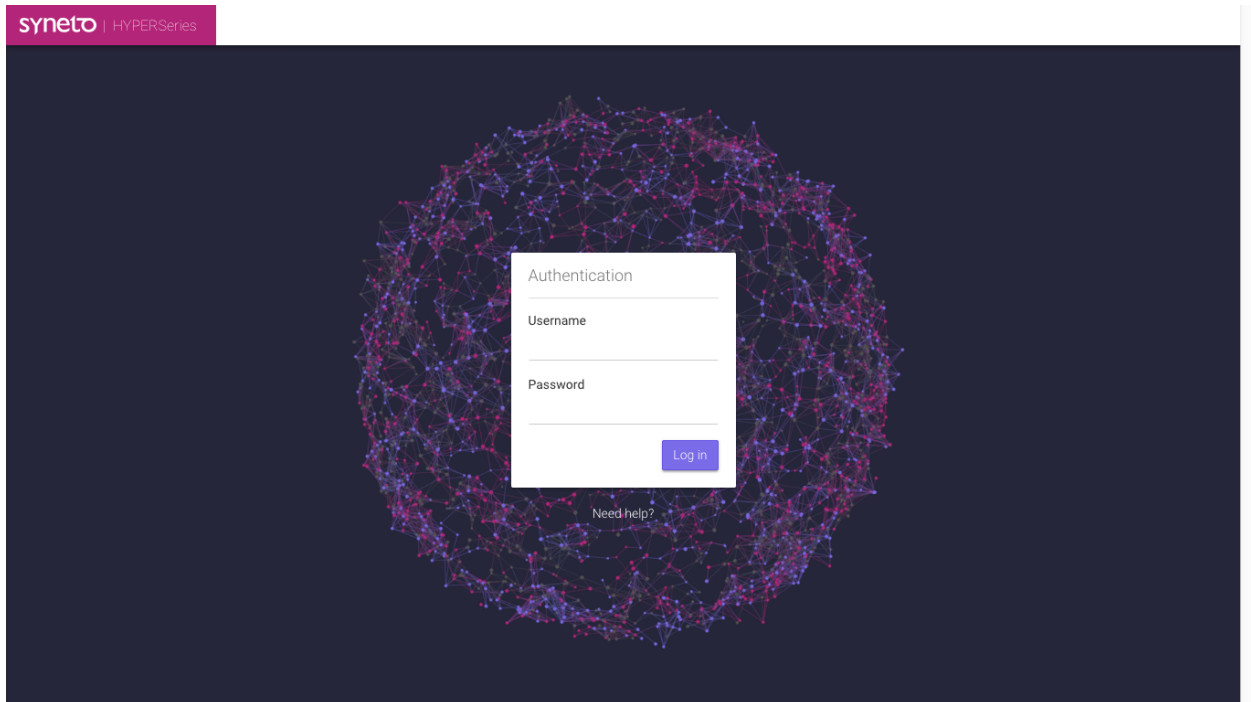
After creating the account, enter its email address and password and click *Activate*.

After the product is successfully activated, you will be able to log in to the management interface.

If the activation fails, it's likely that HYPER can't connect to the Central online service. Please make sure that a valid DNS and gateway are configured, and that there are no firewall rules blocking the access to `central.api.syneto.eu` on port 443.

Access the management interface

Open a web browser(Chrome, Firefox, Safari etc) and write the management IP address. At this point, you should have already activated your product with Central, as shown in the chapter above.



Configure using IPMI

Each Syneto HYPER comes with IPMI (or similar) remote management support.

In order to have access to the IPMI interface, make sure a network cable is connected to the IPMI ethernet port.



The IPMI interface will acquire a network address over DHCP.

Make sure a display is connected to the Syneto HYPHER and note the management IP address shown during boot. See bottom right in the image below: `BCM IP:`

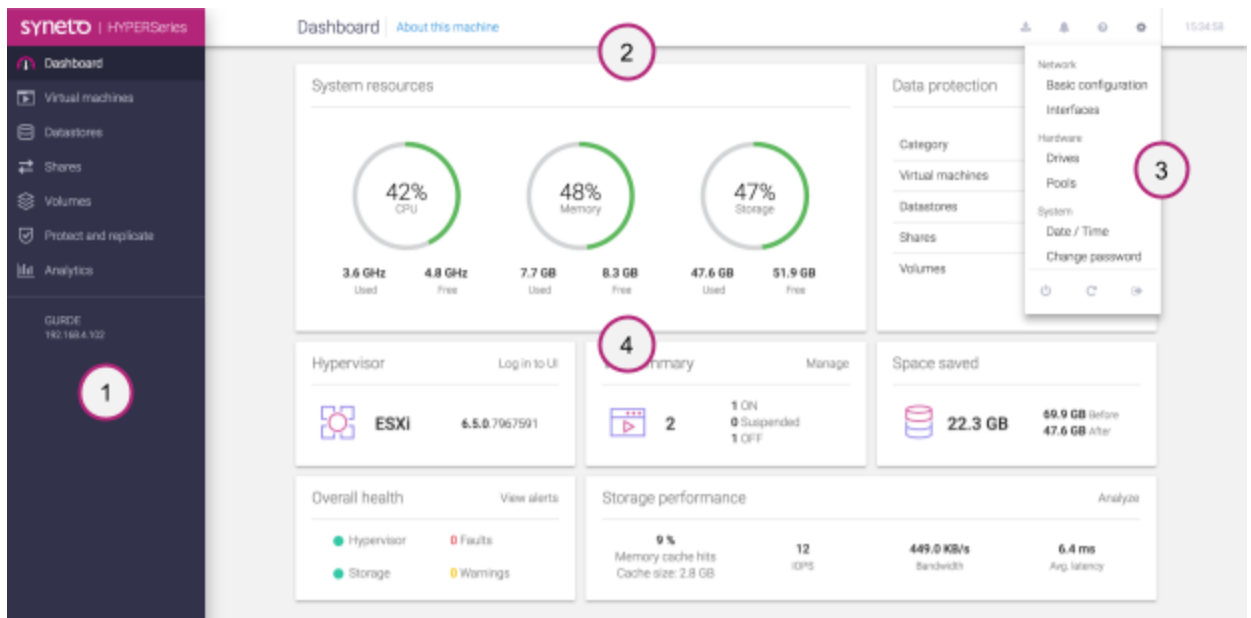


Default credentials

Once you know the IP, open a browser on your computer, and connect to it over the HTTP or HTTPS protocol. The default IPMI credentials are: ADMIN/ADMIN (all caps).

The screenshot shows the Supermicro logo at the top left. Below it is a login form with the title "Please Login". The form contains two input fields: "Username" with the value "ADMIN" and "Password" with the value "*****". A "login" button is located below the password field.

Management interface main components



These are the main components of the management interface:

1. **Main menu** - contains the most important items related to the Syneto HYPER management. Expect to use it frequently.
2. **Action bar** - presents actions related to the current page, and a set of notifications on the right for update, alarms, expired license, missing central connection. Finally there is a button to open the secondary menu.
3. **Secondary menu** - expect to find here the rarely used options. We included the things that you will usually do once.
4. **Main area** - the biggest portion of the screen is dedicated for the area where you will be performing all the administrative tasks of the HYPER machine.

INITIAL CONFIGURATION

Configure ESXi networking

NOTE: Syneto Hyper series comes pre-configured with a virtual switch - *vSwitch1* - that has no uplink defined, and has two port groups: '*Syneto HC VMKernel Network*' and '*Syneto HC API Network*'. The '*Syneto HC VMKernel Network*' has a *vmk1* VMkernel interface attached and configured with a static IP address of *172.16.254.1*. The SynetoOS VM has the '*Syneto HC API Network*' attached to *vmxnet3s1* and configured with a static IP address of *172.16.254.2*. The vSwitch, port groups, vmk1 and static IP addresses documented here **must not be modified**. If you believe that in your case you must change any of these, please contact Syneto customer support first.

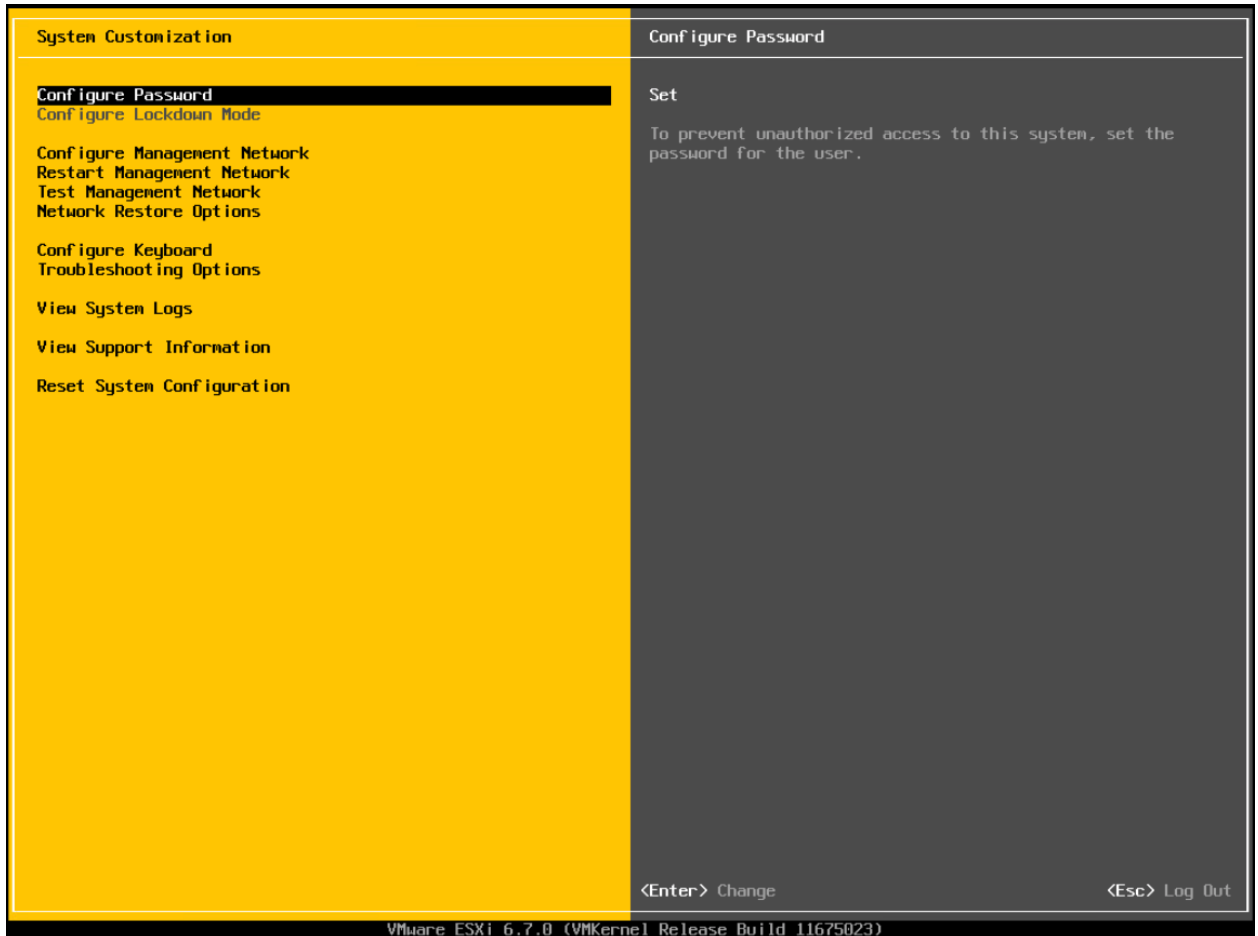
Access the ESXi console. This can be done by manually connecting a monitor and keyboard to the physical machine or through IPMI.

```
VMware ESXi 6.7.0 (VMKernel Release Build 11675023)
VMware, Inc. VMware Virtual Platform
Intel(R) Xeon(R) CPU E5-2620 v2 @ 2.10GHz
12 GiB Memory

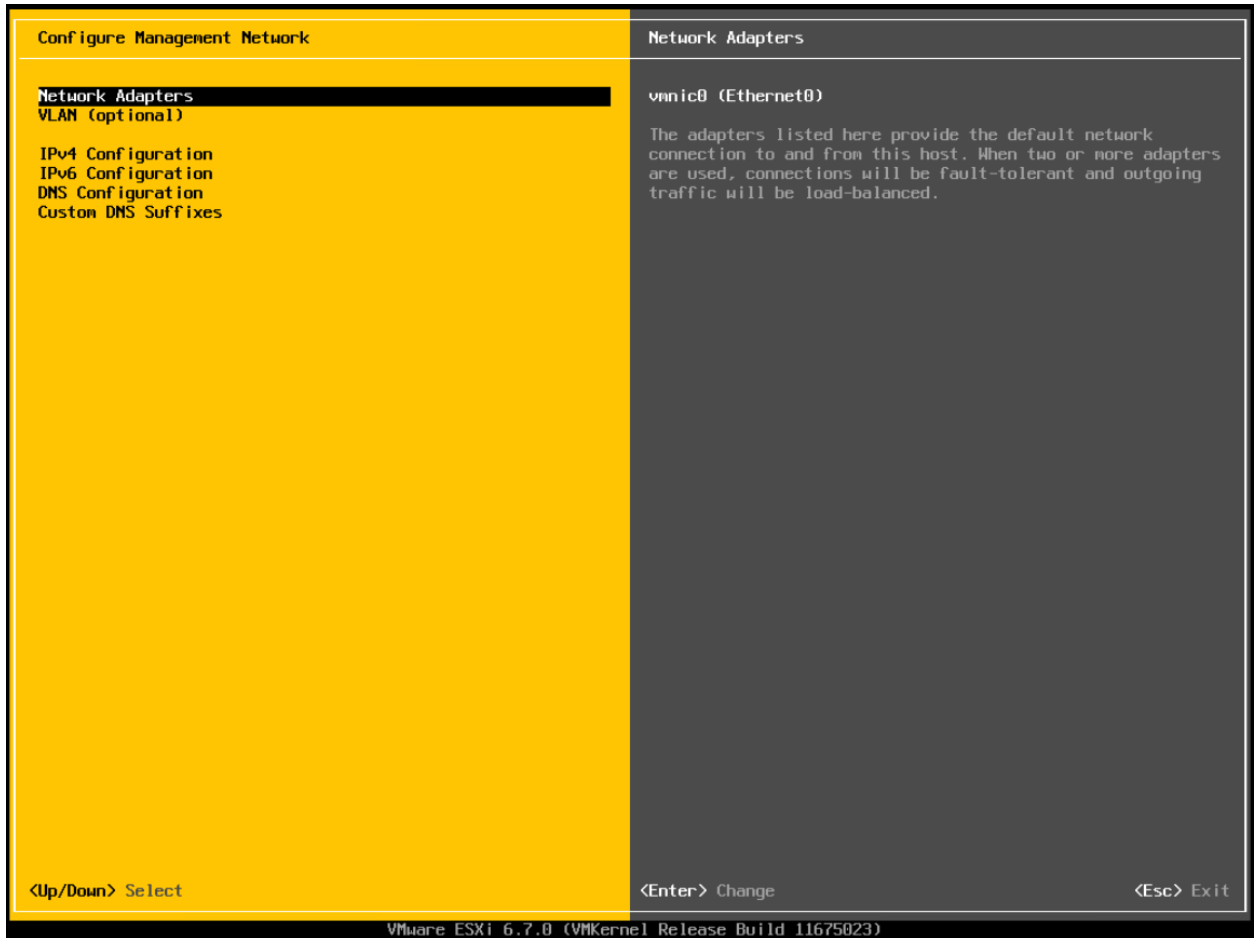
To manage this host go to:
http://syneto-esxi-ba87c8c4/
http://192.168.5.52/ (DHCP)
http://[fe80::250:56ff:fe87:b0b21]/ (STATIC)

<F2> Customize System/View Logs
<F12> Shut Down/Restart
```


Press *F2* and then enter the *username* and *password* of the ESXi. A menu with different actions will open.



Select *Configure Management Network* and press *Enter*. A menu with network configuration options will open.



Select *IPv4 Configuration* and press *Enter*. A dialog for IPv4 Configuration will open.



You can select a dynamic or static IP. By default, the dynamic IP option is selected. Also, the *gateway* can be set from this dialog. After you have finished configuring , press *Enter* to save and exit. Select *DNS Configuration* and press *Enter*.

```
DNS Configuration

This host can only obtain DNS settings automatically if it also obtains
its IP configuration automatically.

( ) Obtain DNS server addresses and a hostname automatically
(o) Use the following DNS server addresses and hostname:

Primary DNS Server      [ 192.168.1.15          ]
Alternate DNS Server    [                    ]
Hostname                [ syneto-esxi-039a3038 ]

<Up/Down> Select  <Space> Mark Selected      <Enter> OK  <Esc> Cancel
```

You can select a dynamic or static DNS and hostname. If you select static, fill the *Primary DNS Server* and *Hostname* fields. The *Alternate DNS Server* field is optional. Press *Enter* to save and exit.

Install ESXi license


Go to the ESXi management interface. Log in with your username and password.

```
vmware

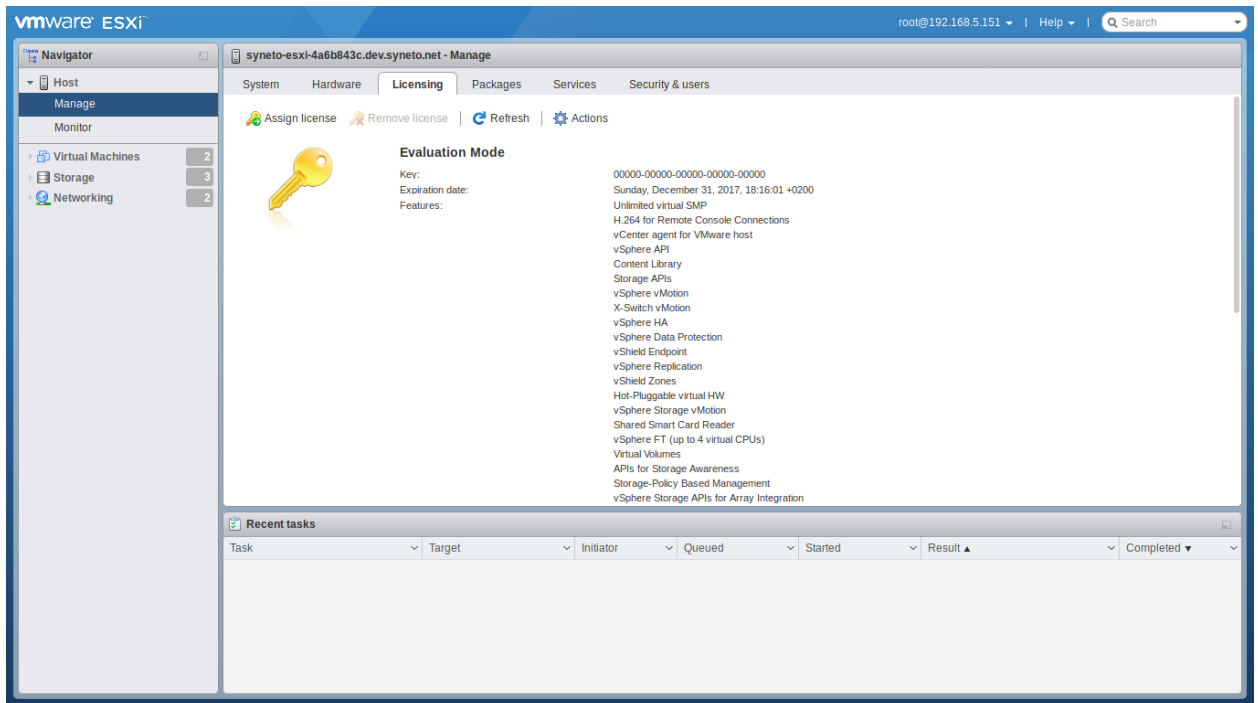
User name: root
Password: *****
Log in

vmware ESXi


Open the VMware Host Client documentation
```



Select *Host* from the *Navigator* on the left and click on *Manage*.



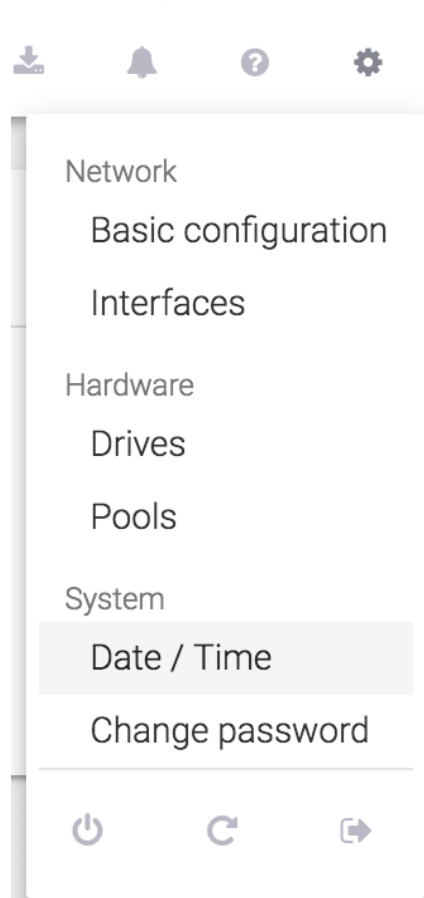
Select *Licensing* tab → *Assign License*.



Type or paste the ESXi license key. Click *Check License*.

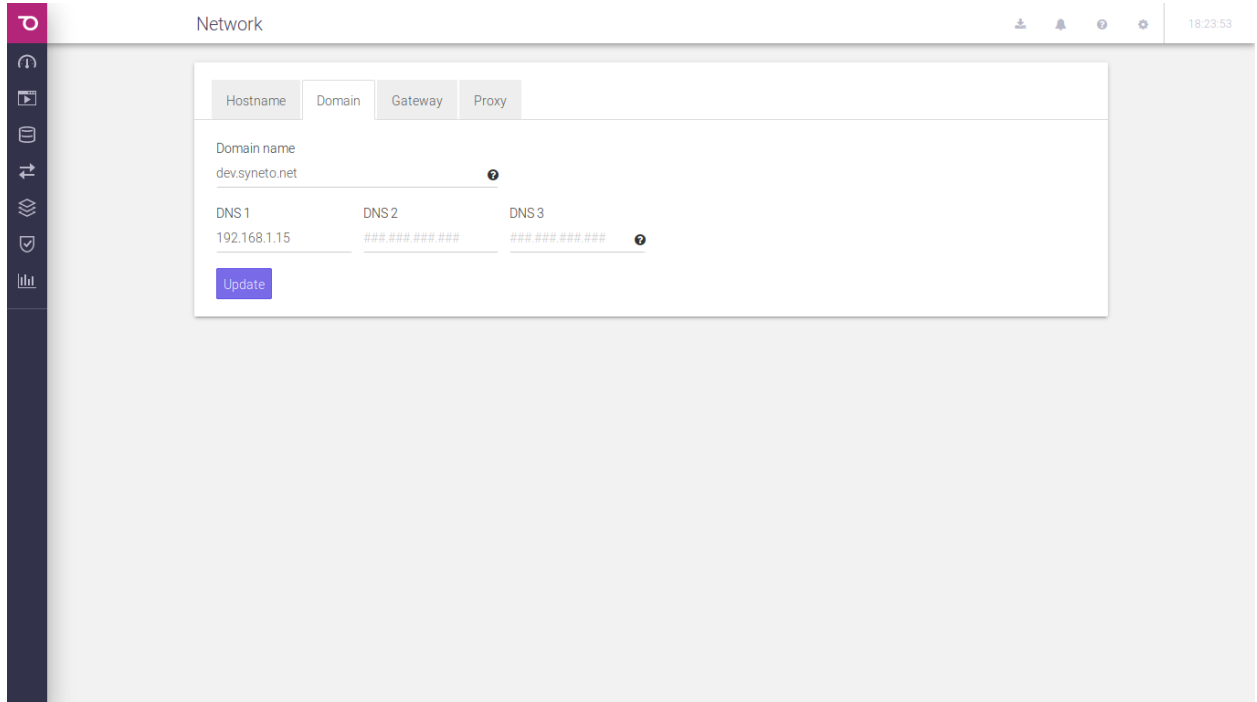
Configure Management VM networking

Log in to SynetoOS Management Interface. Go to *Secondary Menu* → *Network* → *Basic*



Configuration.

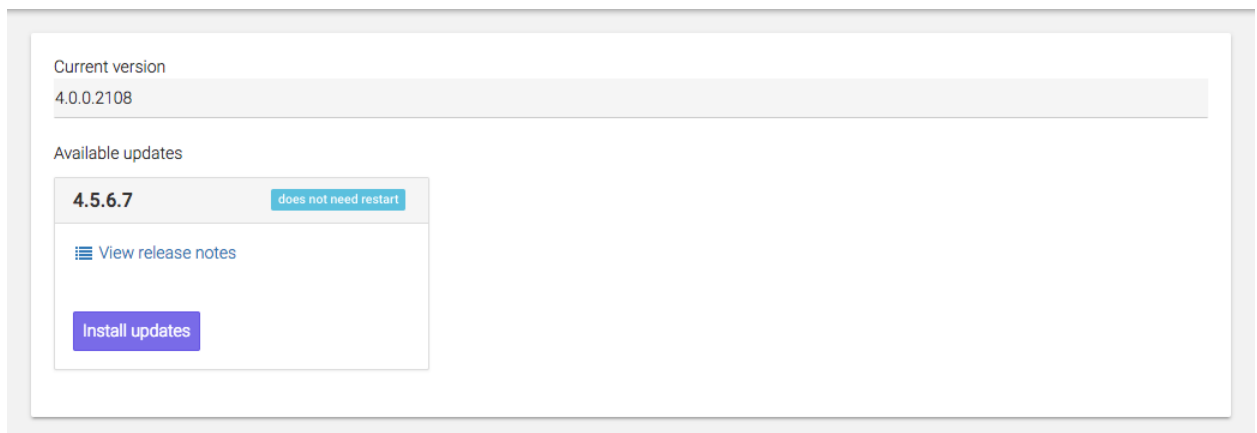
The *Network* page will open. Here you can set the hostname, domain name, DNS, gateway and proxy server.



For more information, [check the networking section](#).

Update Syneto HYPER to latest version

Click on the *Software Update* button, in the action bar. It is the first button from the left. A new page will be opened. SynetoOS will automatically check for a new version. If there is a new version, a button *Update* will appear. Click on it and the update will start.




For more information, check [the software update section](#).

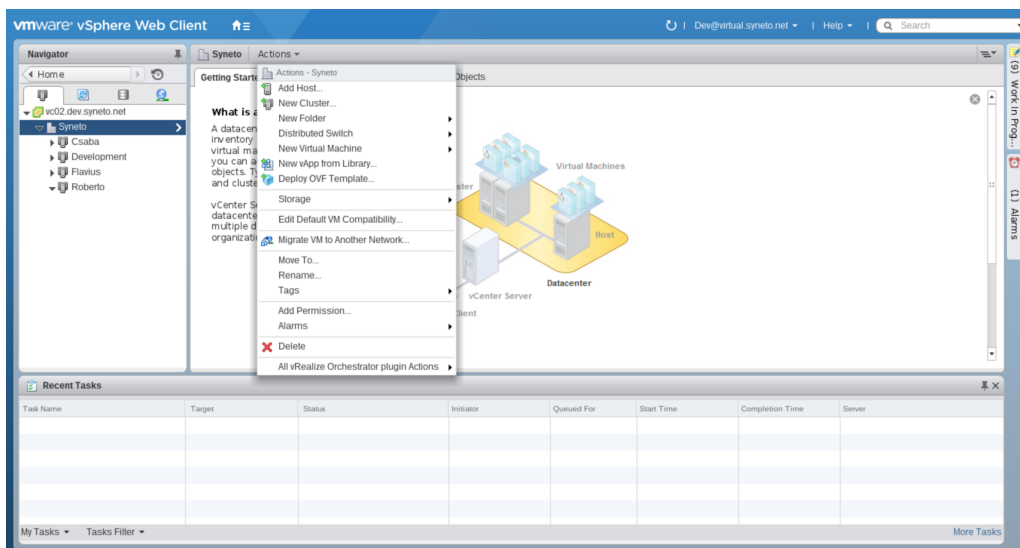
Join vCenter

WARNING: This feature is not available on Syneto HYPER 2000 Series.


If you have a vCenter you can join our solution to it for easier management.
Log in to vCenter Management Interface. In the left menu, select *Hosts and Clusters*.



Select *Actions* → *Add Host*.



A dialog will open which will guide you through the process of adding a host. Enter the hostname or IP of the Syneto ESXi. Click **Next**.



The screenshot shows the 'Add Host' dialog box in the vSphere Web Client. The dialog is titled 'Add Host' and has a progress indicator on the right. The left sidebar shows a list of steps: 1 Name and location (selected), 2 Connection settings, 3 Host summary, 4 VM location, and 5 Ready to complete. The main area contains the following text and fields:

Enter the name or IP address of the host to add to vCenter Server.


Host name or IP address:

Location:

Type: ⓘ

At the bottom right, there are four buttons: Back, Next, Finish, and Cancel.

Enter the username and password of the Syneto ESXi. Click **Next**. If a security alert opens, click **OK**.



The screenshot shows the 'Add Host' dialog box in the vSphere Web Client, now at Step 2: Connection settings. The left sidebar shows the progress: 1 Name and location (checked), 2 Connection settings (selected), 3 Host summary, 4 VM location, and 5 Ready to complete. The main area contains the following text and fields:


Enter the administrative account information for the host. The vSphere Web Client will use this information to connect to the host and establish a permanent account for its operations.

User name:

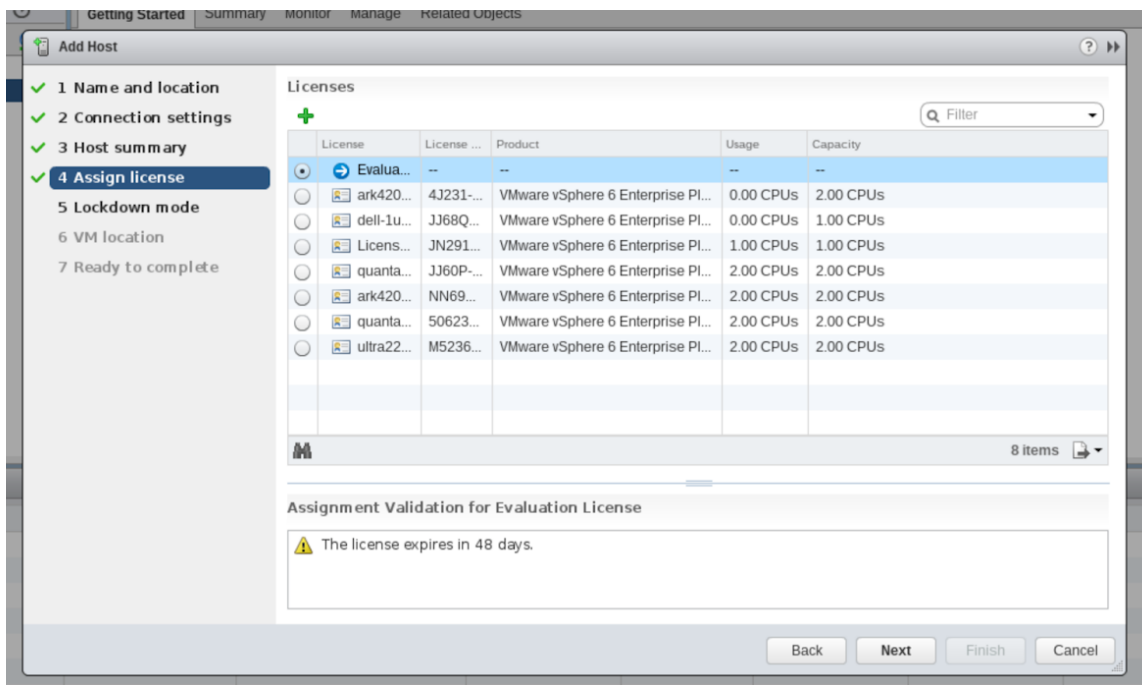
Password:

At the bottom right, there are four buttons: Back, Next, Finish, and Cancel.

In this view, you can review the host configuration. After that, click *Next*.



You will see a list of available licenses. You can assign one of those or add a new one by clicking on the *green plus*.




This window lets you configure the lockdown mode.

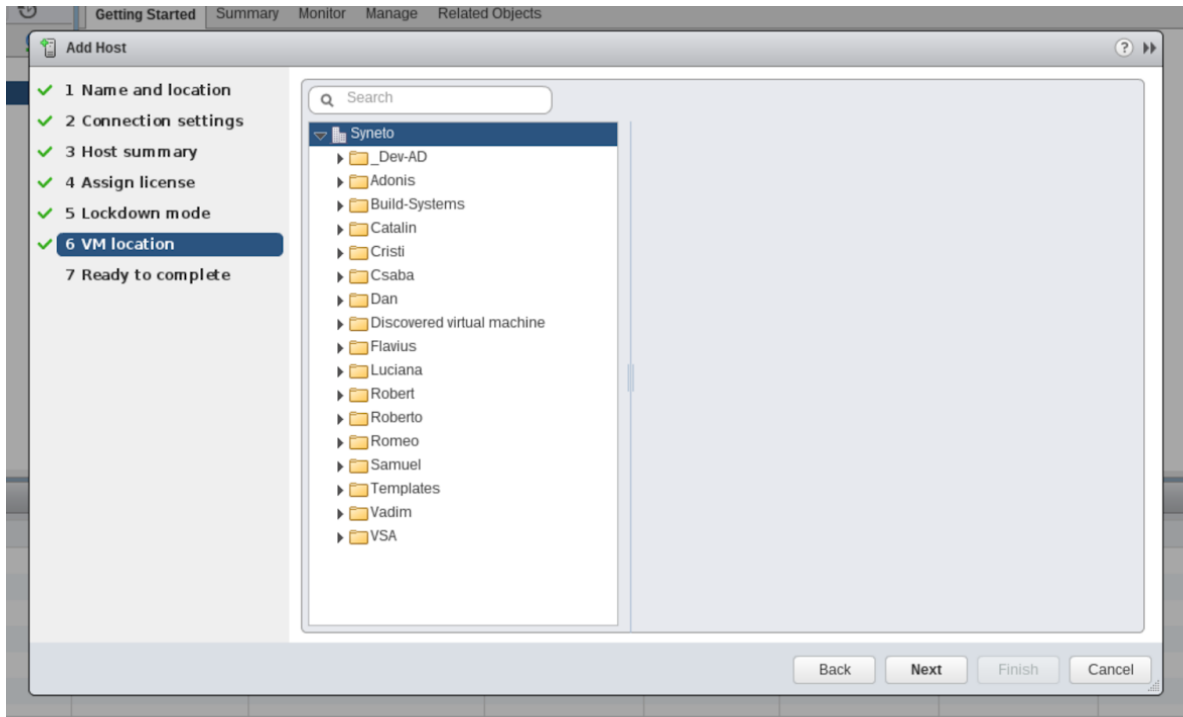
When enabled, lockdown mode prevents remote users from logging directly into this host. The host will only be accessible through local console or an authorized centralized management application.

If you are unsure what to do, leave lockdown mode disabled. You can configure lockdown mode later by editing Security Profile in host settings.

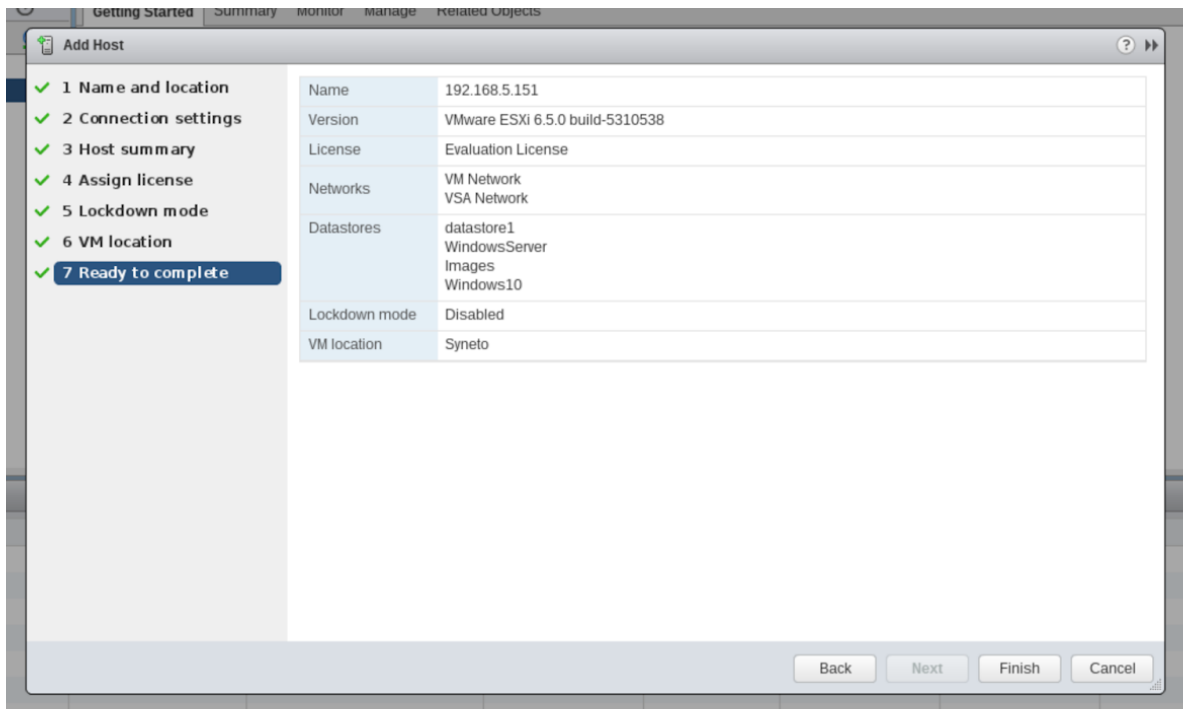
After you finish, click *Next*.



Select the virtual machine location. Click *Next*.



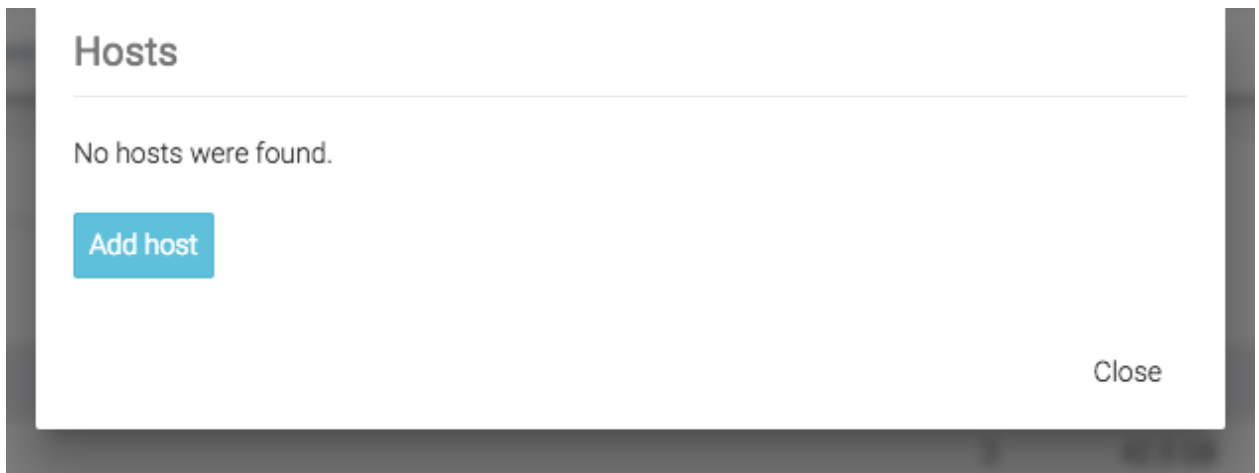
Review your selections and if you do not want to modify anything, click *Finish*.




Now, we need to add vCenter to SynetoOS. On the SynetoOS Management Interface, go to the *Datstores* page.



Click on *External hosts* in the action bar, top of the page. A dialog listing the registered external hosts will open.



Click on *Add host*. A form will open. In the *host* field, enter the hostname or IP of the vcenter. Enter a *username* and its *password*. Click *Add*.

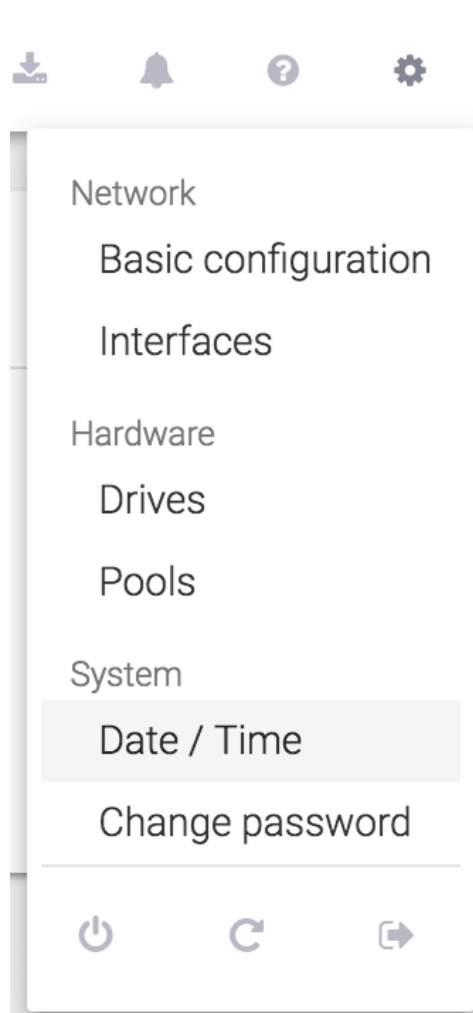
The screenshot shows a form titled 'Add host'. It contains three input fields: 'Host', 'Username', and 'Password'. The 'Host' field has a blue underline. At the bottom of the form, there are two buttons: 'Cancel' and 'Add'.

This will add the vCenter to SynetoOS. Now, you will be able to mount datastores on any of the vCenter's hosts.

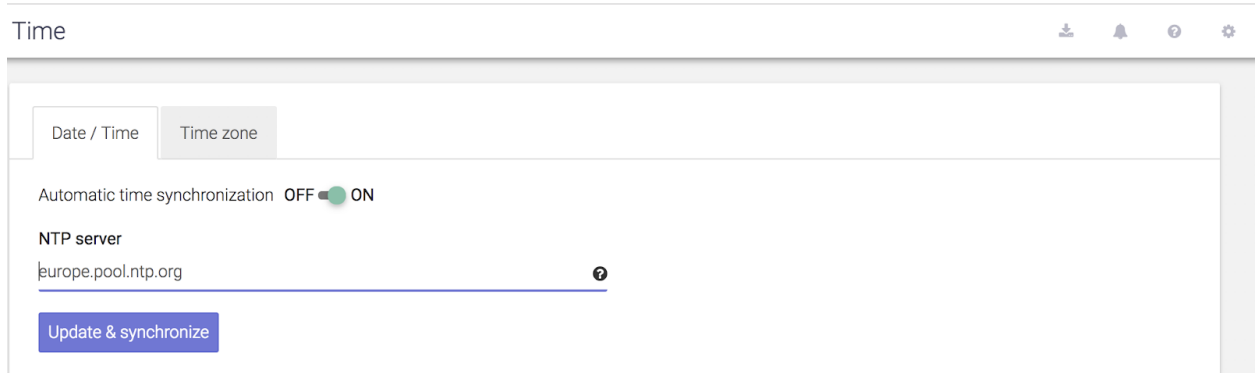
For more information about managing non-Syneto hosts, see chapter [Serving non-Syneto hosts](#)

Configure date / time

Log in SynetoOS Management Interface. Go to *Secondary Menu*.



Click on *Date / Time* option. The *Time* page will open. Here you can set the date, time, and timezone. We recommend setting *Automatic time synchronization* on.




For more information, [check the date/time section](#).

Email server and Alert thresholds

Syneto wants to keep you well informed about what goes around with your machine. Alerts and email notifications are very useful in managing the few errors that may appear.

It is highly recommended that you introduce at least one email address in order to receive emails in case there are any issues on your machine.

You can either use your local email service, if there is no firewall blocking it, or you can use an SMTP server.

To add an email addresses, go to the *Alerts* page, the bell icon  from the action bar.

To learn more about adding email addresses to receive fast and accurate alerts and notifications, please go to [Alerts and notification](#) chapter.

NOTE: Some email servers may mark these emails as spam, please also check your spam folder.

Alternatively, an SMTP server can be used to send notifications emails. In case you have an SMTP server please use this option.

Recommended thresholds

Keep disk pool space under 80% utilization for best performance. Disk pool performance can degrade when it is close to full.

WARNING: Keep in mind that even with mostly static content in the 95-96% range, write, read, and resilvering performance might suffer.

Define recurring integrity check intervals

Syneto products have a built-in mechanism for protecting against silent data corruption. Data corruption means that data being read is different from the data previously written to a drive. Some errors go unnoticed, without being detected by the drive firmware or the host operating system; these errors are known as silent data corruption.

To prevent silent data corruption, Syneto uses self-healing which happens behind the scenes as blocks are being read into memory.

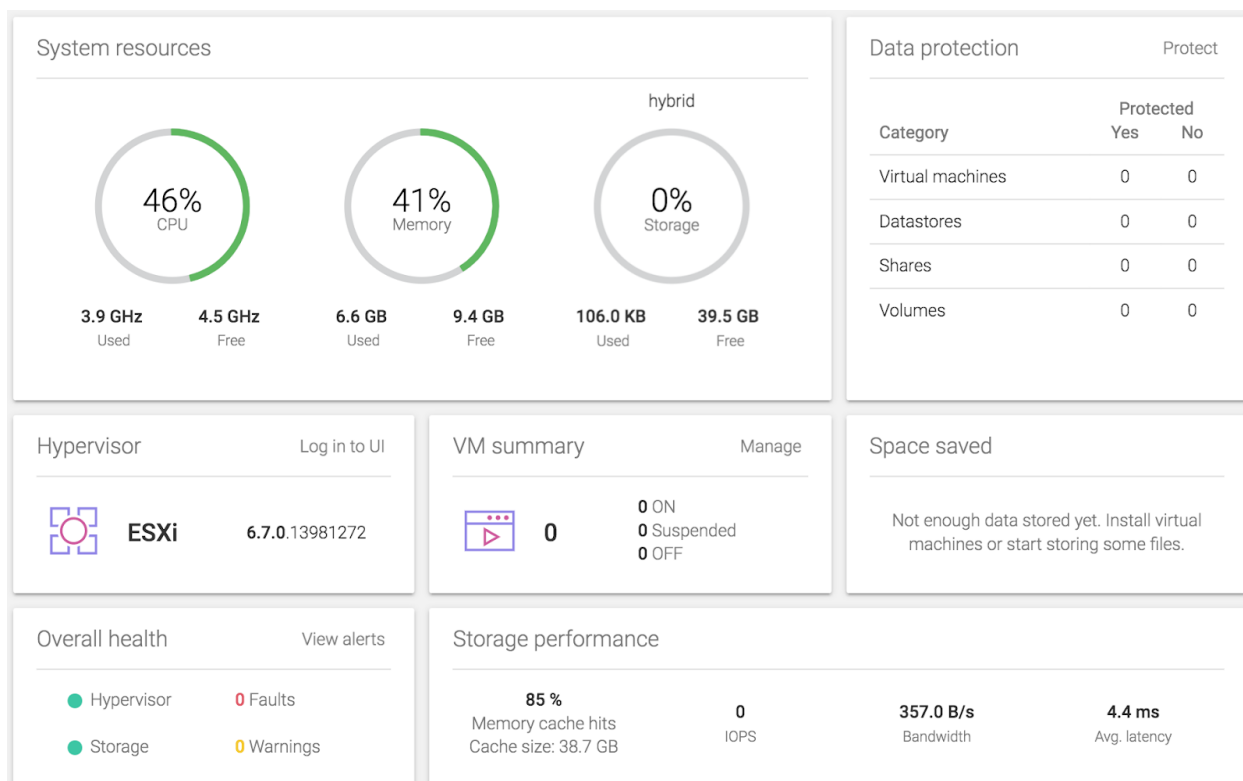
Self-healing protects especially hot data, which is accessed frequently and thus checked frequently for its consistency/integrity. Less frequently accessed (cold) data is at a larger risk because its blocks may accumulate multiple errors, which could render checksums useless. To protect cold data, the system is configured by default to periodically run a scan of the drives and fix corruptions. This process can also be initiated manually.

To see more about Integrity check, please go to [Protecting against silent data corruption](#) chapter.

THE DASHBOARD

After you log in to the UI, you will be greeted by the *Dashboard*.

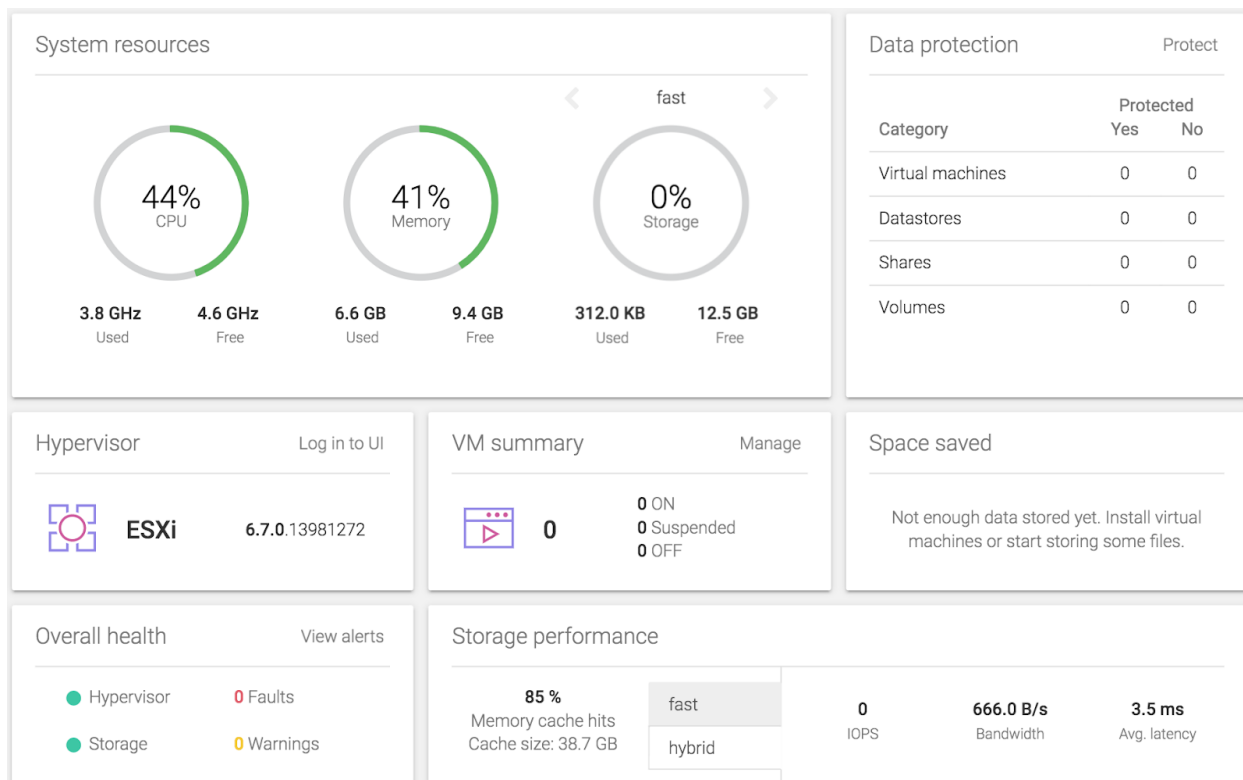
Here you will have a quick overview of the *System resources*, *Data protection*, *Hypervisor*, *VM summary*, *Space saved*, *Overall health*, and the *Storage performance*. By following the links on each widget, you can explore all the functionalities of Syneto HYPER.



NOTE: The data protection widget does count clones.

If you have more than one pool, then a pair of angle brackets will appear at the top (to the left and to the right of the pool name), which will allow you to switch among the pools.

In the image below, a new pool was created. So now there are two pools: *hybrid* and *fast*. Only the stats of one pool can be seen at a time. To switch from the *fast* pool to the *hybrid* and vice versa you use the angle brackets.



As you populate your machine with data (*Virtual machines, Volumes* (This feature is not available on Syneto HYPER 2000 Series), *Shares, Datastores*), the *Dashboard* will reflect the real time changes.

Here you can easily track the protection level of your data, the health of the system, as you will receive real time feedback from the *Hypervisor* and the storage. You can track the *CPU* and easily manage the space handled by the machine.

On the *Actions bar*, you can see information about the machine with a simple click on *About this machine*: The Central status and detailed software and hardware info will be shown.

About this machine

Central [↗](#)
Linked with account **name.surname@company.eu** [Unlink](#)

Hardware

ID	9f784abf9c979e77f191457b99ebb453
Serial number	VMware-42 07 3c 32 25 54 47 23-ef 00 f5 67 b1 08 ce 9c
Processor	4x GenuineIntel Intel(R) Xeon(R) CPU E5-2618L v2 @ 2.00GHz
Memory	4.0 GB
Data drives	1
Data storage (raw)	2.0 TB
Network	2 x 10 Gb/s

Software

Software version	4.2.0.2857
Last system boot	2018-07-13 14:37
Uptime	16:02:26 up 1:25, 1 user, load average: 1.75, 1.30, 1.04

[Close](#)

From the *About this machine* dialog, you can also *Unlink* from Central by clicking *Unlink*. A confirmation dialog will pop out saying: “You will be logged out and need to activate this machine with a new Central account in order to access the UI again.” Should you agree with the terms, the product will be unlinked from the account.


Confirm

Are you sure you want to unlink your product from Central?

You will be logged out and need to activate this machine with a new Central account in order to access the UI again.

[Cancel](#) [Unlink](#)

You will be redirected to a form in order to insert the new central account you want the product to be linked to. Complete the new email address and the password. Press *Activate*. You can now log in to your product again.



Welcome to HYPERSeries

To begin, please link your product with a **Central** account. The **Central cloud platform** tracks details about your **technical support, hardware warranty and software updates**.

If this is your first purchase, you need to create a new account. If you already have one, please enter your credentials below.

Serial number
SYN0001

Central account
name.surname@company.eu

Password

[Forgot password?](#) [Create account](#) [Activate](#)

WARNING: If there is no gateway server set, a working DNS configured, or network connection, the Central portal will not be reachable.

DATASTORES


Introduction

Datstores are entities that allow you to share storage with VMware. When you create a new datastore on SynetoOS, a series of entities are created automatically. There will be underlying file system, a network share, and a VMware datastore.


Provisioning and managing

Provision new datastore

Your new storage comes with no predefined datstores. When you go to *Main menu* → *Datstores* you will see an empty list.



To create a new datastore click *New datastore* in the action bar.

A screenshot of the "New datastore" form. The form has a title "New datastore" and a "Name" field with the value "Windows10". Below the name field is a "Workload type" section with three radio button options: "Database", "Application server", and "File storage". The "Application server" option is selected. At the bottom left, it shows "39.7 GB free space". At the bottom right, there are "Cancel" and "Create" buttons.

Provide a name for the datastore, choose a workload type, and then click *Create*.

The workload types are as follows:

- *Database* - for virtual machines running database servers. The datastore will be configured with a block size of 4 KB.

- *Application Server* - for virtual machines that are application servers (ie. email server), workstations, or virtual desktops. The datastore will be configured with a block size of 128 KB.
- *File storage* - for virtual machines providing file storage. The dataset will be configured with a block size of 1MB.

The screenshot shows the SynetoOS interface with a list of datastores and a detailed view for 'WindowsServer'.

Name	VMs	Used space	Protection
EmailServer-Linux	0	23.0 KB	NOT PROTECTED
OracleServer	0	23.0 KB	NOT PROTECTED
Windows10	1	7.8 GB	M, D
Windows10_2018_12_12_15_23	1	75.0 KB	NOT PROTECTED
WindowsServer	0	23.0 KB	NOT PROTECTED

WindowsServer Details:



- Virtual machines: none
- Workload: Application server (Change)
- Protection: NOT PROTECTED (Protect)
- Space effective used: 11.50 KB (Data: 23.00 KB, Snapshots: 0.00 B)
- Compression: 1.00x (Saved 0.0 B)
- Mounted on: syneto-esxi-059473bb.dev.syneto.net (Change)
- Description: [2018-12-13 17:56:19] Datastore Windows-Server created. (Edit)


NOTE: If you have recovered files from a snapshot, then a new datastore will be created - it will not appear in this list of datastores / from SynetoOS, you will only be able to see it from your ESXi.

The new datastore will be added to the list, and preselected for you. Under the lists you will see the details about it. In the details section you will be able to perform a set of actions:




- Change on which ESXi hosts the datastore is mounted on
- Change *Workload* type
- *Protect* datastore
- Browse snapshots related to this datastore
- Delete datastore
- View and edit the datastore's description. This field shows by default the time when the datastore was created.

At the top of the page you will see, next to the *Search* field, the *Show* button. Clicking on it will open up a context menu, from which you can select which kinds of items to show, like:

- *Originals* - datastores created on this machine
- *Clones* - datastores cloned from snapshots, marked with 
- *Replicas* - datastores received from other HYPER machines, marked with 




The screenshot shows a table with columns: Name, VMs, Used space, and Protection. A context menu is open over the table, showing options: Originals (checked), Clones (checked), and Replicas (checked). The table contains the following data:

Name	VMs	Used space	Protection
EmailServer-Linux	0	23.0 KB	NOT PROTECTED
OracleServer	0	23.0 KB	NOT PROTECTED
Windows10	1	7.8 GB	 
 Windows10_2018_12_12_15_23	1	75.0 KB	NOT PROTECTED
WindowsServer	0	23.0 KB	NOT PROTECTED

Provisioning new datastore when multiple pools are available

Your device comes preconfigured with either only a hybrid pool, or a flash pool, or both.




The *New datastore* form will allow you to choose on which of these to create the datastore.



The screenshot shows the 'New datastore' form with the following fields and options:

Name

Workload type

Database  Application server  File storage 

Store on

flash /19.9 GB free space hybrid /2.0 TB free space

Cancel

Each pool will have its available space specified. Usually, flash pools are very fast but small. Hybrid pools are slower but provide a large amount of space.

Here are some tips to help you choose based on your expected workload:

- Flash pool
 - High IO virtual machines
 - VDI
 - Database servers with many write operations
- Hybrid pool
 - Normal workload virtual machines
 - Email servers
 - Active directory
 - DHCP servers
 - Database servers with many read and few write operations

Deleting a datastore

There may be cases when you want to completely remove a datastore and all its data. Open the details view of the datastore and click on the *Delete* button.

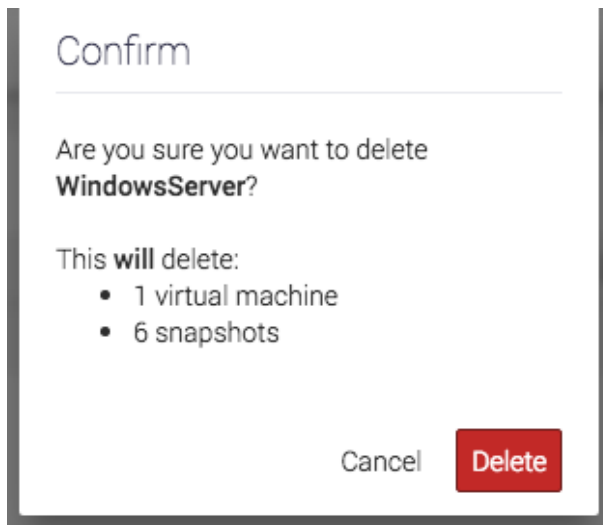
The screenshot shows the SynetoOS interface. At the top, there is a search bar and a filter dropdown set to 'Show Originals, Clones, » Replica:'. Below this is a table of datastores:

Name	VMs	Used space	Protection
EmailServer-Linux	0	23.0 KB	NOT PROTECTED
OracleServer	0	23.0 KB	NOT PROTECTED
Windows10	1	7.8 GB	M D
Windows10_2018_12_12_15_23	1	75.0 KB	NOT PROTECTED
WindowsServer	0	23.0 KB	NOT PROTECTED

The 'WindowsServer' datastore is selected, and its details view is shown below. The 'Delete' button is highlighted with a red box. The details view includes:

- Virtual machines:** none
- Workload:** Application server (Change)
- Protection:** NOT PROTECTED (Protect)
- Space effective used:** 11.50 KB (Data: 23.00 KB, Snapshots: 0.00 B)
- Compression:** 1.00x (Saved 0.0 B)
- Mounted on:** syneto-esxi-059473bb.dev.syneto.net (Change)
- Description:** [2018-12-13 17:56:19] Datastore Windows-Server created. (Edit)

After you click *Delete* a confirmation dialog will explain you that snapshots and virtual machines related to this datastore will be also deleted. You will need to confirm it in order to proceed further.



NOTE: Only datastores with no clones can be deleted

NOTE: Only datastores with no running virtual machines can be deleted

Protecting

Please refer to [Configuring the snapshot schedules for the datastore](#) for more details.

Monitoring

You can list all your datastores by going to *Main menu* → *Datastores*. The page will load a table containing various information about the datastores.

Name	VMs	Used space	Protection
EmailServer-Linux	0	23.0 KB	NOT PROTECTED
OracleServer	0	23.0 KB	NOT PROTECTED
Windows10	1	7.8 GB	M D

The list contains several columns highlighting the most important information about each datastore: number of virtual machines, used space, protection.

Virtual machines

The *VMs* column shows how many virtual machines are running on each datastore. By clicking this column title you can sort the list ascending or descending by the number of virtual machines.

Name	VMs	Used space	Protection
Windows10	1	7.8 GB	M D
Windows10_2018_12_12_15_23	1	75.0 KB	NOT PROTECTED

NOTE: We recommend one virtual machine per datastore.

Used space

This column allows you to spot the datastores, and virtual machines that are using the most space. The column represents the space allocated by the virtual machines by VMware on the datastore.

Name	VMs	Used space	Protection
IsolImages	4	23.8 GB	H

NOTE: This helps you spot virtual machines, or groups of virtual machines that are using up a lot of space.

Recovering data from a snapshot

Get access to a previous version of the datastore

It sometimes happens that you lose data. This can happen because of a virus penetrating your LAN, a user deleting a file by accident, your virtual machine having crashed and not starting anymore, or a database upgrade going terribly wrong.

Or perhaps you'd like to get access to a previous version of the data for development or testing purposes. For example, you want to run a report on the database, but don't want to accidentally delete live production data.

The screenshot shows the SynetoOS interface. At the top, there is a search bar and a dropdown menu set to 'Show Originals, Clones, » Replica:'. Below this is a table of datastores:

Name	VMs	Used space	Protection
EmailServer-Linux	0	23.0 KB	NOT PROTECTED
OracleServer	0	23.0 KB	NOT PROTECTED
Windows10	1	7.8 GB	M D
Windows10_2018_12_12_15_23	1	75.0 KB	NOT PROTECTED
WindowsServer	0	23.0 KB	NOT PROTECTED

The 'WindowsServer' datastore is selected, and its details are shown below. The 'Snapshots' tab is highlighted with a red box. The details view includes:

- Virtual machines:** none
- Workload:** Application server (Change)
- Protection:** NOT PROTECTED (Protect)
- Space effective used:** 11.50 KB (Data: 23.00 KB, Snapshots: 0.00 B)
- Compression:** 1.00x (Saved 0.0 B)
- Mounted on:** syneto-esxi-059473bb.dev.syneto.net (Change)
- Description:** [2018-12-13 17:56:19] Datastore Windows-Server created. (Edit)

For all these, and many more cases, HYPER clones come in very handy. You can recreate the data exactly as it was at the time of the snapshot, using the *Clone* function. To create a clone, start by clicking *Snapshot* tab in the details view of a datastore.

WindowsServer Delete Details Snapshots

December 2018

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

December 14, 2018 4 snapshots

▶	ⓘ	11:08	23.0 KB
▶	ⓘ	11:00	23.0 KB
▶	ⓘ	10:57	23.0 KB
▶	ⓘ	09:26	23.0 KB

Summary

Oldest snapshot on 2018/12/13

Total snapshots 7

Pick a date from the calendar, and click on the contextual menu for a snapshot.

WindowsServer Delete Details Snapshots

December 2018

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

December 14, 2018 4 snapshots

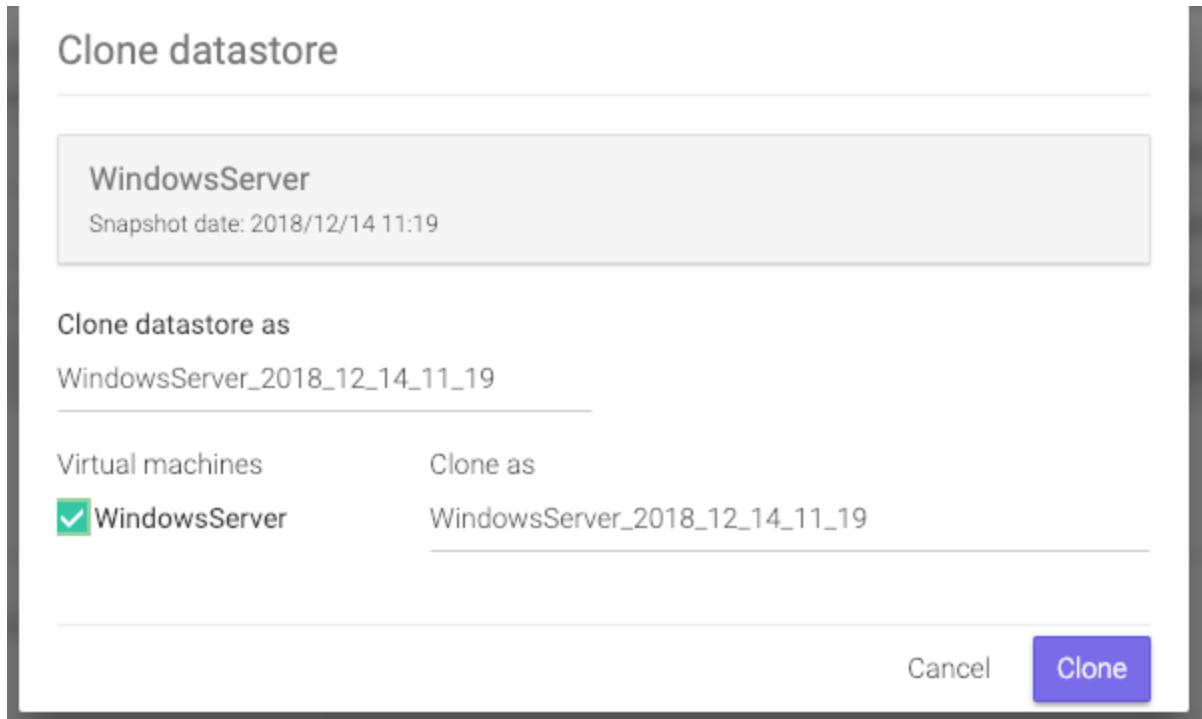
▶	ⓘ	11:08	23.0 KB
▶	ⓘ	11:00	23.0 KB
▶	ⓘ	10:57	23.0 KB
▶	ⓘ	09:26	23.0 KB

Summary

Oldest snapshot on 2018/12/13

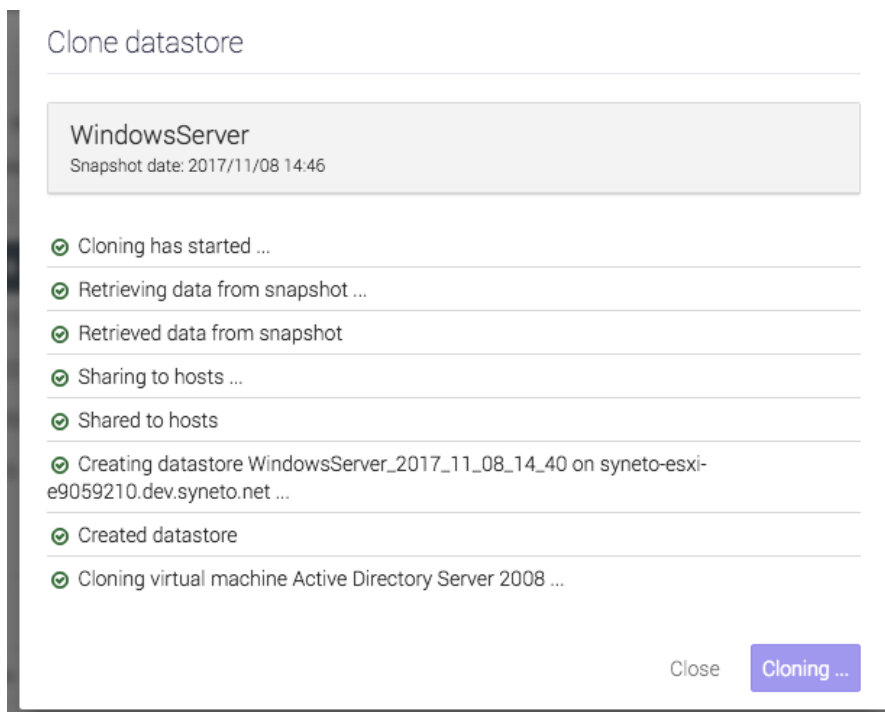
Total snapshots 7

Click *Clone* and select the virtual machines you want to recover.




If the snapshot stored more than 1 virtual machine, you can select all of them or just the virtual machines you need. Only the selected VMs will be registered with the hypervisor.

Click *Clone* and wait for the process to finish.



Connect to your vSphere configuration interface and note that a new virtual machine was created for you.



The screenshot shows the vSphere Virtual Machines interface for the host 'syneto-esxi-e9059210.dev.syneto.net'. A table lists several virtual machines, with one selected. Below the table, the details for the selected VM are shown.

Virtual machine	Status	Used space	Guest OS	Host name	Host CPU	Host mem
Active Directory Server 2008	Normal	3 KB	Microsoft Windows Serv...	Unknown	1.8 GHz	2.29 GB
Active_Directory_Server_2008_2017_11_08_14_40	Normal	5.55 GB	Microsoft Windows Serv...	Unknown	0 MHz	0 MB
Email Server	Normal	1.49 KB	CentOS 7 (64-bit)	Unknown	0 MHz	0 MB
Email Server MX	Normal	1.49 KB	CentOS 7 (64-bit)	Unknown	0 MHz	0 MB
Oracle Server	Normal	1.49 KB	Oracle Solaris 11 (64-bit)	Unknown	0 MHz	0 MB

Active_Directory_Server_2008_2017_11_08_14_40

Guest OS	Microsoft Windows Server 2008 R2 (64...
Compatibility	ESXi 6.5 and later (VM version 13)
VMware Tools	No
CPUs	1
Memory	4 GB

CPU: 0 MHz
MEMORY: 0 B
STORAGE: 5.55 GB

From here you can start the virtual machine and recover/copy the data you need.

After the cloning has finished, the cloned datastore can be identified by the icon in the table row, or the blue information box in the details view. In the information box you will find details about the clone source (original datastore): datastore name and time when the cloned snapshot was taken.

Search 1—6 of 6

Show Originals, Clones, » Replica: ▾

Name	VMs	Used space	Protection
OracleServer	0	23.0 KB	NOT PROTECTED
Windows10	1	7.8 GB	M D
Windows10_2018_12_12_15_23	1	75.0 KB	NOT PROTECTED
WindowsServer	1	224.5 KB	M H
WindowsServer_2018_12_14_11_19	1	71.0 KB	NOT PROTECTED

WindowsServer_2018_12_14_11_19 Revert Delete Details

Cloned from WindowsServer (2018/12/14 11:19)

Virtual machines 1/1 [View](#) Register additional Space effective used 653.50 KB

Workload Application server [Change](#) 71.00 KB

Data: 71.00 KB — Snapshots: 0.00 B

Compression 10.71x Saved 582.5 KB

Mounted on • syneto-esxi-059473bb.dev.syneto.net [Change](#)

Description [2018-12-14 11:40:41] Datastore WindowsServer_2018_12_14_11_19 created. [Edit](#)

It is also possible to see all the clones of a datastore by selecting the original datastore. In the details view there will be a list of all the clones made from it.

WindowsServer Delete

Virtual machines 1 [View](#) Space effective used 18.0 GB

Protection D 9.9 GB

Snapshots 3 [Browse](#) Data: 6.1 GB — Snapshots: 3.8 GB

Clones

Name	VMs	Snapshot date	Description
WindowsServer_2018_02_09_14_40	1/1	2018/02/09 14:40	[2017-07-22 11:56:03] Datastore WindowsServer created. Edit
WindowsServer_2018_02_08_16_40	1/1	2018/02/08 16:40	

Compression 1.81x Saved 8.1 GB

Mounted on • syneto-esxi-70d3bdb4.dev.syneto.net [Change](#)

Clean up when you are done

One of the more tedious tasks after recovering data is cleaning up. Finding, stopping, deregistering virtual machines, removing VMware datastores, deleting storage clones are all complex operation. And because they are time consuming and difficult, many people ignore these tasks. The data rots, takes up unused space and in the long term slows us down due to clutter.

With Syneto HYPER the whole cleanup process is just one click away. And we take care to clean up only what we created.

Go to *Main menu* → *Datastores* and select a datastore that has cloned snapshots. In its details view you can see a section named *Clones*. Open its contextual menu and select *Delete clone*.

The screenshot shows the details view for a datastore named 'WindowsServer'. At the top, there is a table listing VMs and snapshots. Below this, there are tabs for 'Delete', 'Details', and 'Snapshots'. The 'Details' tab is active, showing various metrics and settings. A contextual menu is open over the 'Clones' section, with 'Delete clone' selected.

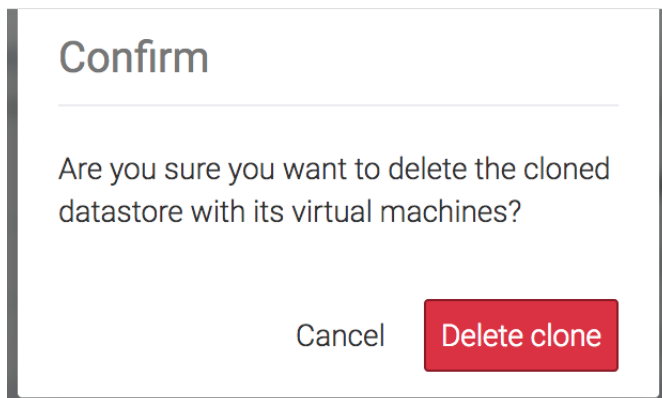
Name	VMs	Snapshot date
WindowsServer_2018_12_14_11_19	1/1	2018/12/14 11:19

Or select the cloned datastore and click *Delete*.

The screenshot shows the details view for a cloned datastore named 'WindowsServer_2018_12_14_11_19'. At the top, there is a table listing VMs and snapshots. Below this, there are tabs for 'Revert', 'Delete', and 'Details'. The 'Delete' tab is active, and the 'Delete' button is highlighted with a red box.

Name	VMs	Snapshot date
WindowsServer_2018_12_14_11_19	1/1	2018/12/14 11:19

It will ask you to confirm the process and it will explain all the things that will happen, such as virtual machines will be removed, datastore unmounted from VMware, etc.



The process takes a little time, you will see a step by step progress.

That's it, all that was created, shared, connected, imported for that snapshot clone was reverted. No more mess left behind by the recovery process.

Reverting a clone

Once you created one or more clones from your snapshots, you may find a specific clone that contains exactly the data and virtual machines you wish to recover. The *revert* functionality allows you to replace your current dataset with the clone.

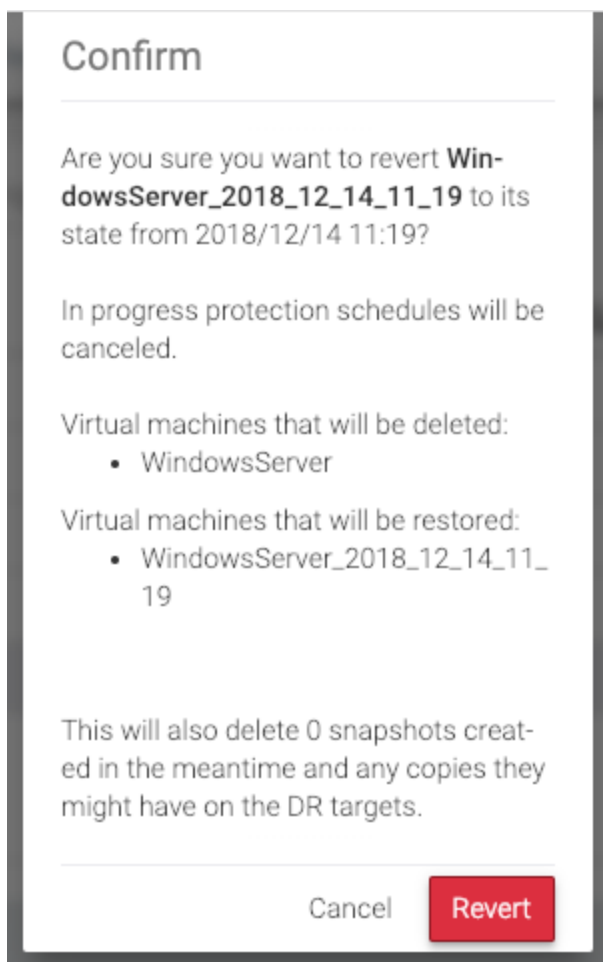
To revert a clone, select a clone, then click *Revert*.

A clone revert process will do the following steps:

- Force stop and unregister all virtual machines from the datastore
- Unmount the original datastore
- Delete all snapshots created after the snapshot from which the clone was created
- Replace the original datastore with the clone both on SynetoOS and vSphere.
- Keep virtual machines restored from the clone with their cloned names

WARNING: Once a clone is restored, all data and snapshots newer than the clone will be permanently lost.

NOTE: If you have several clones of the same datastore, only the clone from the most recent snapshot can be reverted.



During the revert process you will see a step-by-step progress. If something goes wrong, an appropriate message will be shown.

Datastore clone revert

- ✔ Started revert process ...

- ✔ Unmounting original datastore from hypervisors ...

- ✔ Removing virtual machine WindowsServer_2018_07_17_12_00 ...

- ✔ Removed virtual machine WindowsServer_2018_07_17_12_00.

- ✔ Datastore unmounted.

- ✔ Reverting clone ...

Close

WARNING: On DR unit, on a received datastore, you can create clones but you cannot revert the clone.

After the revert process finishes, the system will update the datastore's *Description* field with the relevant information.

Description	[2018-02-07 18:26:21] VMware Datastore reverted from clone WindowsServer_2018_02_07_18_25	Edit
	[2018-02-07 18:25:57] Datastore WindowsServer_2018_02_07_18_25 created.	

Malware / Ransomware use case

Clone revert is a very powerful tool to protect you from malware and ransomware. Let's say you have a datastore, *DS1*, with a snapshot schedule every 15 minutes.

The schedule was running for the past 2 hours, and with a retention policy of 6 snapshots, you have 6 snapshots for *DS1*.

Suddenly, users are starting calling you that something is wrong with their virtual machines on *DS1*. They doesn't seem to have access to their machines, or they were infected with a ransomware.

You, as a sysadmin, can start creating clones of the snapshots you have on *DS1*. You will create *DS1_clone1* from the newest snapshots. Let's see if things were alright 15 minutes ago. You realize they were already broken.

You go on with the second, third, and fourth snapshots. Finally, the clone created from the fourth snapshot contains healthy virtual machines.

You decide to revert the data to *DS1_clone4*. You know you will not need any newer data as all is consumed by malware.

But you cannot click revert just yet. You will have to first delete all newer clones. When done deleting, you can go on with the revert process.

The datastore will be reverted. All virtual machines registered into vSphere will be preserved. If you need to register additional VMs, you will need to use vSphere management interface.

Registering additional virtual machines from a clone

A datastore snapshot can contain multiple virtual machines. In the recovery process a datastore clone is made and you are presented with the option to register the virtual machines from the snapshot. If in this process you did not select all the virtual machines you don't need to worry, after the cloning has finished you can register additional virtual machines.

Select the datastore from which the clone was made. In the details view identify the clone and open the contextual menu and select *Register additional VMs*.

Alternatively you can select the clone and click *Register additional* in the *Virtual machines* sections.

The screenshot displays the vSphere management interface for a datastore clone. At the top, a table lists clones with columns for name, count, size, and protection status. Below this, the details for the selected clone 'Windows10_2018_12_12_15_23' are shown, including options to 'Revert', 'Delete', and 'Details'. A section indicates the clone was created from 'Windows10 (2018/12/12 15:23)'. The 'Virtual machines' section shows 0/1 VMs with a 'Register additional' link. The 'Workload' section shows 'Application server' with a 'Change' link. A progress bar indicates 'Space effective used' with a total of 404.50 KB and 86.00 KB of data. Other details include 'Compression 5.86x', 'Saved 318.5 KB', 'Mounted on' (syneto-esxi-059473bb.dev.syneto.net), and a 'Description' of the clone creation.

Clone Name	Count	Size	Protection
Windows10_2018_12_12_15_23	0	86.0 KB	NOT PROTECTED
WindowsServer	0	315.5 KB	M H

Windows10_2018_12_12_15_23 [Revert](#) [Delete](#) [Details](#)

Cloned from *Windows10 (2018/12/12 15:23)*

Virtual machines: 0/1 [Register additional](#) Space effective used: 404.50 KB (86.00 KB Data, 0.00 B Snapshots)

Workload: Application server [Change](#) Compression 5.86x Saved 318.5 KB

Mounted on: syneto-esxi-059473bb.dev.syneto.net [Change](#)

Description: [2018-12-13 18:01:56] Datastore Windows10_2018_12_12_15_23 created. [Edit](#)

A dialog box will appear and provide a list of all the virtual machines you can register.

Register virtual machines

WindowsServers_2018_07_17_16_18
Cloned from: WindowsServers
Snapshot date: 2018/07/17 16:18

Virtual machines	Clone as
<input checked="" type="checkbox"/> Microsoft SQL Server	Microsoft_SQL_Server_2018_07_17_16_18
<input type="checkbox"/> Microsoft Active Directory	

Cancel **Register**

Select the desired ones, and click *Register*. A nice progress will show you the steps taken by the system to register the virtual machine to vSphere.

Register virtual machines

WindowsServers_2018_07_17_16_18
Cloned from: WindowsServers
Snapshot date: 2018/07/17 16:18

- Cloning has started ...
- Cloning virtual machine Microsoft SQL Server ...

Cancel **Still working ...**

After the process finished you can see the virtual machines that were registered by hovering over *View* in the clone's details view.

The screenshot displays a table of virtual machines with the following data:

Name	VMs	Used space	Protection
EmailServer-Linux	0	23.0 KB	NOT PROTECTED
OracleServer	0	23.0 KB	NOT PROTECTED
Windows10	1	7.8 GB	M, D
WindowsServer	2	420.0 KB	M, H
WindowsServer_2018_12_14_14_40	1	132.0 KB	NOT PROTECTED

The details for 'WindowsServer_2018_12_14_14_40' are shown below:

- Cloned from:** WindowsServer (tooltip: Microsoft_SQL_Server_2018_12_14_14_40)
- Virtual machines:** 1/2 (View, Register additional)
- Workload:** Application server (Change)
- Space effective used:** 1.27 MB (132.00 KB) (Data: 132.00 KB, Snapshots: 0.00 B)
- Compression:** 11.12x (Saved 1.1 MB)
- Mounted on:** syneto-esxi-059473bb.dev.syneto.net (Change)
- Description:** [2018-12-14 14:42:09] Datastore Windows-Server_2018_12_14_14_40 created. (Edit)

Serving non-Syneto hosts

In the bottom right of a newly created datastore's details section you can see that the newly created datastore is automatically mounted on the local hypervisor (ESXi host).

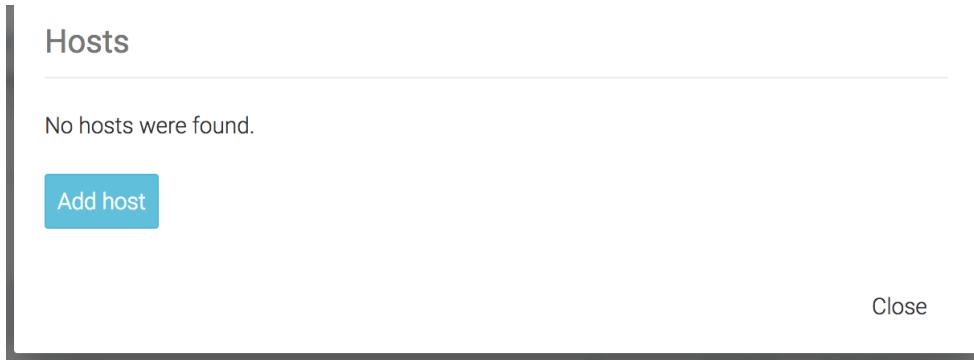
If you click *Change* you will be able to unmount it from this local hypervisor.

But there are infrastructures that already have ESXi hosts, independent or managed by a vCenter. For these architectures, you can mount a datastore on any ESX(i) host.

WARNING: This feature is not available on Syneto HYPER 2000 Series.

Defining external hosts


Go to *Main menu* → *Datastores*. On the action bar you will see a button named *External hosts*.



By default the dialog will show no host.


NOTE: Local ESXi hypervisor is not displayed on the External hosts list.

Clicking the *Add host* button will show a form where you can add another host.



The *Host* field specifies an IP address or a DNS resolvable FQDN (host name).
The *Username* field must be filled with user having full VMware administrative privileges.


WARNING: If your ESXi is managed by a vCenter, you must specify the vCenter at the *Host* field. Adding the ESXi host may lead to partial or unexpected interaction between the Syneto HYPER and VMware.



The 'Add host' dialog box contains the following fields and controls:

- Add host** (Title)
- Host**: 192.168.5.25
- Username**: root
- Password**: masked with seven dots
- Cancel** button
- Add** button

After clicking *Add* the host will appear in the list



The 'Hosts' list displays the following information:

- Hosts** (Title)
- Host entry: 192.168.5.29 (with a menu icon to its left)
- Add host** button
- Close** button

Mounting datastore to additional hosts

Go to *Main menu* → *Datastores* and click a datastore. In the details section click on *Change* (bottom right).

Search 1—5 of 5

Show Originals, Clones, » Replica: ▼

Name	VMs	Used space	Protection
EmailServer-Linux	0	23.0 KB	NOT PROTECTED
OracleServer	0	23.0 KB	NOT PROTECTED
Windows10	1	7.8 GB	
WindowsServer	2	420.0 KB	
WindowsServer_2018_12_14_14_40	1	132.0 KB	NOT PROTECTED

WindowsServer Delete Details Snapshots

Virtual machines 2 [View](#) Space effective used 842.50 KB
420.00 KB

Workload Application server [Change](#) Data: 105.00 KB — Snapshots: 315.00 KB

Protection [Protect](#) Compression 3.51x Saved 422.5 KB

Clones Mounted on • syneto-esxi-059473bb.dev.syneto.net [Change](#)

Name	VMs	Snapshot date	Description
WindowsServer_2018_12_14_14_40	1/2	2018/12/14 14:40	[2018-12-14 14:16:41] VMware Datastore reverted from clone Windows-Server_2018_12_14_14_12 [2018-12-14 14:15:24] Datastore Windows-Server_2018_12_14_14_12 created. Edit

If you added a standalone ESXi to *Hosts*, it will be on the list. You will be able to mount the datastore on this host by checking it and clicking *Update*.

If you added a vCenter server to *Hosts*, you will find in this list all the ESXi hosts managed by the specified vCenter server. You will be able to mount the datastore on these hosts by checking then and clicking *Update*.

Mounted on hosts

syneto-esxi-e9059210.dev.syneto.net

syneto-esxi-bfbd346e.dev.syneto.net


Close [Update](#)

NOTE: All ESXi hosts are listed by their name. In case you didn't set up a hostname for your hosts and you add several of them, you may see "localhost" listed here several times. We recommend you set up a unique hostname on each of your ESXi hosts.

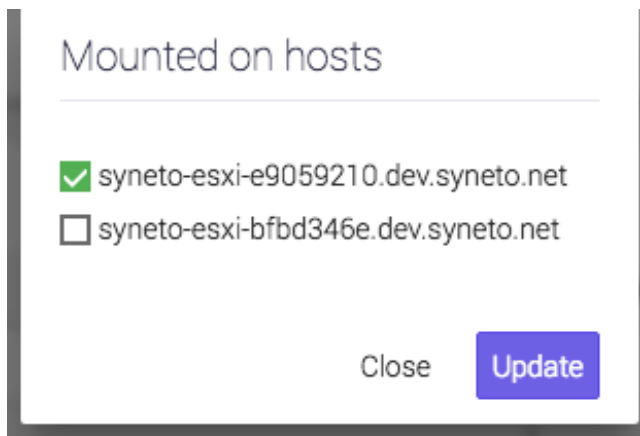
NOTE: In some cases VMware reports the host name as an IP address. In these cases you will see the IP address of the host, regardless of the way you specified it in the *Hosts* dialog.

Unmounting datastore from hosts

Using the same workflow as for mounting a datastore to additional hosts, you can unmount them as well.



Simply uncheck the hosts you want the datastore to be unmounted from.



Click *Update*.

VIRTUAL MACHINES

Introduction

Virtual machines are virtual computers that run on a hypervisor. Syneto HYPER comes pre-configured with a VMware ESXi local hypervisor. You can provision, manage, run, protect, analyze virtual machines on the device itself.

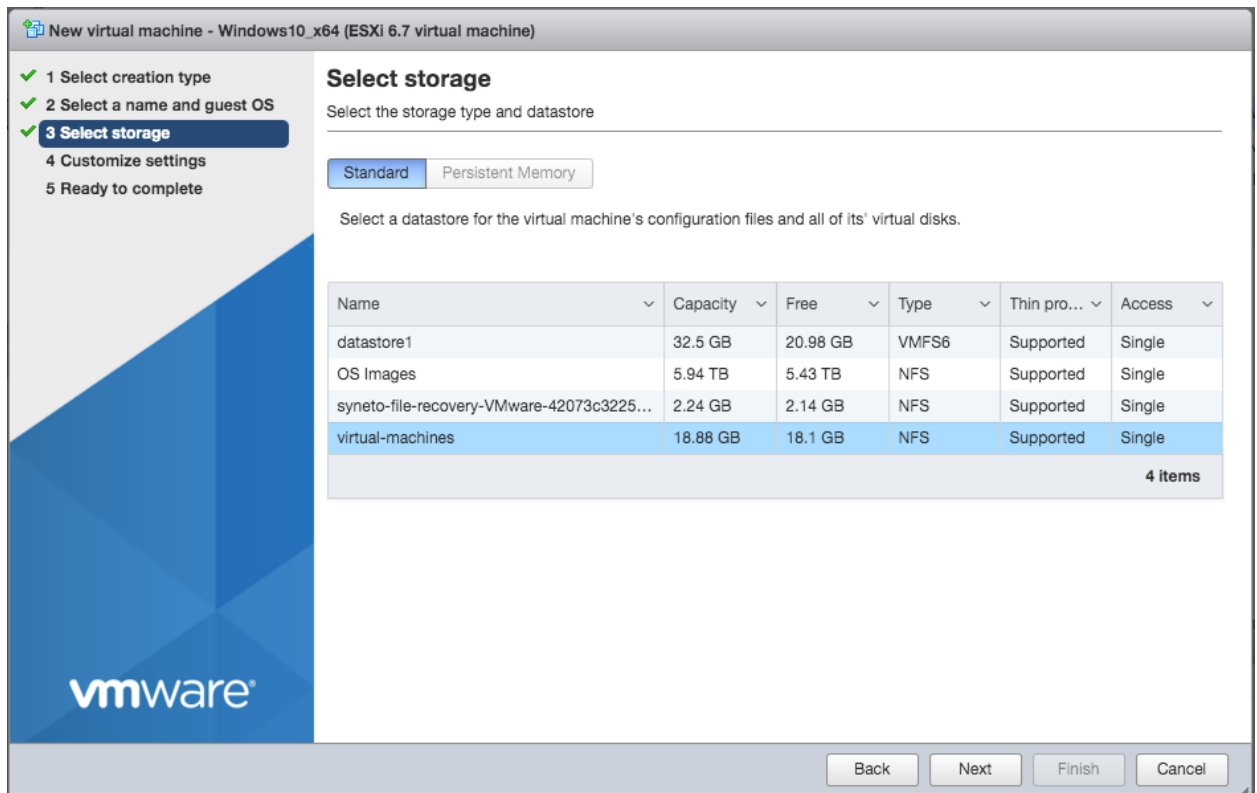
Provisioning and managing

Adding virtual machines to your infrastructure is done using vSphere management interface.

NOTE: You can quickly connect to your local ESXi hypervisor web management interface by clicking *Login to UI* on the *Hypervisor* widget on the dashboard. See [THE DASHBOARD](#) for more details.

After completing the previous chapter you will see one or several new datastores provided by Syneto HYPER.

All you have to do is provision a new virtual machine on vSphere management interface. Specify that it resides on a datastore provided by Syneto HYPER.



NOTE: For granular protection and performance analyses we recommend creating **one virtual machine per datastore**.

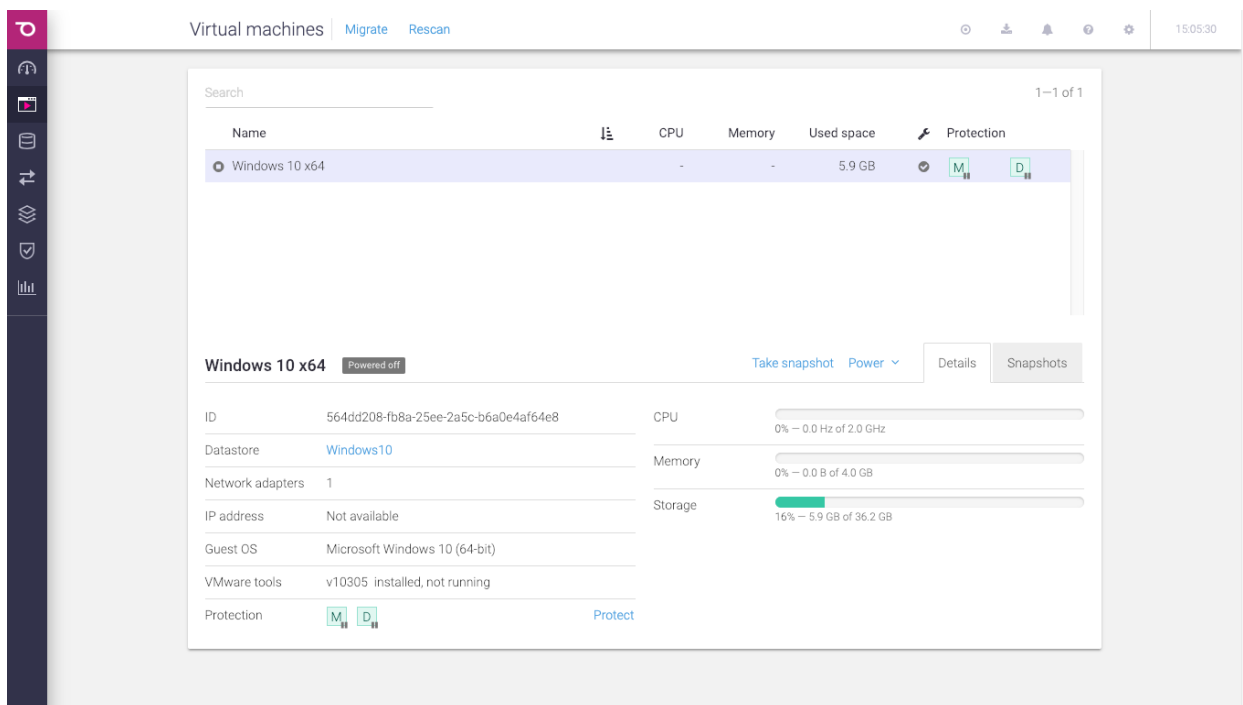
WARNING: While it is technically possible to attach several datastores to a single virtual machine, Syneto HYPER doesn't support this configuration. Data protection will not work properly on virtual machines connected to several datastores.

Power Management


There are several management actions that can be performed on a virtual machine:

- Power on
- Power off
- Reset
- Suspend
- Resume

Let's start with a virtual machine that is powered off. After clicking it in the VMs list, its details pane opens. In the top right corner of the details pane, click *Power On*. This kickstarts the boot sequence of the operating system.

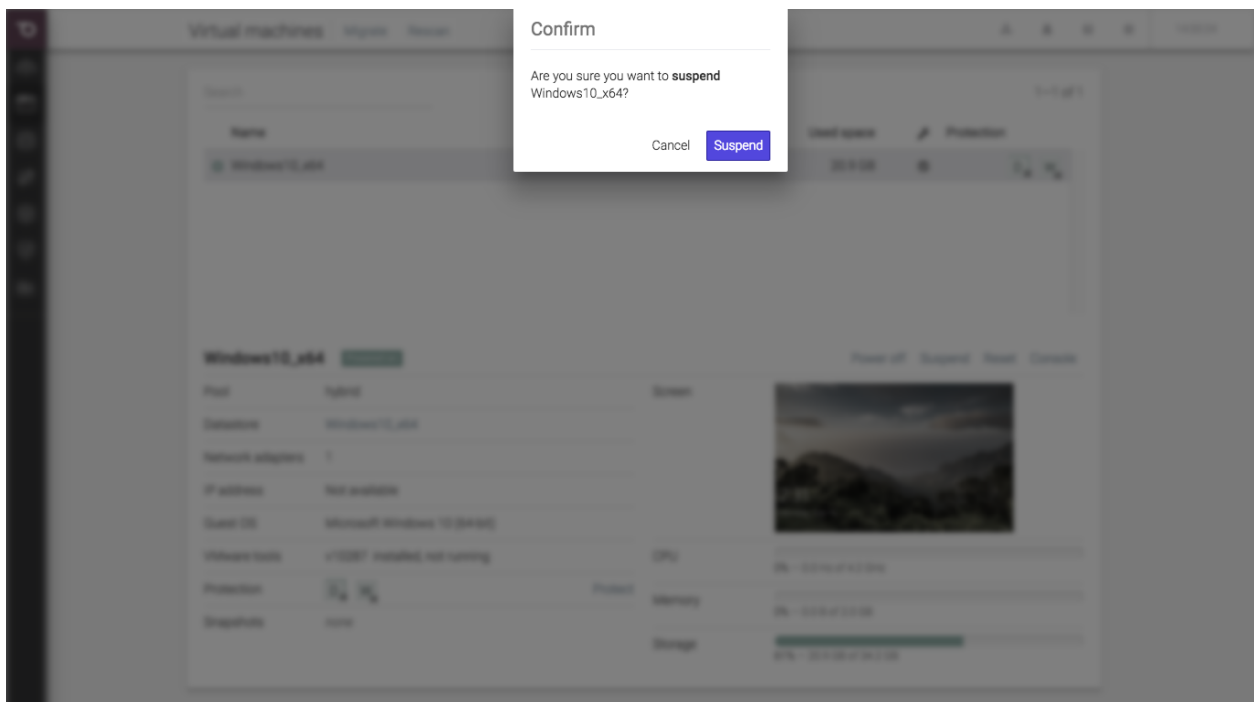


After a VM was powered on, new actions become available: *Power off*, *Suspend* and *Restart*. Also, the virtual machine's console is displayed on screen.



WARNING: *Power off* and *Reset* will perform hard operations, it will behave exactly as if you switched off a physical machine by turning off the power.

If you want to “pause” a virtual machine, click on *Suspend*. This action is similar to closing the lid on a laptop: it dumps the VM's memory to drive and then deallocates the CPU and memory resources it used.



A confirmation dialog will open. Click on *Suspend* if you want to continue.


The transition message *Suspending...* will be displayed while the memory is being written to drive, and the other power buttons will disappear. *Power off* is still an option if the virtual machine is suspended or is being suspended.



You can *Resume* or *Power off* the suspended virtual machine. Resuming a suspended virtual machine will load it from drive in the state it was before the suspend. No boot sequence is performed.

Remote Console

If a virtual machine is powered on, you can see its console in the details section.



The screenshot shows the SynetoOS interface for managing virtual machines. The top navigation bar includes "Virtual machines", "Migrate", and "Rescan". A search bar is present, and a table lists the VMs. The selected VM, "Windows 10 x64", is shown in the details section. The VM is powered on and has the following specifications:

Name	CPU	Memory	Used space	Protection
Windows 10 x64	1.5 GHz	1.6 GB	5.9 GB	M D

The details section for "Windows 10 x64" (Powered on) includes the following information:


- ID: 564dd208-fb8a-25ee-2a5c-b6a0e4af64e8
- Datastore: Windows10
- Network adapters: 1
- IP address: 192.168.5.154
- Guest OS: Microsoft Windows 10 (64-bit)
- VMware tools: v10305 running, managed by VMware
- Protection: M D Protect

Performance metrics are displayed as progress bars:

- CPU: 76% - 1.5 GHz of 2.0 GHz
- Memory: 41% - 1.6 GB of 4.0 GB
- Storage: 18% - 5.9 GB of 32.0 GB

A "Screen" preview shows the Windows 10 desktop environment. Action buttons include "Take snapshot", "Console", "Power", "Details", and "Snapshots".

If you want to interact with the virtual machine, click on the image of the console. This will open a new tab where you are able to do actions on your virtual machine.



The screenshot shows the remote console view for the "Windows 10 x64" VM. The interface includes the SynetoOS logo and "HYPERSeries" branding. The console is connected to the VM, and the status bar shows "Connected (encrypted) to Windows 10 x64" and "Send Ctrl+Alt+Del". The main display area shows the Windows 10 desktop environment, including the Recycle Bin, File Explorer, and the Start menu. The File Explorer window is open, showing the "This PC" view with folders (Desktop, Downloads, Pictures, Music, Videos, Documents) and devices/drives (Local Disk (C:) with 20.6 GB free of 31.5 GB, and DVD Drive (D:)).

Protecting

Please refer to [Configuring the type of snapshots for the virtual machines](#) for more details.

Consistency levels

Each virtual machine can be protected by three types of snapshot consistency levels:

1. *Crash consistent* - When restoring, the virtual machine will be in a state similar after a power outage.
Recommended for virtual machines resilient to forced reboots. (ie. Linux or Solaris servers or workstations)
2. *Application consistent* - Before taking the snapshot, a message will be sent to the operating system. This will instruct compatible applications to flush all their data to the drive. The snapshot will be taken after the flush operation finishes. Drive consistency for the application supporting the protocol will be consistent. When restoring, the virtual machine will be powered off.
 - a. Windows VSS - Volume Shadow Copy is recommended for Microsoft servers (Active Directory, MS-SQL, etc)
 - b. Hooks only - Allows running a custom pre-freeze and post-thaw script on the virtual machine. This works on any operating system with VMware tools enabled.
 - c. Windows VSS and hooks - Uses both options simultaneously
3. *Live snapshot* - Before taking a snapshot all operations on the virtual machine will be suspended and saved to the drive. This includes waiting for disk flush operations, persisting RAM, persisting current running state. Unfortunately, during this operation the virtual machine is not accessible. The process can take up to 5 minutes, or even longer for large servers.
Recommended for mission critical system and only on a daily or weekly schedule, run overnight, so that work is not interrupted.

WARNING: Live snapshot is not available on Syneto HYPER 2000 Series.

WARNING: While it is technically possible to place a VM's virtual disks on different datastores, HYPER doesn't support this configuration. Data recovery will not work properly when a virtual machine is stored on multiple datastores.

Monitoring

The main virtual machines page can be accessed by *Main menu* → *Virtual machines*. This page offers a list of virtual machines. The list contains virtual machines from all connected ESXi hosts. You will see that only the virtual machines that are on datastores served by the Syneto HYPER will be included in the list.

Name	CPU	Memory	Used space	Protection
Test	-	-	21.6 GB	✓
Windows10_x64	366.0 MHz	1.4 GB	20.9 GB	✓

The list has several columns with information about the virtual machine: CPU usage, memory usage, used space, VMware tools status, and protection status.

CPU

This column shows how much processing power is used right now by the virtual machine. The value is presented in GHz.

For detailed information click the virtual machine and in the details section on the bottom you will see current CPU usage in GHz, maximum allocated to the virtual machine in GHz, and the percentage of the current value.

Name	CPU	Memory	Used space	Protection
Windows10_x64	366.0 MHz	1.4 GB	20.9 GB	✓
Test	-	-	21.6 GB	✓

NOTE: it is useful to sort the list descending by *CPU* when you want to spot virtual machines using too much CPU. For example when you observe a virtual machine is very slow, it is a good practice to come here and quickly check its CPU usage.

Memory

This column show how much memory is currently used by the virtual machines. The value is presented in GB (or other multiples of bytes).

For detailed information click the virtual machine and in the details section on the bottom you will see current Memory usage in GB, maximum memory that the virtual machine can use, and a percentage of the the current value.

Name	CPU	Memory	Used space	Protection
Windows10_x64	366.0 MHz	1.4 GB	20.9 GB	✓
Test	-	-	21.6 GB	✓

NOTE: it is useful to sort the list descending by *Memory* when you want to spot virtual machines using too much memory. For example when you observe a virtual machine is very slow, it is a good practice to come here and quickly check its memory usage.

NOTE: virtual machines with all their memory reserved will always show 100% memory usage.

Used space

This column show how much drive space is currently used by the virtual machines. The value is presented in GB (or other multiples of bytes).

For detailed information click the virtual machine and in the details section on the bottom you will see current drive usage in GB, maximum drive space available for the virtual machine, and a percentage of the current value.

Name	CPU	Memory	Used space	Protection
Test	-	-	21.6 GB	✓
Windows10_x64	366.0 MHz	1.4 GB	20.9 GB	✓

NOTE: it is useful to sort the list descending by *Used space* when you want to spot virtual machines using too much disk space. For example when you observe a virtual machine cannot write new data or it is very slow with disk operation you may want to check this column.

Additional information

The detailed information box for each virtual machine show essential data about like operating system, datastore, disk pool, network adapter, ip address, protection type and snapshots with quick access to clone and restore. The current version of SynetoOS does not allow the modification of these properties. To change them, please login to the ESXi web console.

From the Details information box, you can not only see the Datastore the VM is mounted on, but you can also access it by clicking on the hyperlinked Datastore Name.

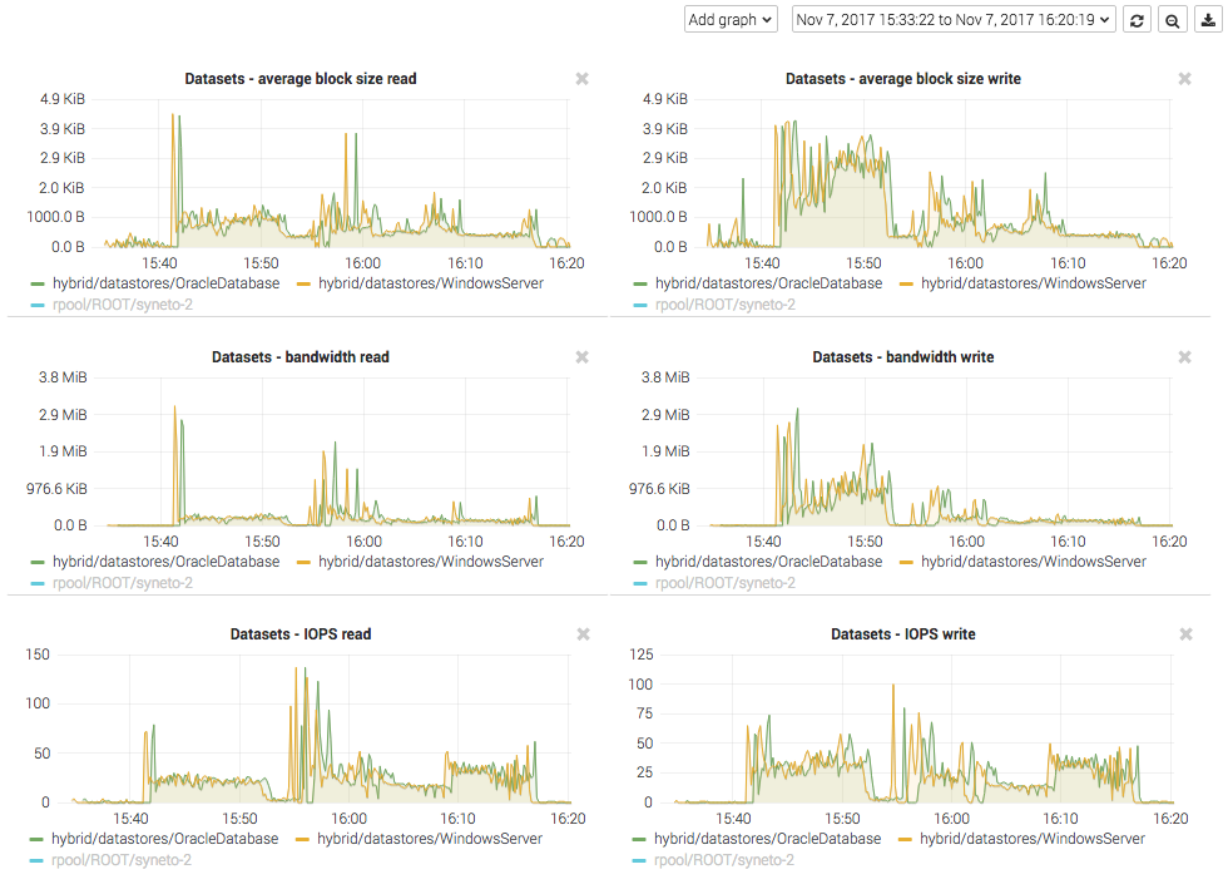
Windows 10 x64 Powered on Take snapshot Console Power Details Snapshots

ID	564dd208-fb8a-25ee-2a5c-b6a0e4af64e8	Screen	
Datastore	Windows10	CPU	46% – 922.0 MHz of 2.0 GHz
Network adapters	1	Memory	92% – 3.7 GB of 4.0 GB
IP address	192.168.5.154	Storage	18% – 5.9 GB of 32.0 GB
Guest OS	Microsoft Windows 10 (64-bit)		
VMware tools	v10305 running, managed by VMware		
Protection	Protect		

Analyzing performance

With one virtual machine per datastore you can retrieve some interesting data to analyze.

Dataset IOPS




Go to *Main menu* → *Analytics*, click *Add graph*, select *Dataset* → *IOPS*.

These graphics show the IOPS happening on the datastore from the Syneto HYPER's perspective.

NOTE: If you create one virtual machine per datastore, this is the IOPS produced by the virtual machine.

Dataset bandwidth

Go to *Main menu* → *Analytics*, click *Add graph*, select *Dataset* → *Bandwidth*.

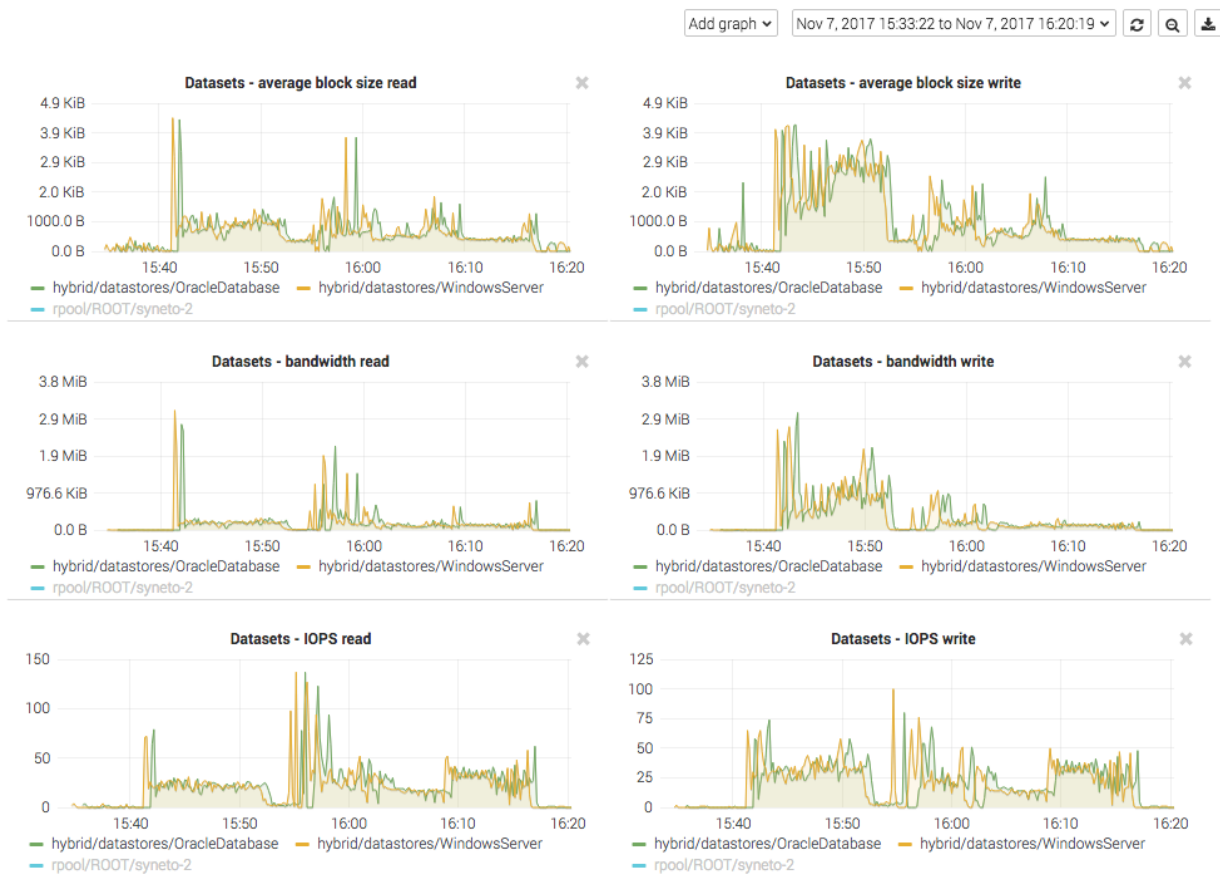


These graphics show the data transfer speed on the datastore from the Syneto HYPER's perspective.

NOTE: If you create one virtual machine per datastore, this is the bandwidth produced by the virtual machine.

Dataset average block size

Go to *Main menu* → *Analytics*, click *Add graph*, select *Dataset* → *Average block size*.



These graphics show the average block size of the data on the datastore from the Syneto HYPERS's perspective.

NOTE: If you create one virtual machine per datastore, this is the average block size of the data produced by the virtual machine.

Recovering data from a snapshot

Get access to a previous version of the virtual machine

Sometimes, it happens that you lose some data. Maybe a virus corrupted something, maybe a user deleted a file. Or simply your virtual machine crashed and it won't start any more.

The screenshot displays the SynetoOS interface for managing virtual machines. At the top, there is a search bar and a '1-1 of 1' indicator. Below this is a table listing VMs with columns for Name, CPU, Memory, Used space, and Protection. The 'Windows 10 x64' VM is selected and highlighted in blue. Below the table, the VM's name 'Windows 10 x64' is shown with a 'Powered off' status. To the right of the name are buttons for 'Take snapshot', 'Power' (with a dropdown arrow), 'Details', and 'Snapshots'. The 'Details' tab is active, showing a list of system properties and their values. On the right side of the details view, there are progress bars for CPU, Memory, and Storage usage.

Name	CPU	Memory	Used space	Protection
Windows 10 x64	-	-	5.9 GB	✓ M D

Windows 10 x64 Powered off Take snapshot Power Details Snapshots

ID	564dd208-fb8a-25ee-2a5c-b6a0e4af64e8	CPU	0% – 0.0 Hz of 2.0 GHz
Datastore	Windows10	Memory	0% – 0.0 B of 4.0 GB
Network adapters	1	Storage	16% – 5.9 GB of 36.2 GB
IP address	Not available		
Guest OS	Microsoft Windows 10 (64-bit)		
VMware tools	v10305 installed, not running		
Protection	M D		Protect

For all these, and many more cases, you can simply click on the Snapshots tab in the details view of a virtual machine.

Windows 10 x64 Powered off Take snapshot Power Details Snapshots

Automatic Manual

December 12, 2018 3 snapshots

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

Summary

Oldest snapshot on 2018/12/12

Total snapshots 3

▶	⋮	15:23	C	5.2 MB
▶	⋮	15:23	C	4.8 MB
▶	⋮	14:33	C	798.6 MB

Pick a date from the calendar, and click on the contextual menu for a snapshot.

Windows 10 x64 Powered on Take snapshot Console Power Details Snapshots

Automatic Manual

December 12, 2018 3 snapshots

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

Summary

Oldest snapshot on 2018/12/12

Total snapshots 3

▶	⋮	15:23	C	5.2 MB
▶	⋮	15:23	C	4.8 MB
▶	⋮	14:33	C	798.6 MB

Click clone. The virtual machines will be preselected for you.

Clone virtual machine

Windows 10 x64
Snapshot date: 2018/12/12 15:23

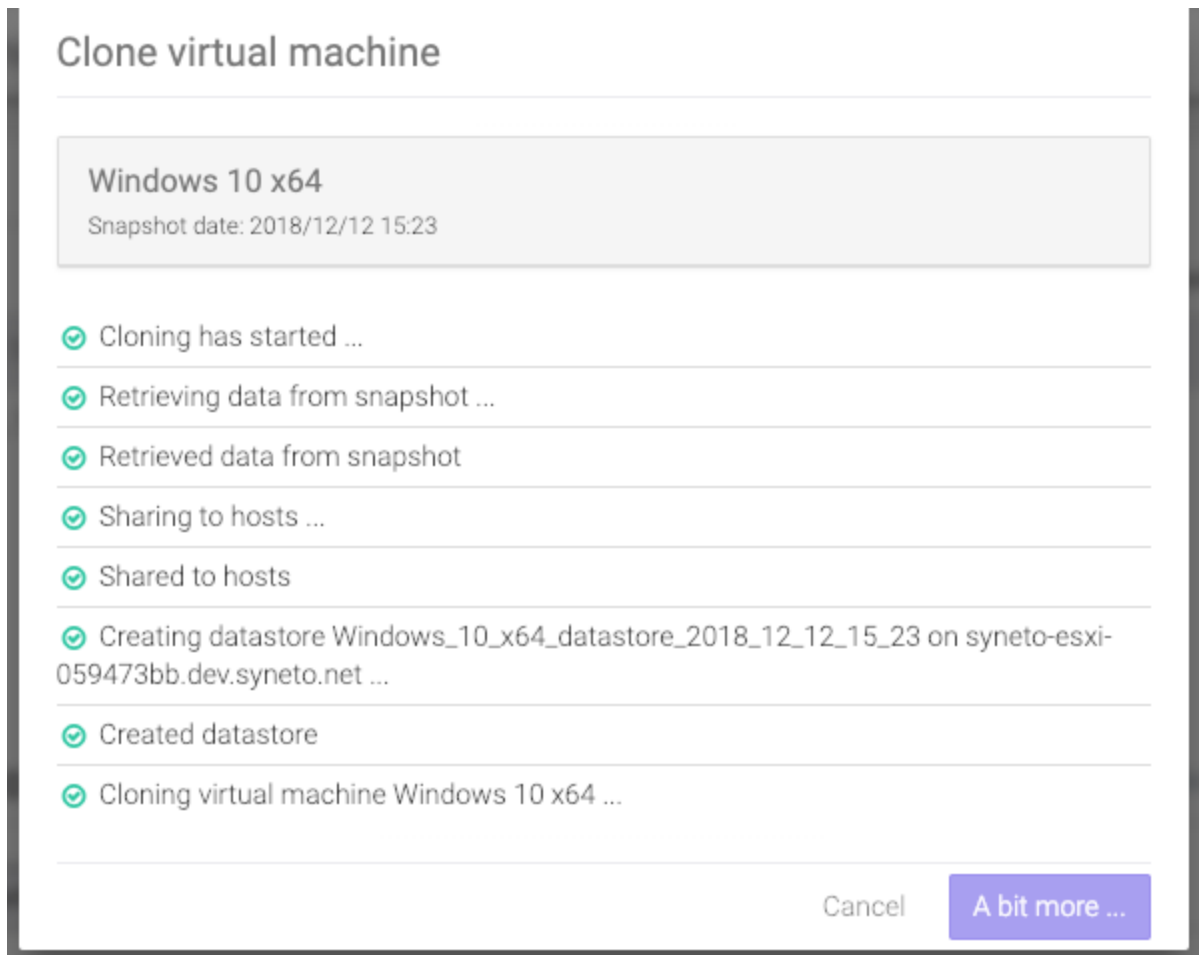
Clone virtual machine as
Windows_10_x64_2018_12_12_15_23

Access on host
syneto-esxi-059473bb.dev.syneto.net

Clone datastore as
Windows_10_x64_datastore_2018_12_12_15_23

Cancel **Clone**

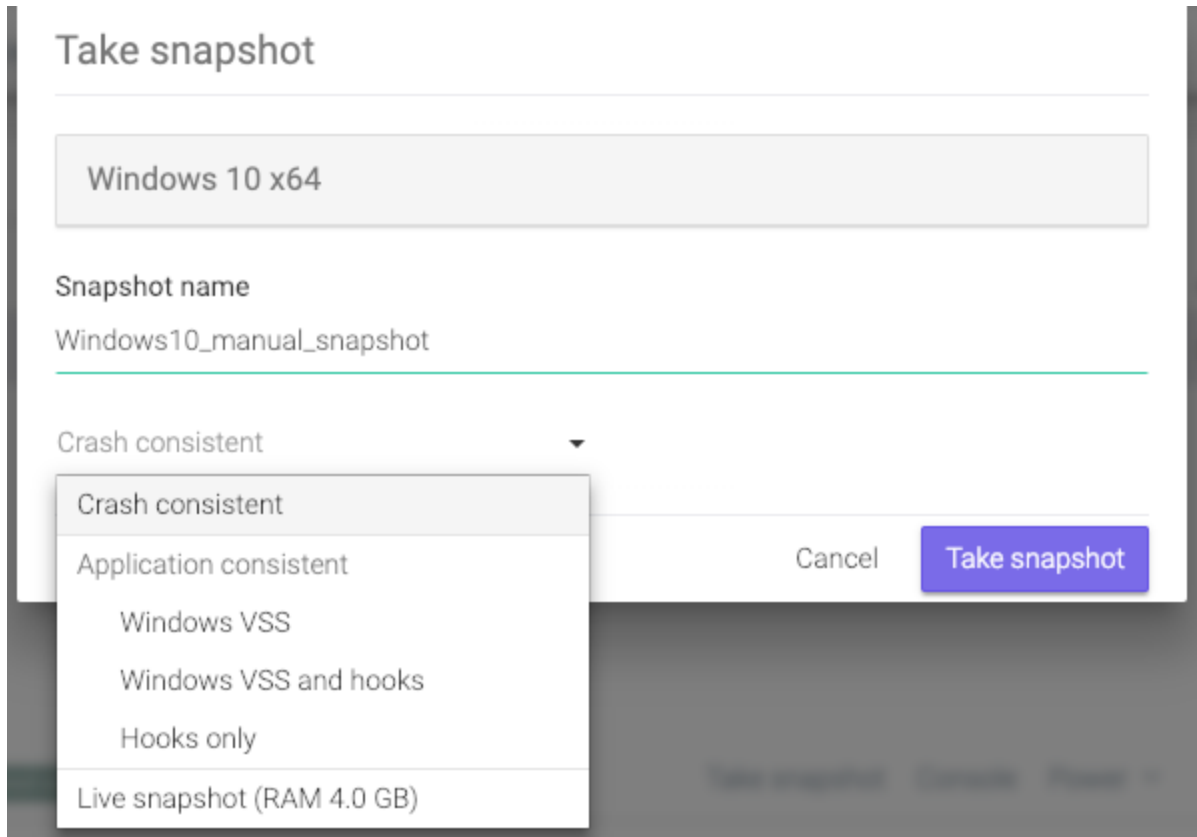
Click clone and wait for the process to finish.




On the *Virtual machines* page we note that a new virtual machine was created. Using the ESXi web interface you can start the virtual machine and recover/copy the data you need.

Take a manual snapshot

For manual snapshots click on *Take snapshot* Insert a unique snapshot name and select consistency type.



To view manual snapshots click on the Snapshots tab in the details view of a virtual machine → *Manual*.



The screenshot displays the SynetoOS Virtual Machines management interface. At the top, there are control buttons for 'Create / Register VM', 'Console', 'Power on', 'Power off', 'Suspend', 'Refresh', and 'Actions'. A search bar is located on the right. Below these is a table listing virtual machines with columns for 'Virtual machine', 'Status', 'Used space', 'Guest OS', 'Host name', 'Host CPU', and 'Host memory'. The table contains six entries, with 'Active_Directory_Server_2008_2017_11_09_11_12' selected. Below the table, a 'Quick filters...' dropdown is visible. The main area shows details for the selected VM, including a thumbnail, 'Guest OS' (Microsoft Windows Server 2008 R2 (64-bit)), 'Compatibility' (ESXi 6.5 and later (VM version 13)), 'VMware Tools' (No), 'CPUs' (1), and 'Memory' (1 GB). On the right, resource usage is shown: CPU (0 MHz), MEMORY (0 B), and STORAGE (5.82 GB). At the bottom, a 'Recent tasks' table lists actions like 'Power On VM', 'Register VM', and 'Create Nas Datastore' with their respective targets, initiators, and completion statuses.

Virtual machine	Status	Used space	Guest OS	Host name	Host CPU	Host memory
Active Directory Server 2008	✓ Nor...	8.45 GB	Microsoft Windows Ser...	Unknown	38 MHz	1.03 GB
Active_Directory_Server_2008_2017_11_09_11_12	✓ Nor...	5.82 GB	Microsoft Windows Ser...	Unknown	0 MHz	0 MB
Email Server	✓ Nor...	4.47 KB	CentOS 7 (64-bit)	Unknown	18 MHz	538 MB
Email Server MX	✓ Nor...	4.44 KB	CentOS 7 (64-bit)	Unknown	7 MHz	24 MB
Oracle Server	✓ Nor...	2.93 KB	Oracle Solaris 11 (64-bit)	Unknown	0 MHz	0 MB
SynetoOS	✓ Nor...	43.59 GB	Oracle Solaris 11 (64-bit)	syneto-os-e9059210	1 GHz	3.54 GB

Task	Target	Initiator	Queued	Started	Result	Completed
Power On VM	Active_Directory_Server_2...	root	11/09/2017 11:17:16	11/09/2017 11:17:16	Failed - Transport (VMDB) error -45: Fa...	11/09/2017 11:17:16
Register VM	vm	root	11/09/2017 11:17:04	11/09/2017 11:17:04	Completed successfully	11/09/2017 11:17:15
Create Nas Datastore	syneto-esxi-e9059210.dev.syne...	root	11/09/2017 11:17:04	11/09/2017 11:17:04	Completed successfully	11/09/2017 11:17:04

Clean up when you are done

One of the more tedious tasks after recovering data is cleaning up after you. Finding, stopping, deregistering virtual machines, removing VMware datastores, deleting storage clones can be a complex operation. And because it is time consuming and difficult, many people just let it there... forever.

With Syneto HYPER the whole cleanup process is just one click away. And we take care to clean up only what we created.

Go to *Main menu* → *Virtual machines* and select a virtual machine that has cloned snapshots. Click *Browse* in the snapshot section.

The screenshot displays the SynetoOS interface for a virtual machine named "Windows 10 x64". At the top, there is a search bar and a "1-1 of 1" indicator. Below this is a table with columns for Name, CPU, Memory, Used space, and Protection. The "Windows 10 x64" VM is listed with a CPU of "-", Memory of "-", and Used space of "5.9 GB". The Protection column shows a checkmark and two icons labeled "M" and "D".

Below the table, the VM's status is "Powered off". There are buttons for "Take snapshot", "Power" (with a dropdown arrow), "Details", and "Snapshots".

The details section is divided into two columns. The left column lists various attributes:

ID	564dd208-fb8a-25ee-2a5c-b6a0e4af64e8
Datastore	Windows10
Network adapters	1
IP address	Not available
Guest OS	Microsoft Windows 10 (64-bit)
VMware tools	v10305 installed, not running
Protection	Protect

The right column shows resource usage with progress bars:

CPU	0% – 0.0 Hz of 2.0 GHz
Memory	0% – 0.0 B of 4.0 GB
Storage	16% – 5.9 GB of 36.2 GB

You will see that all cloned snapshots are marked with a label. Click the context menu for a cloned snapshot and select *Clean up clone*.

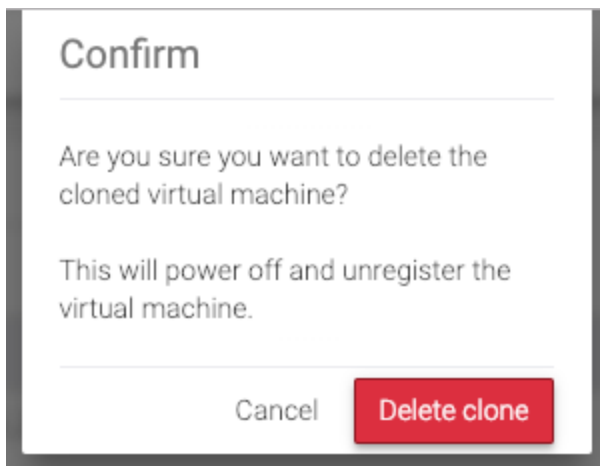
The screenshot shows the SynetoOS interface for a virtual machine named "Windows 10 x64" which is "Powered on". The "Snapshots" tab is active, displaying a calendar for December 2018 with the 12th highlighted. To the right, a list of snapshots is shown for December 12, 2018, with 3 snapshots in total. The first snapshot is at 15:23, labeled "C cloned" with a size of 5.2 MB. A context menu is open over this snapshot, showing "Recover files" and "Delete clone" options. Below the list, a "Summary" section indicates the oldest snapshot is from 2018/12/12 and there are 3 total snapshots.

Time	Label	Size
15:23	C cloned	5.2 MB
14:33	C	798.6 MB

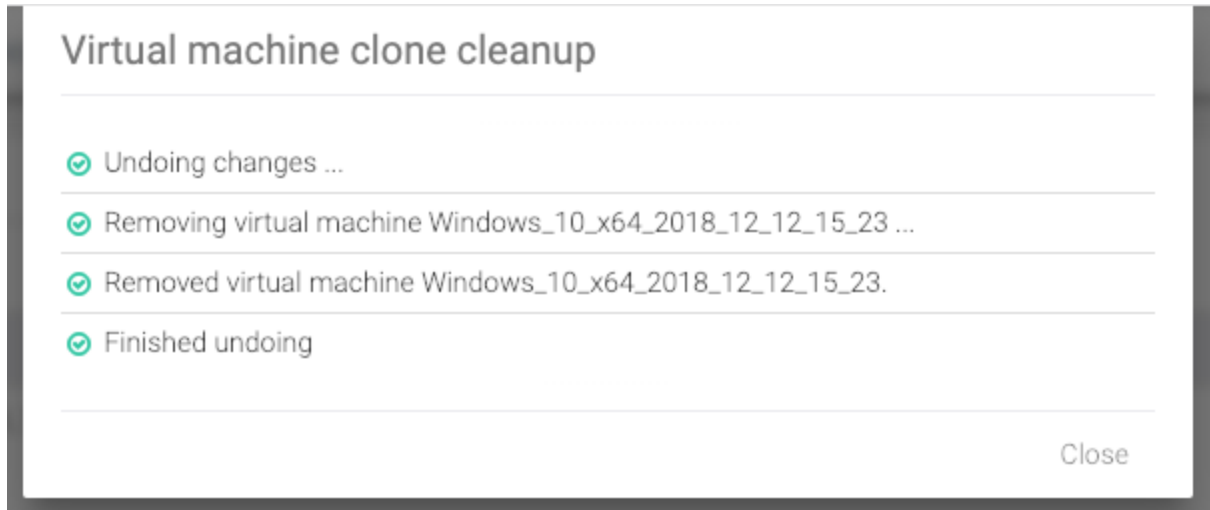
Summary

Oldest snapshot on	2018/12/12
Total snapshots	3

The system will ask you to confirm the process and it will explain all the things that will happen, such as virtual machines will be removed, datastore unmounted from VMware, etc.



The process takes a little time, you will see a step by step progress.



That's it, all that was created, shared, connected, imported for that snapshot clone was reverted. No more mess left behind by the recovery process.

Migrating Virtual Machines


SynetoOS offers the possibility to migrate Virtual Machines from your existing, non-Syneto infrastructure. That is, if you have Virtual Machines running on a different hardware, inside an ESXi or vCenter, SynetoOS provides you with a simple mechanism to migrate some or all of your Virtual Machines to your HYPER hardware.

NOTE: The scope of this feature is limited to existing installations running on VMware. Hyper-V, KVM and other hypervisors are not supported.

NOTE: Though not intended for this purpose, you could also use SynetoOS to import Virtual Machines from existing Syneto infrastructure (other HYPER machines).

How to migrate virtual machines

After updating your SynetoOS to version 4.4 or higher (or, after installing a new SynetoOS 4.4 or higher) you will be notified with the following message:




This way, you can easily begin the process of migrating Virtual Machines from existing infrastructure as soon as you have Syneto OS 4.4 (or higher) available.

If you do not want to migrate virtual machines at this time, you can do it later, by accessing the *Migrate* link from the *Virtual machines* page.

Virtual machines | [Migrate](#) [Rescan](#)     

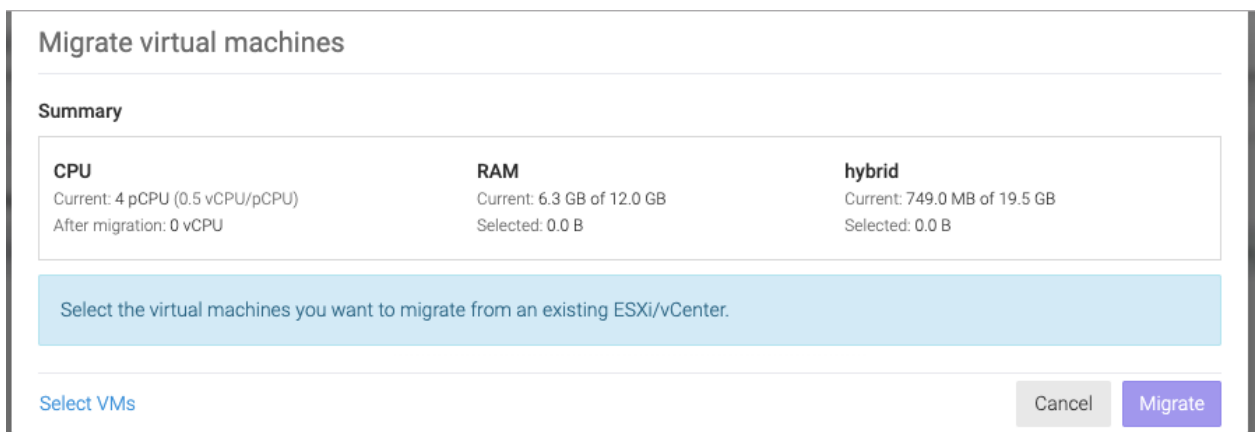
If you click *Migrate* (either the button in the notification, or the link in the page), a modal will appear, as you can see in the image below.



The preparations may take a few minutes.

NOTE: The Migration feature requires a running instance of *SynetoFileRecovery* Virtual Machine. You should not remove this virtual machine from your HYPER installation.


As soon as the migration service is ready, the modal will change into something like below.



Here you can see various resource information about the local hypervisor, which will help you determine how many and the kinds of Virtual Machines you can migrate onto the machine.

In the image above, there are 4 physical CPUs and 12 GB of RAM installed. Since there are 2 Virtual machines running at the time, each with 1 virtual CPUs, the ratio of CPU usage is 0.5 (i.e. 2 virtual CPUs running on 4 physical CPUs). Also, the current RAM usage is displayed as 6.3 GB. There is one disk pool set up, *hybrid*, which has a capacity of 19.5 GB and a current usage of 749 MB.

If you believe that your HYPER system has enough resources, you can go ahead and click *Select VMs*. The following modal will appear:



Connect to host

Host
192.168.5.40

Username
root


Password
.....

Cancel Connect






Here you will specify the IP address or hostname of the ESXi / vCenter you wish to migrate Virtual Machines from, along with the credentials.

NOTE: Make sure that you specify an external host (ESXi / vCenter), and not the IP / hostname of the local hypervisor. Migrating virtual machines from your HYPER machine onto the same machine is not allowed.

Attempting to connect to your local hypervisor from here will result in an error:





After you filled in the connection details, and clicked *Connect*, SynetoOS will connect to the specified ESXi / vCenter and fetch the list of Virtual Machines and details from there.


Select virtual machines					Refresh
Virtual machine		State	vCPU	RAM	Storage
<input type="checkbox"/> vm1  		Powered on	1	4.0 GB	7.2 GB
<input type="checkbox"/> vm2  		Powered on	1	1.0 GB	5.1 GB
<input type="checkbox"/> centos7		Powered off	1	2.0 GB	2.0 KB
<input type="checkbox"/> windows 		Powered off	1	2.0 GB	126.9 KB

Cancel

In the case above, there were 4 Virtual Machines found on the host: 2 are powered on and 2 are powered off (as seen in the *State* column). In the *vCPUs* and *RAM* columns you can see the associated values configured for each specific Virtual Machine - for instance, *vm1* was configured on the hypervisor to have 1 vCPU and 4 GB RAM. The *Storage* column displays the amount of disk space associated with the Virtual Machine:

- If all the disks of the VM are thin provisioned, then the value represents the total disk usage of the VM (i.e. the sum of disk usage of each disk)
- If all the disks of the VM are thick provisioned, then the value represents the total disk space configured for the VM (i.e. the sum of the disk capacity value of each disk)
- If some of the disks of the VM are thin provisioned and others are thick provisioned, then the value represents the combined sum of disk usage and disk capacity, for thin and thick provisioned disks, respectively.

NOTE: The  symbol means that the migration can happen, but there are things you should take into consideration. The  symbol gives information why the migration cannot be performed on the Virtual Machine. Hover the cursor over the symbols to learn the details.

NOTE: If you want to migrate a Virtual Machine that has thick provisioned disks, the disks of that Virtual Machine will become thin provisioned on your local HYPER. Also, that Virtual Machine will have a warning () symbol next to it.

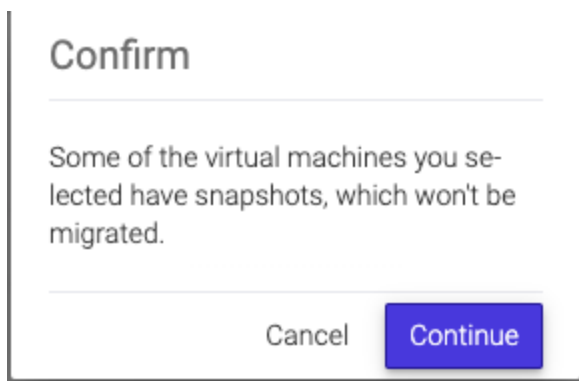
NOTE: The checkbox will become greyed out if it cannot be checked.

There are situations when the Virtual Machine can't be migrated. For example:

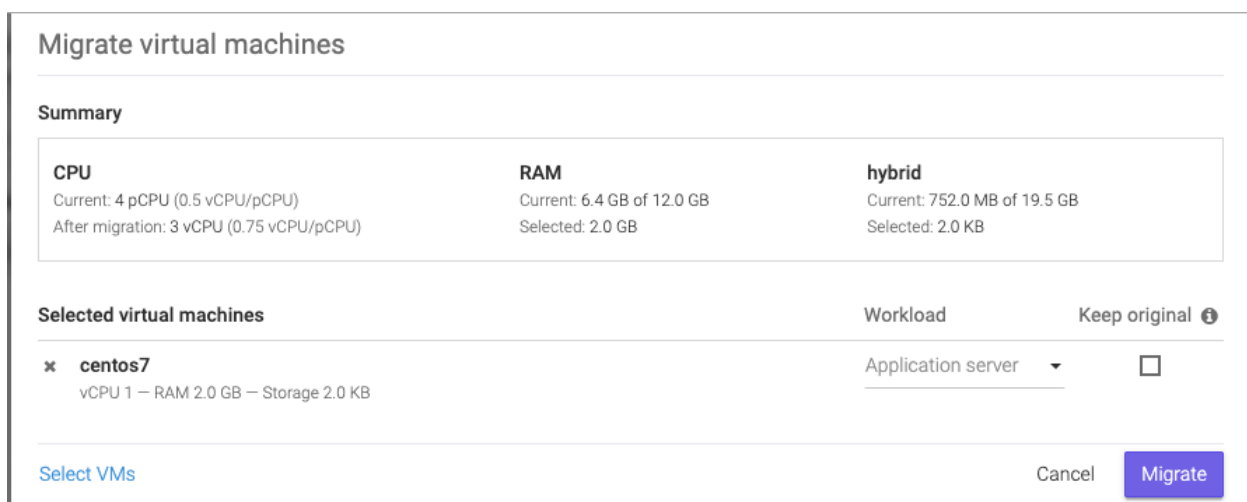
- The VM is in a *Suspended* state.
- The VM is *Running*.
- The VM is already being imported.
- The hardware version of the external host is greater than that of the local hypervisor.

After you checked the virtual machines you want to migrate, click *Select*.


If one or more of the virtual machines has snapshots, then, upon clicking *Select*, a dialog will appear, telling you about the implications, and asking you to confirm.



Upon confirming (or, if no confirmation dialog appeared, due to no potential issues), you will return to the first modal, where you will see the system resources of your local hypervisor, along with the Virtual Machines you have selected.



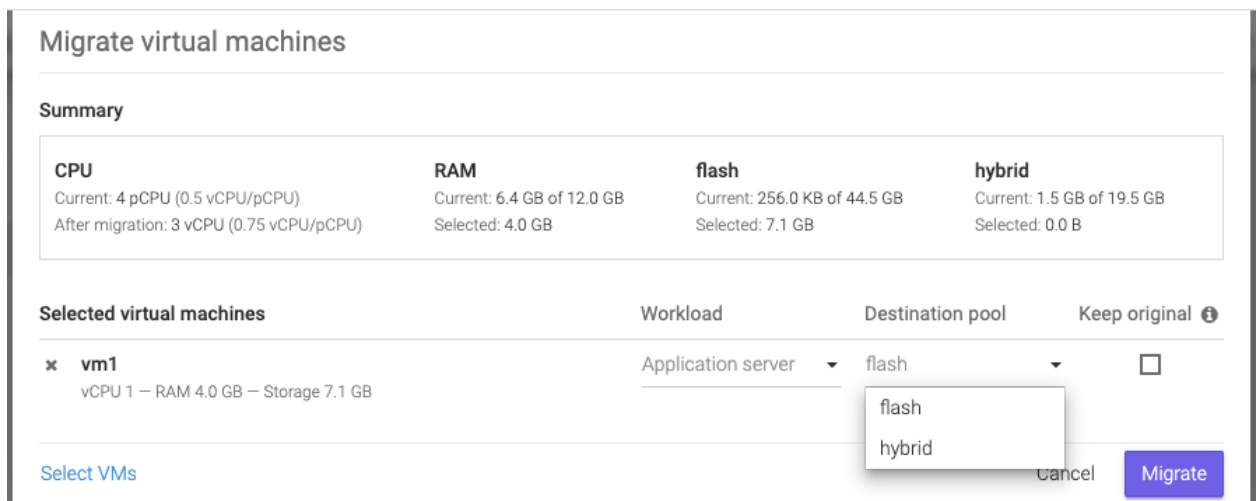
In the image above, the Virtual Machine *centos7* was added for migration. The workload was automatically set to *Application server*, but it can be changed.



Also, since there is only one pool on this SynetoOS machine, *hybrid*, the Virtual Machine will automatically be migrated to this pool.

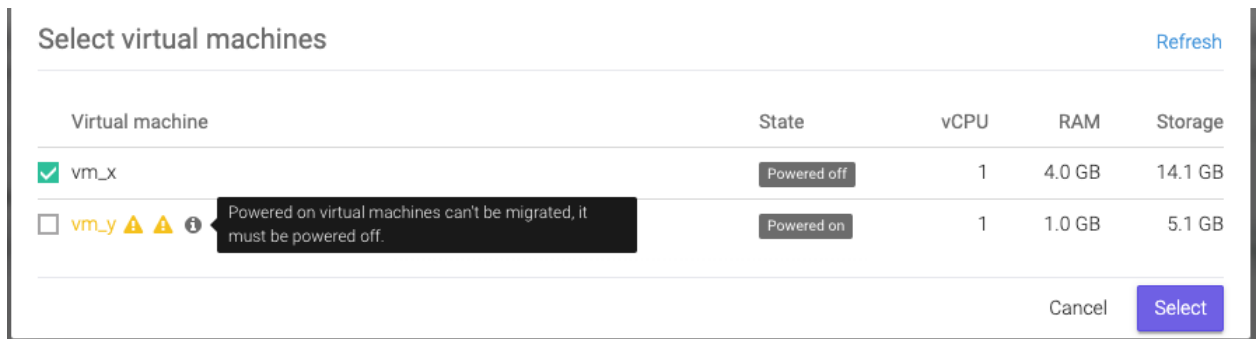
NOTE: If the *Keep original* field is unchecked, the Virtual Machine will be removed from the external host (ESXi / vCenter) after the migration is completed. Otherwise, the Virtual Machine will simply be copied here.

If you have more than one pool configured on your SynetoOS, then you will be able to select which of them you want to migrate the Virtual Machine on.



In the image above, there are 2 pools: *flash* and *hybrid* - they both appear in the *Summary* section. The Virtual Machine *vm1* was selected for migration. However, since there are more than one pools, a new column has appeared for *Selected virtual machines*: *Destination pool* - from here you can choose which of the two disk pools to migrate the Virtual Machine on.

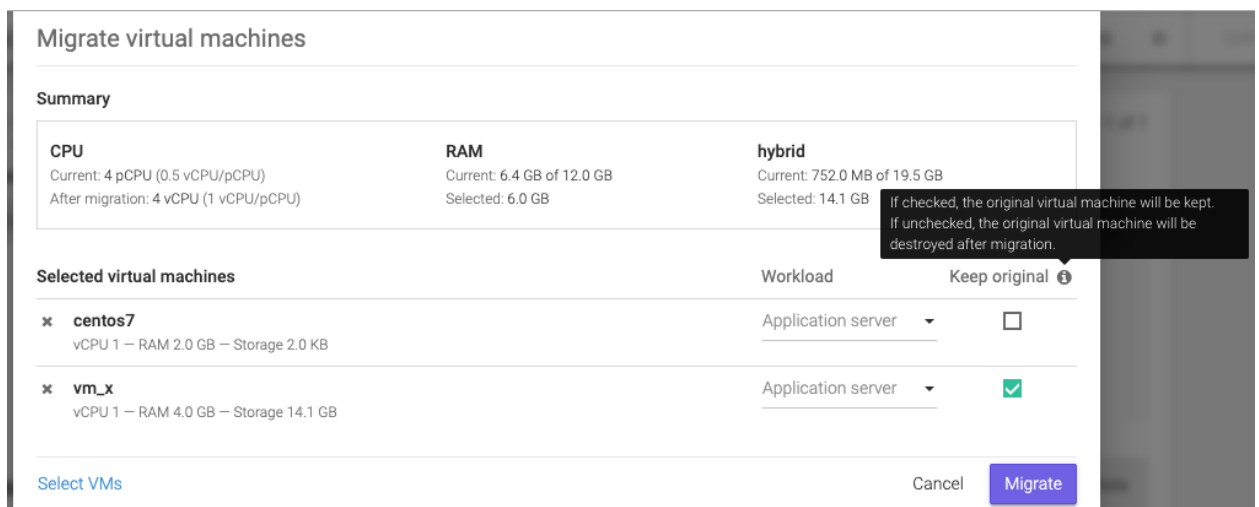
After you configured the *Workload*, *Destination pool*, and whether to keep the original Virtual Machine or not, you can click *Select VMs* to add more Virtual Machines to the list - say, from a different ESXi or vCenter.



In the image above, SynetoOS has connected to a different host, where other virtual machines are available.

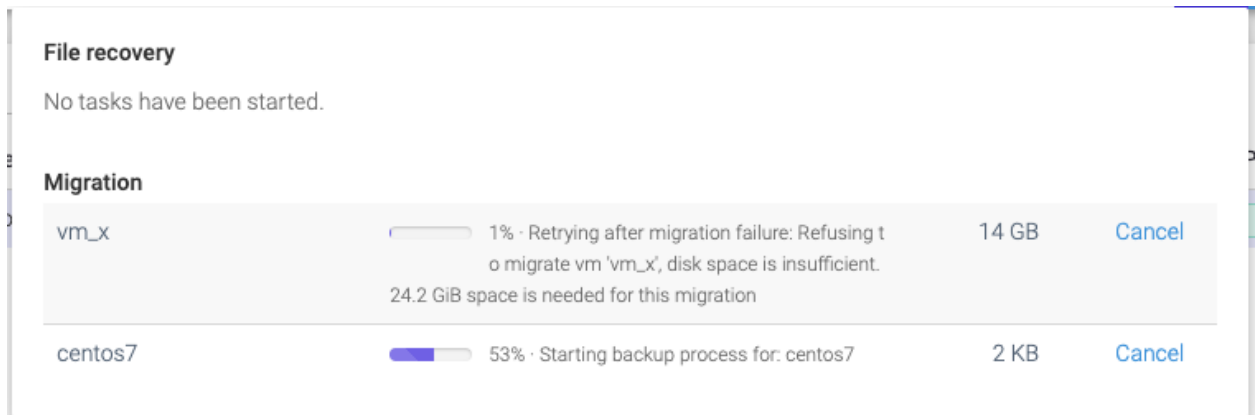
NOTE: If you hover over the information or warning icons, you can read the associated information.

Click *Select* to add the additional Virtual Machines. This will bring you back to the first modal, where you can see all the Virtual Machines you have selected, from whichever external hosts.

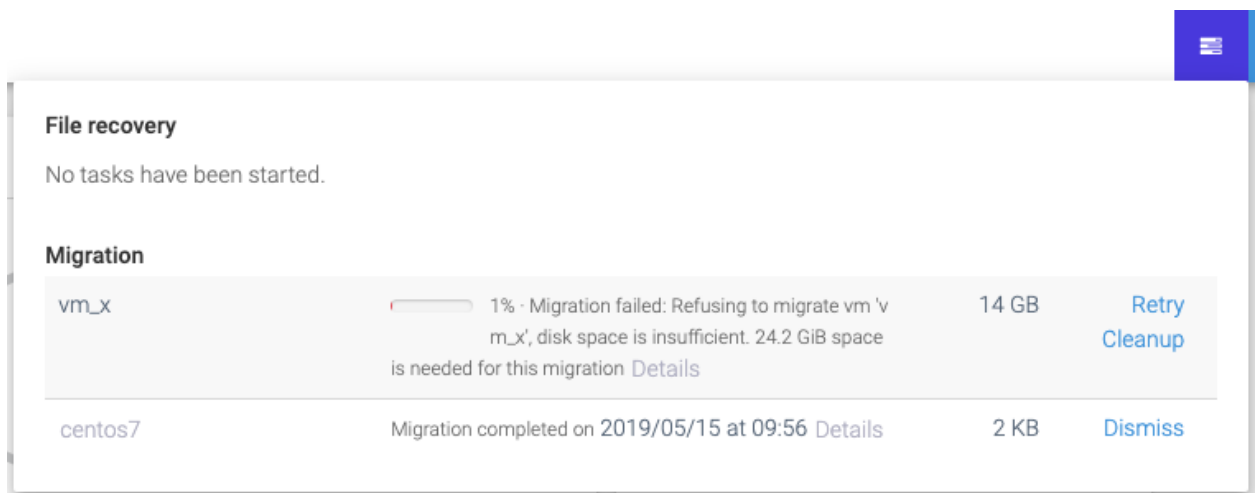


In the image above, *vm_x* was configured to have the original preserved, while *centos7* is set to be moved (i.e. after migration, the VM will be destroyed from the external host).

Once you hit *Migrate*, the migration process will begin.



In the image above, *centos7* is being migrated while *vm_x* is failing. You can cancel the migration of any virtual machine at any time.



In the image above, the migration process has finished.

NOTE: You can anytime access the tasks dialog from the button at the right side of the top bar.

After the migration of a virtual machine finishes, the name of the successfully migrated virtual machine (in the image above, *centos7*) becomes a link: if you click the name (in this case, *centos7*), the page will change to *Virtual machines* where you will see it in the list of virtual machines installed on your HYPER system.

After the migration of a virtual machine finishes or fails, a link will appear in the description of the task, *Details*. Clicking on it will open a modal where you will see the list of operations it performed.

Migration log

2019/05/15 09:41:52	Starting VM migration
2019/05/15 09:42:02	Creating datastore: centos7
2019/05/15 09:42:06	Datastore created at: hybrid/datastores/centos7
2019/05/15 09:42:06	Preparing for full backup
2019/05/15 09:42:07	Datastore setup complete
2019/05/15 09:42:11	Mounting datastore 'hybrid/datastores/centos7'
2019/05/15 09:42:16	Datastore mounted at '/tmp/vdb-centos7'
2019/05/15 09:42:17	Starting backup process for: centos7
2019/05/15 09:56:01	Datastore 'hybrid/datastores/centos7' unmounted
2019/05/15 09:56:04	Starting snapshot process
2019/05/15 09:56:04	VMware Datastore name: centos7
2019/05/15 09:56:04	Host: syneto-os-7919f26c
2019/05/15 09:56:04	Taking datastore snapshot
2019/05/15 09:56:05	Datastore snapshot completed
2019/05/15 09:56:05	Operation complete. Ingested 16.0 GiB of data for VM 'centos7'
2019/05/15 09:56:20	Completed VM migration

Close

In the image above, the *Details* link was clicked for the *centos7* virtual machine.

SHARES

Syneto HYPER products offer traditional file sharing over the SMB and AFP protocols. We know that hyperconverged solutions are the trend, but classic, old school sharing isn't dead either. So, let's see how to configure file sharing on Syneto HYPER.

User management

When we need to share files to several users, the first question that comes to mind is "Who can access the files?" Syneto HYPER offers user management and file access management in two contexts: Microsoft Active Directory integration or local users on the device itself.

Active Directory integration

To join Active Directory go to *Main menu* → *Shares* and select *Workgroup / Domain* in the action bar.

Join to

Workgroup **joined** | Active Directory

Domain name
DN.MYCOMPANY ?

Domain controller
dc.dn.mycompany ?

Username
Administrator

Password
.....

LAN manager authentication level 0 1 2 3 4 5 ?
Select lmauth level 4 if using Windows 2008 SP2.

Limit organizational units access

Cancel **Join**

Click on the *Active Directory* tab. Fill in the form with the required data and click *Join*. Once joined you manage your users from the Active Directory server. File permissions are also set from Windows.

NOTE: *Username* must be an Active Directory username with sufficient rights to join hosts to Active Directory.

WARNING: You must make sure DNS is set to the *Domain controller's* IP address. Otherwise joining the Active Directory will fail. See [Networking](#) for details about how to configure DNS server on the Syneto HYPER.

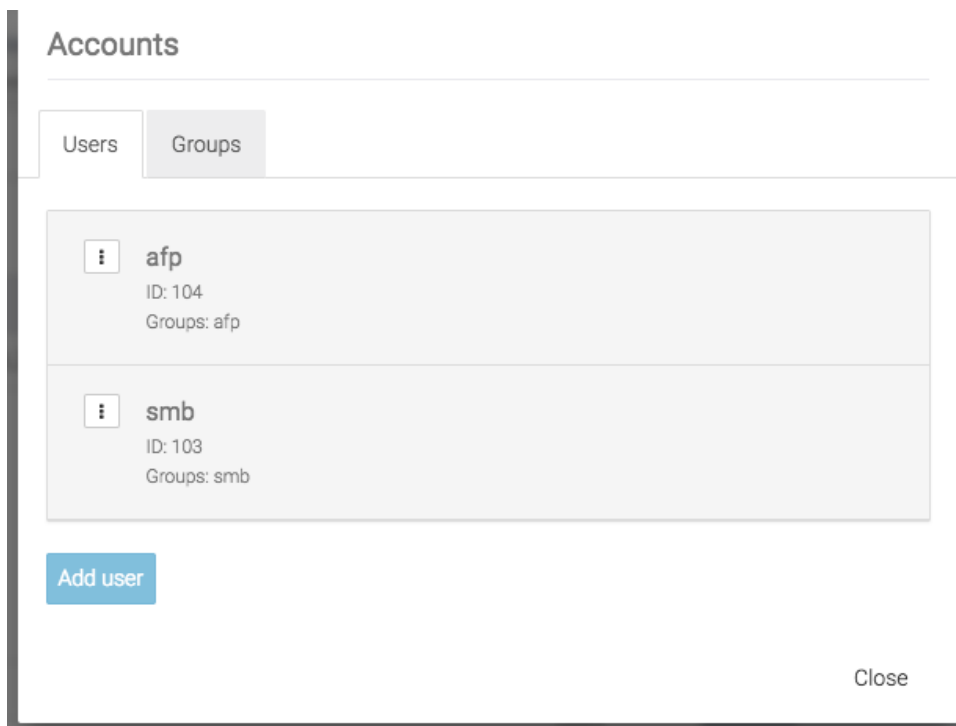
Managing local users and groups

In case you are not using an Active Directory server, or you simply wish to control access to shares with users local to the Syneto HYPER, you have the option to do so.

First, go to *Main menu* → *Shares* and select *Workgroup / Domain*. Click the *Workgroup* tab and join a workgroup.

NOTE: *Workgroup* is the default setting.

When you are joined to a workgroup, you can manage local users and groups. Click on *Accounts* in the action bar. The dialog will present two tabs: *Users* and *Groups*.



NOTE: By default there are two users created: **afp** with default password **afp** and user **smb** with default password **smb**.

The default users are created for simple networks. They provide a simple preconfigured security option. Just make your users use the smb or afp user with the default passwords.

If this is not enough for your case, just click *Add user* and add a new user.

New user

Username
jane

Password
.....

Re-type password
.....

Groups

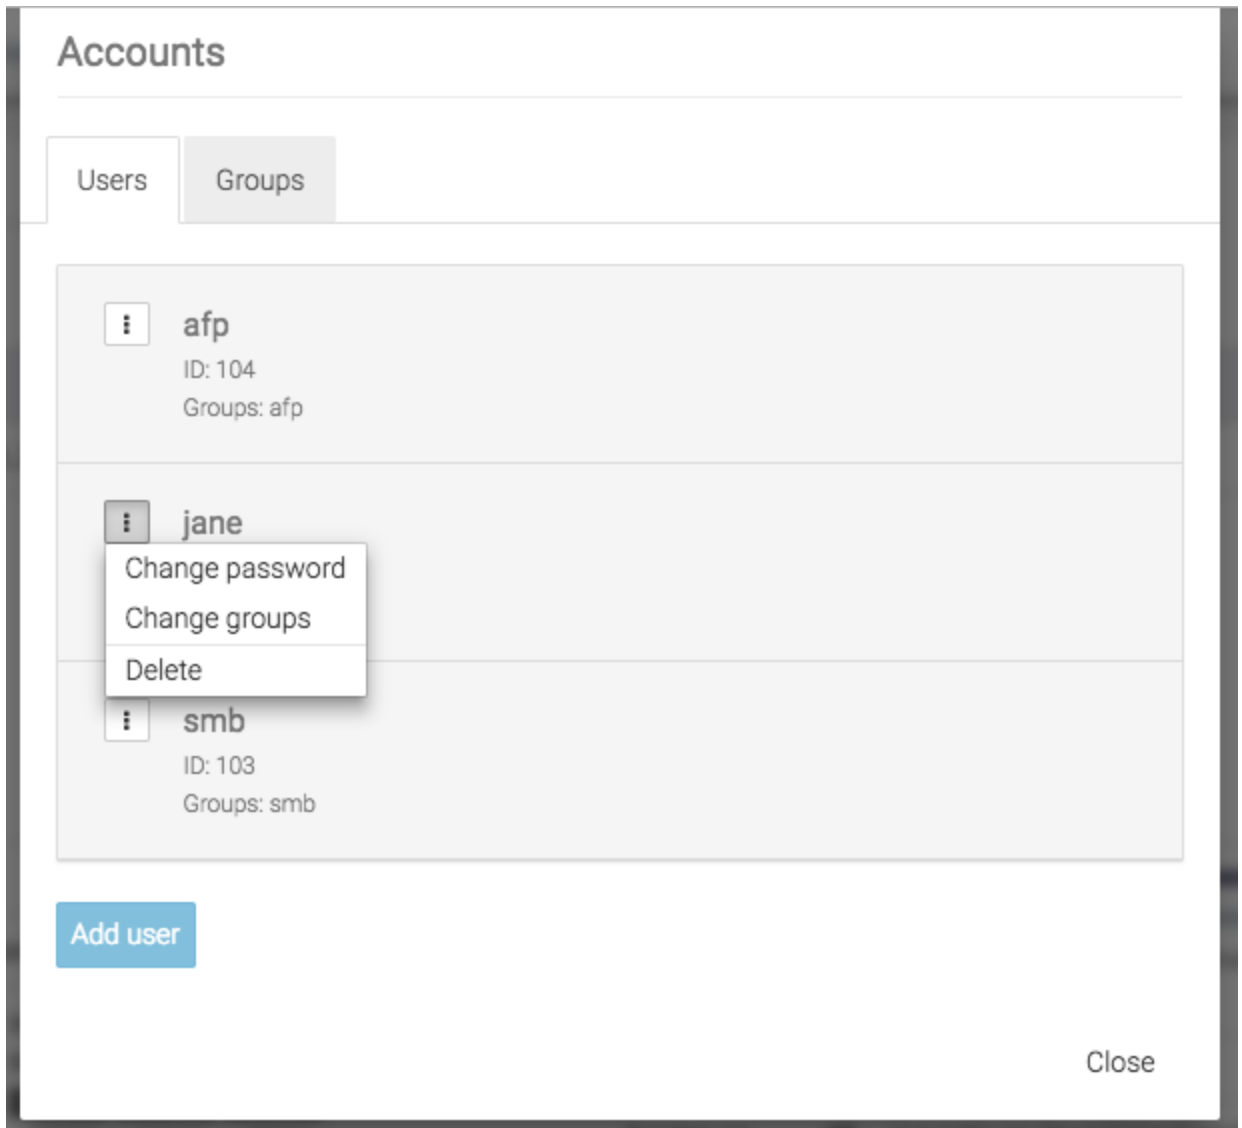
afp

smb

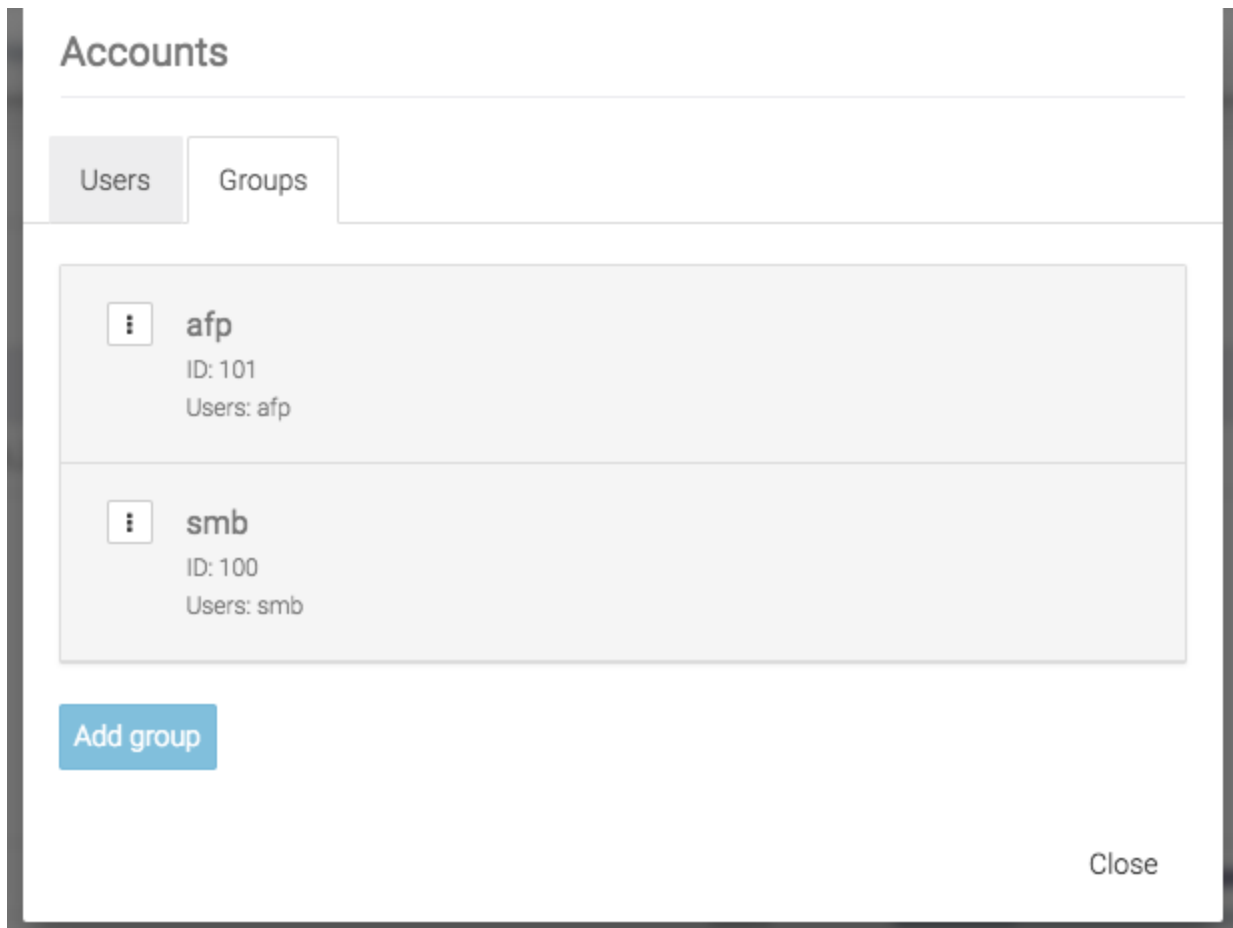
Cancel

NOTE: Each user will be created with a default group having the same name as the user. You can select secondary groups for each user from the existing groups.

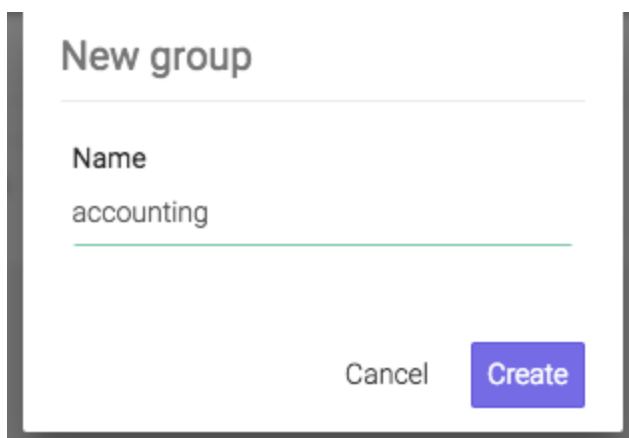
In the user list you can click the context menu and manage the user's groups, password, and delete the user.



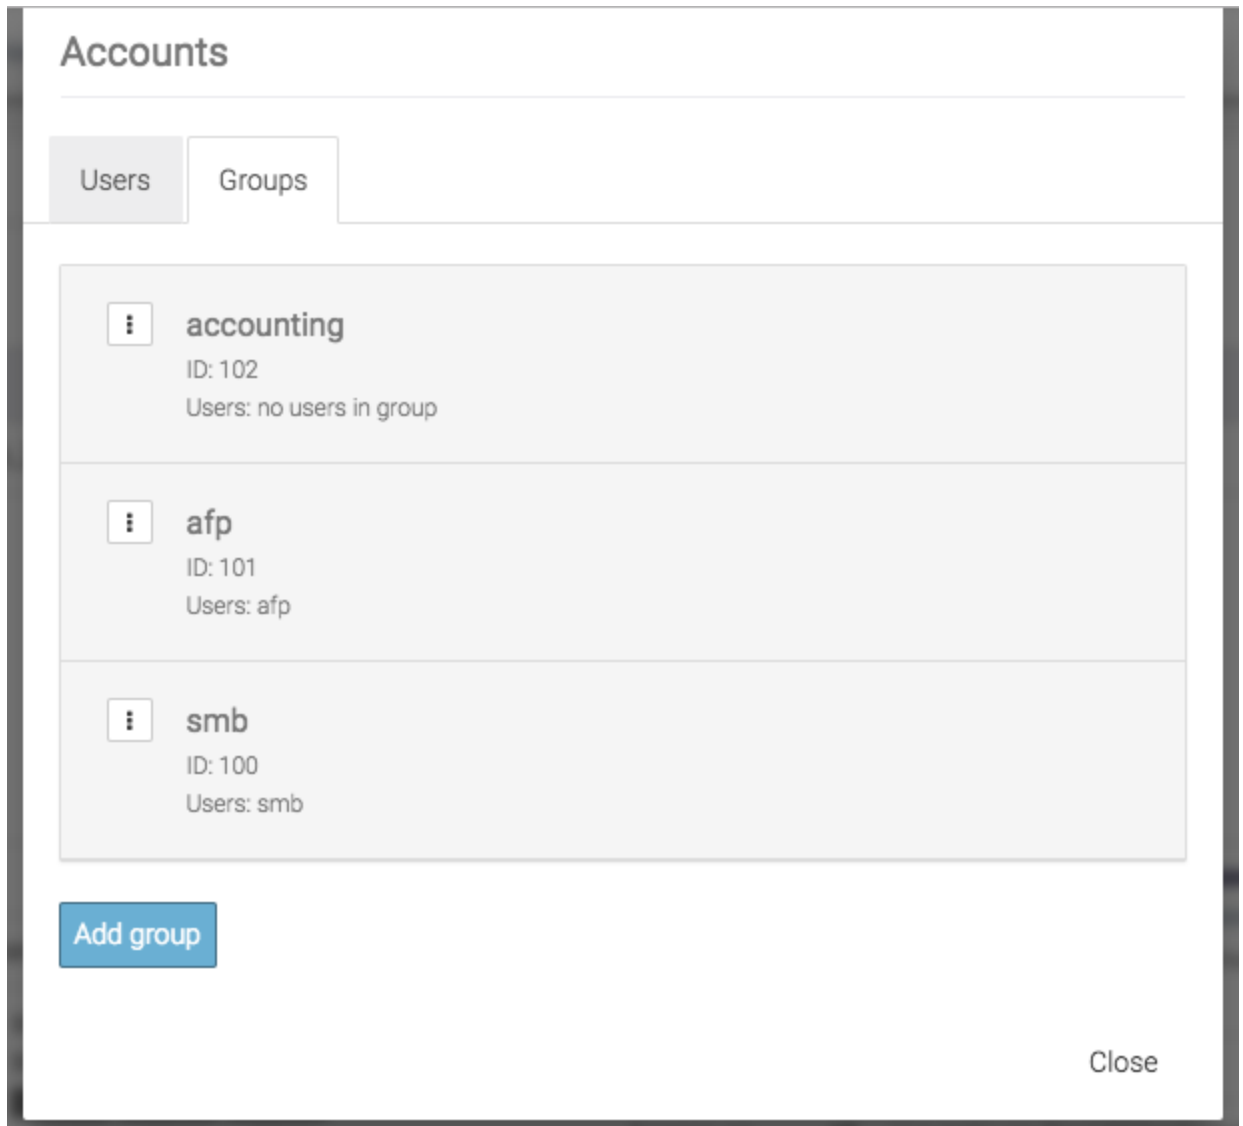
Adding and removing groups is similar to users.



Just click *Add group* and fill in a name for the group.



And then, click *Create* to create the group.



The new group will appear in the list.

Provisioning and managing

SMB - Create share

To create a new SMB share go to *Main menu* → *Shares* and click on *New SMB share*.

New SMB share

Name
document

Store on

hybrid /2.0 TB free space

testeee /79.5 GB free space

Allow guest access

Set quota

Change permissions (owner, group, rights)

Cancel Share

Provide a name for the share and optionally set one or more of the additional parameters:

- *Allow guest access* - makes the share public. Anybody can read, write, and modify anything.

NOTE: Creating or setting an SMB share's guest access, will automatically set the rights to everybody read/write/execute.

- *Set quota* - sets a quota on a share. In other words, how much data can be written into it.
- *Change permissions* - allows setting a user, group, and UNIX type file permissions.

New SMB share

Name
document

Store on
 hybrid /2.0 TB free space
 testeee /79.5 GB free space

Allow guest access

Set quota _____ GB

Change permissions (owner, group, rights)

Owner Group
smb ▾ smb ▾

Rights	Read	Write	Execute
Owner	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Group	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Others	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Cancel

SMB - Manage share

After the share was created you can see it in the list.

Search Show Originals, Clones, » Replica: ▾ 1-1 of 1

Name	Type	Used space	Quota	Protection
● document	SMB	24.0 KB	1.0 GB	NOT PROTECTED

document [Pause](#) [Delete](#) [Details](#) [Snapshots](#)

Type	SMB share	Space effective used	<div style="width: 100%;"><div style="width: 12.00 KB; height: 10px; background-color: #007bff;"></div><div style="width: 24.00 KB; height: 10px; background-color: #007bff;"></div></div> 12.00 KB 24.00 KB
Permissions	Owner smb Group smb Rights rwx rw- r-- owner group others	Change	<div style="width: 100%;"><div style="width: 24.00 KB; height: 10px; background-color: #007bff;"></div><div style="width: 0.00 B; height: 10px; background-color: #17a2b8;"></div></div> Data: 24.00 KB — Snapshots: 0.00 B
Quota	1.0 GB	Change	Compression 1.00x Saved 0.0 B
Protection	NOT PROTECTED	Protect	Access on ■ \\syneto-os-Dec1027e-adi\document 📄 🔗 smb://syneto-os-0ec1027e-adi/document 📄
	Guest access	OFF <input checked="" type="checkbox"/> ON	
	Description	[2018-12-14 15:19:19] SMB share document created.	Edit

Click on it and you can see its details. From the details section you can perform a set of management actions:

- *Pause* - will temporarily disable the share. You won't lose any data, but access to it will be impossible.

The screenshot displays the SynetoOS interface for managing shares. At the top, there is a search bar and a dropdown menu showing 'Show Originals, Clones, » Replica:'. Below this is a table with columns: Name, Type, Used space, Quota, and Protection. A single row is visible for the share 'document', which is marked as 'NOT PROTECTED'.

Below the table, the details for the 'document' share are shown. The share is an SMB share with a quota of 1.0 GB and is currently 'NOT PROTECTED'. The permissions are set to Owner: smb, Group: smb, and Rights: rwx for owner, rw- for group, and r-- for others. The space effective used is 24.00 KB, with 12.00 KB of data and 0.00 B of snapshots. The guest access is currently OFF. The description is '[2018-12-14 15:19:19] SMB share document created'.

When a share is paused its icon changes from the green triangle into an orange pause sign.

- *Delete* - will permanently remove the share and all its data.
- *Permissions* → *Change* - allows changing owner (user), group, and UNIX file permissions.
- *Quota* → *Change* - allows changing the quota of the share. You can increase or decrease the value. Minimum is 1GB. Set it to 0 for unlimited.
- *Guest access* → *OFF / ON* - toggles the guest access state of the share.
- View and edit the share's description. This field will shows by default the time when the share was created.

NOTE: After setting an SMB share's guest access, the rights to everybody read/write/execute will be set automatically.

AFP - Create share

To create a new AFP share go to *Main menu* → *Shares* and click on *New AFP share*.

New AFP share

Name

Store on

hybrid /2.0 TB free space

space /79.5 GB free space

Allow time machine backups

Enable spotlight search

Set quota

Change permissions (owner, group, rights)

Cancel

Provide a name for the share and optionally set one or more of the additional parameters:

- *Allow time machine backups* - enables support for MacOS to send Time Machine backups to this location.
- *Enable spotlight search* - enables server side indexing of this folder so you can search for files in your Mac's Spotlight.
- *Set quota* - sets a quota on a share. In other words, how much data can be written into it.
- *Change permissions* - allows setting a user, group, and UNIX type file permissions.
- View and edit the share's description. This field will show by default the time when the share was created.

New AFP share

Name

Store on

hybrid /2.0 TB free space

space /79.5 GB free space

Allow time machine backups

Enable spotlight search

Set quota

Change permissions (owner, group, rights)

Owner Group

Rights	Read	Write	Execute
Owner	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Group	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Others	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Cancel

AFP - Manage share

After the share was created you can see it in the list.

Search Show Originals, Clones, » Replica: ▾ 1-2 of 2

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
video	AFP	23.0 KB	200.0 GB	NOT PROTECTED

video [Pause](#) [Delete](#) [Details](#) [Snapshots](#)

Type	AFP share	Space effective used	<div style="width: 100%;"><div style="width: 12.00 KB; height: 10px; background-color: #007bff;"></div></div> 12.00 KB <div style="width: 100%;"><div style="width: 23.00 KB; height: 10px; background-color: #007bff;"></div></div> 23.00 KB ■ Data: 23.00 KB — ■ Snapshots: 0.00 B
Permissions	Owner afp Group afp Rights rwx rw- r-- owner group others	Change	Compression 1.04x Saved 0.0 B
Quota	200.0 GB	Change	Access on afp://syneto-os-0ec1027e-adi/video
Protection	NOT PROTECTED	Protect	Time machine OFF <input checked="" type="checkbox"/> ON Spotlight search OFF <input checked="" type="checkbox"/> ON
		Description	[2018-12-14 15:28:34] AFP share video created. Edit

Click on it and you can see its details. From the details section you can perform a set of management actions:

- *Pause* - will temporarily disable the share. You won't lose any data, but access to it will be prohibited.

Search Show Originals, Clones, » Replica: ▾ 1-2 of 2

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
video	AFP	23.0 KB	200.0 GB	NOT PROTECTED

video Resume Delete Details Snapshots

Type	AFP share	Space effective used	12.00 KB 23.00 KB
Permissions	Owner afp Group afp Rights rwx rw- r-- owner group others		■ Data: 23.00 KB ■ Snapshots: 0.00 B
Quota	200.0 GB	Time machine	OFF <input type="checkbox"/> ON <input type="checkbox"/>
Protection	NOT PROTECTED	Spotlight search	OFF <input type="checkbox"/> ON <input type="checkbox"/>
		Description	[2018-12-14 15:28:34] AFP share video created. Edit

When a share is paused its icon changes from the green triangle into an orange pause sign.

- *Delete* - will permanently remove the share and all its data.
- *Permissions* → *Change* - allows changing owner (user), group, and UNIX file permissions.
- *Quota* → *Change* - allows changing the quota of the share. You can increase or decrease the value. Minimum is 1GB. Set it to 0 for unlimited.
- *Time machine* → *OFF / ON* - toggles the time machine functionality.
- *Spotlight search* → *OFF / ON* - toggles the server side Spotlight integration.

Protecting

Please refer to [Configuring snapshotting without virtual machines](#) for more details.

Monitoring

The main shares page can be accessed by *Main menu* → *Shares*. This page offers a list of shares. The table combines SMB and AFP shares into a single list.

The screenshot displays the SynetoOS interface for managing shares. At the top, there is a search bar and a dropdown menu for 'Show Originals, Clones, » Replica:'. Below this is a table listing shares with columns for Name, Type, Used space, Quota, and Protection. The 'jane-backup' share is highlighted in blue. Below the table, there are buttons for 'Pause', 'Delete', 'Details', and 'Snapshots'. The 'Details' section for 'jane-backup' shows various settings: Type (AFP share), Permissions (Owner: afp, Group: afp, Rights: rwx for owner, rw- for group, r-- for others), Quota (unlimited), Protection (M), Space effective used (30.50 KB total, 58.00 KB used, with 34.00 KB for Data and 24.00 KB for Snapshots), Compression (1.05x, Saved 0.0 B), Access on (afp://syneto-os-0ec1027e-adi/jane-backup), Time machine (OFF), Spotlight search (OFF), and Description ([2018-12-14 15:56:58] AFP share jane-backup created).

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
jane-backup	AFP	58.0 KB	unlimited	M
marketing-materials	SMB	24.0 KB	unlimited	NOT PROTECTED
video	AFP	24.0 KB	200.0 GB	NOT PROTECTED

jane-backup Pause Delete Details Snapshots

Type	AFP share	Space effective used	30.50 KB 58.00 KB
Permissions	Owner afp Group afp Rights rwx rw- r-- owner group others		
Quota	unlimited	Compression 1.05x	Saved 0.0 B
Protection	M	Access on afp://syneto-os-0ec1027e-adi/jane-backup	
		Time machine OFF <input type="checkbox"/>	
		Spotlight search OFF <input type="checkbox"/>	
		Description [2018-12-14 15:56:58] AFP share jane-backup created.	

The list has several columns with information about the share: *Type*, *Used space*, *Quota*.

Type

Specified the share type: SMB or AFP. It can be used to quickly determine a share's type.

Used space

Show the amount of data written to the share. It can be used to determine space pressure on the Syneto HYPER.

Quota


Shows the quota set on the share. It can be used along with *Used space* to determine if space limits are reached.

Analyzing performance

Syneto HYPER collects performance data and saves them for later analyzes. This chapter will describe what analytics can be visualized for shares.

Dataset IOPS


Go to *Main menu* → *Analytics*, click *Add graph*, select *Dataset* → *IOPS*.



These graphics show the IOPS happening on the share from the Syneto HYPER's perspective.

Dataset bandwidth

Go to *Main menu* → *Analytics*, click *Add graph*, select *Dataset* → *Bandwidth*.



These graphics show the data transfer speed on the share from the Syneto HYPER's perspective.

Dataset average block size

Go to *Main menu* → *Analytics*, click *Add graph*, select *Dataset* → *Average block size*.

These graphics show the average block size of the data on the share from the Syneto HYPER's perspective.

Recovering data from a snapshot

Access a previous version of your share

Sometimes, it happens that you lose some data. Maybe a virus corrupted something, maybe a user deleted a file.

The screenshot shows the SynetoOS interface. At the top, there is a search bar and a dropdown menu for 'Show' with options 'Originals', 'Clones', and 'Replica:'. The main area displays a table of shares:

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
jane-backup	AFP	58.0 KB	unlimited	M
marketing-materials	SMB	24.0 KB	unlimited	NOT PROTECTED
video	AFP	24.0 KB	200.0 GB	NOT PROTECTED

Below the table, the details for the 'jane-backup' share are shown. The 'Snapshots' tab is highlighted with a red box. The details view includes:

- Type:** AFP share
- Permissions:** Owner: afp, Group: afp, Rights: rwx (owner), rw- (group), r-- (others)
- Quota:** unlimited
- Protection:** M
- Space effective used:** 30.50 KB (Data: 34.00 KB, Snapshots: 24.00 KB)
- Compression:** 1.05x, Saved 0.0 B
- Access on:** afp://syneto-os-0ec1027e-adi/jane-backup
- Time machine:** OFF (ON)
- Spotlight search:** OFF (ON)
- Description:** [2018-12-14 15:56:58] AFP share jane-backup created.

For all these, and many more cases, you can simply click on the Snapshots tab in the details view of a share.

jane-backup Pause Delete Details Snapshots

December 2018

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

December 14, 2018 2 snapshots

▶	ⓘ	16:01	24.0 KB
▶	ⓘ	16:00	24.0 KB

Summary

Oldest snapshot on	2018/12/14
Total snapshots	2

Pick a date from the calendar, and click on the contextual menu for a snapshot.

jane-backup Pause Delete Details Snapshots

December 2018

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

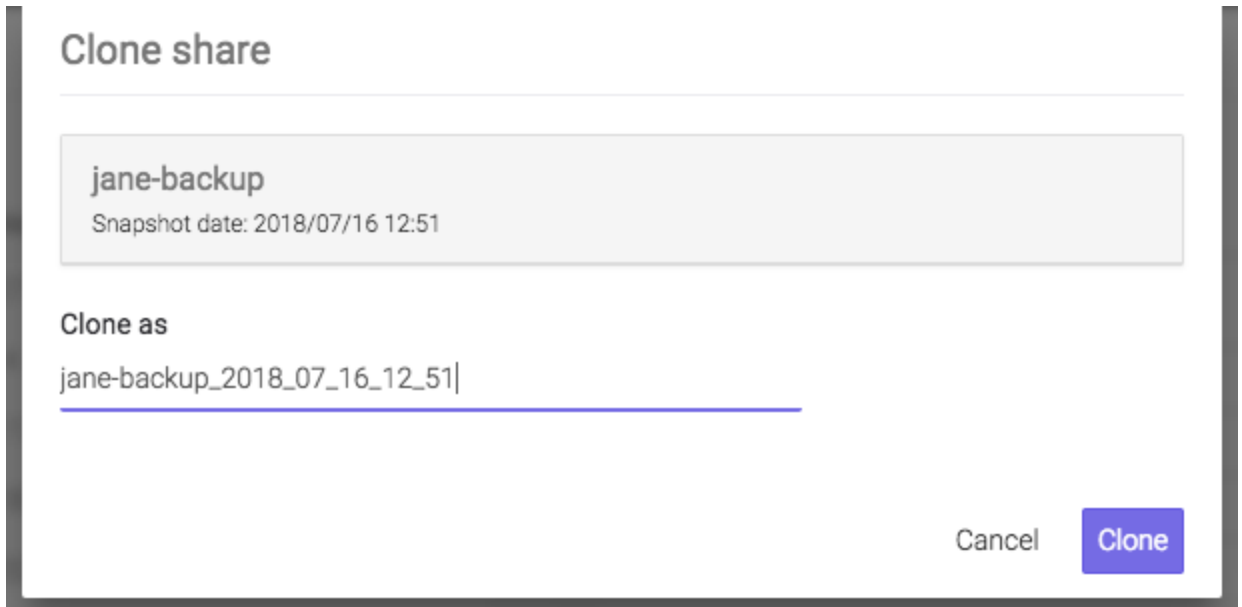
December 14, 2018 2 snapshots

▶	ⓘ	16:01	24.0 KB
▶	ⓘ	16:00	24.0 KB

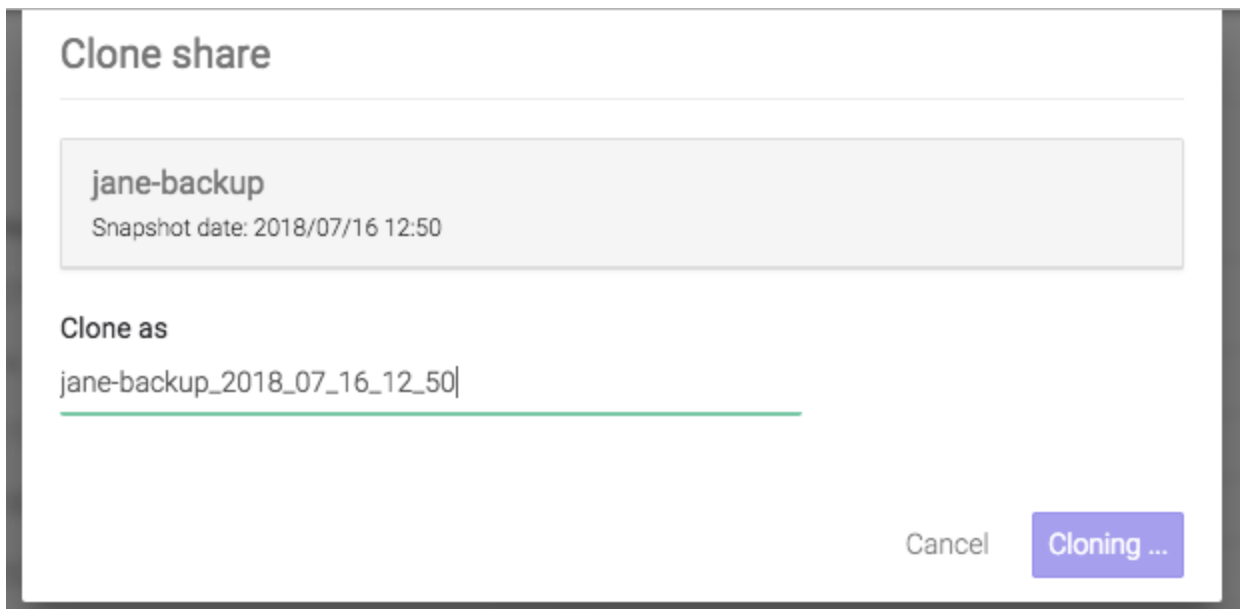
Summary

Oldest snapshot on	2018/12/14
Total snapshots	2

Select clone.



Click clone and wait for the process to finish.



Now you can access a share with the name specified at the clone step. It will contain the data as it was present at the moment of taking the snapshot.

WARNING: Do not change data on the clones. Clones cannot be protected, and after the Cleanup process (see below), all changes will be deleted.

Windows previous versions

Microsoft Windows users can benefit from built-in integration between Syneto HYPER and Microsoft's "Previous Versions" functionality.

On your Windows machine, on any given share, you can right click a folder or file and choose "Previous Versions". This will show a list with all the versions of the selected folder or file that are part of Syneto HYPER snapshots.

You can see the content for these folders or files directly from your Windows machine, no need to do anything on the Syneto HYPER.

Clean up when you are done

One of the more tedious tasks after recovering data is cleaning up after you. Deleting storage clones and shares can be a tedious operation. And because it is time consuming and difficult, many people just let them there... forever.

With Syneto HYPER the whole cleanup process is just one click away. And we take care to clean up only what we created.

Go to *Main menu* → *Shares* and select a share that has cloned snapshots. Click *Browse* in the snapshot section.

Search 1-5 of 5

Show Originals, Clones, » Replica: ▼

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
jane-backup	AFP	82.0 KB	unlimited	
jane-backup_2018_12_14_16_01	AFP	14.0 KB	unlimited	NOT PROTECTED
marketing-materials	SMB	24.0 KB	unlimited	NOT PROTECTED
video	AFP	24.0 KB	200.0 GB	NOT PROTECTED

jane-backup

[Pause](#) [Delete](#) [Details](#) [Snapshots](#)

Type	AFP share	Space effective used	<div style="width: 42.5%;"><div style="width: 42.5%;"></div></div> 42.50 KB						
Permissions	Owner afp Group afp Rights <table border="1"><tr><td>rwx</td><td>rw-</td><td>r--</td></tr><tr><td>owner</td><td>group</td><td>others</td></tr></table>	rwx	rw-	r--	owner	group	others	Change	82.00 KB <div style="width: 100%;"><div style="width: 34%; background-color: #007bff;"></div><div style="width: 48%; background-color: #6c757d;"></div></div> Compression 1.03x Saved 0.0 B
rwx	rw-	r--							
owner	group	others							
Quota	unlimited	Change	Access on afp://syneto-os-0ec1027e-adi/jane-backup						
Protection		Protect	Time machine OFF ON						
Clones			Spotlight search OFF ON						
			Description [2018-12-14 15:56:58] AFP share jane-backup created. Edit						
Clone		Cloned on							
jane-backup_2018_12_14_16_01		2018/12/14 16:01							

You will see that all cloned snapshots are marked with a label. Click the context menu for a cloned snapshot and select *Delete clone*.

Search 1-5 of 5

Show Originals, Clones, » Replica: ▾

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
jane-backup	AFP	82.0 KB	unlimited	M
jane-backup_2018_12_14_16_01	AFP	14.0 KB	unlimited	NOT PROTECTED
marketing-materials	SMB	24.0 KB	unlimited	NOT PROTECTED
video	AFP	24.0 KB	200.0 GB	NOT PROTECTED

jane-backup Pause Delete Details Snapshots

December 2018

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

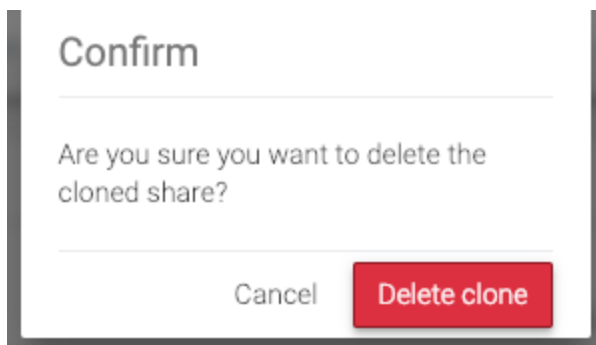
December 14, 2018 2 snapshots

<div style="display: flex; justify-content: space-between;"> 16:01 cloned 24.0 KB </div> <div style="font-size: 0.8em; margin-top: 5px;"> <p>Starting snapshot process ...</p> <p>AFP Share name: jane-backup ...</p> <p>Frequency: every 1 minutes ...</p> <p>Host: syneto-os-0ec1027e-adi ...</p> <p>Taking AFP Share snapshot ...</p> <p style="color: green;">Local snapshot taken.</p> </div>
<div style="display: flex; justify-content: space-between;"> 16:00 24.0 KB </div>

Summary

Oldest snapshot on	2018/12/14
Total snapshots	2

The system will ask you to confirm the process and it will explain all the things that will happen.



The process takes only a second, you will see a little spinner in the list.

That's it, all that was created and shared for that snapshot clone was reverted. No more mess left behind by the recovery process.

Reverting a clone

Once you created one or more clones from your snapshots, you may find a specific clone that contains exactly the data you wish to recover. The *revert* functionality allows you to replace your current share with the clone.

To revert a clone, select a clone, then click *Revert*.

The screenshot shows a table of clones with columns for Name, Type, Used space, Quota, and Protection. The clone 'jane-backup_2018_12_14_16_01' is selected. Below the table, the details for this clone are shown, including its type (AFP share), permissions (Owner: afp, Group: afp, Rights: rwx for owner, rw- for group, r-- for others), space effective used (7.00 KB), and access on (afp://syneto-os-0ec1027e-adi/jane-backup_20). The 'Revert' button is highlighted with a red box.

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
jane-backup	AFP	82.0 KB	unlimited	M
jane-backup_2018_12_14_16_01	AFP	14.0 KB	unlimited	NOT PROTECTED
marketing-materials	SMB	24.0 KB	unlimited	NOT PROTECTED
video	AFP	24.0 KB	200.0 GB	NOT PROTECTED

jane-backup_2018_12_14_16_01 Revert Pause Delete Details

Cloned from *jane-backup* (2018/12/14 16:01)

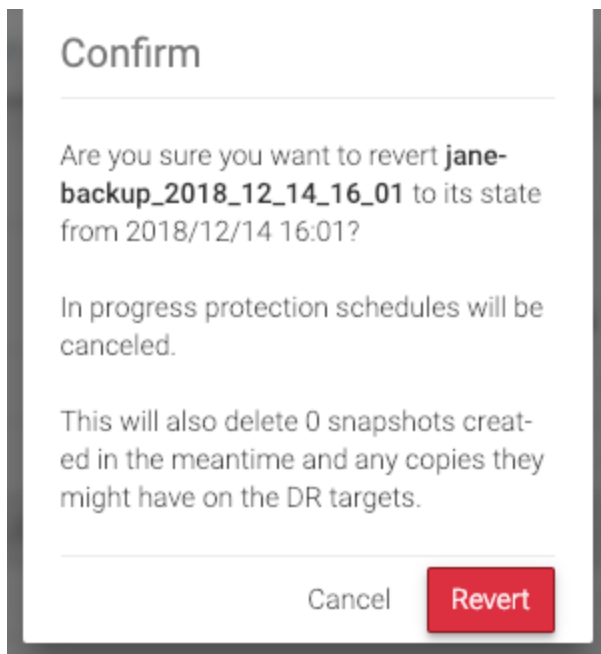
Type	AFP share	Space effective used	7.00 KB
Permissions	Owner afp Group afp Rights rwx rw- r-- owner group others	14.00 KB	14.00 KB
Quota	unlimited	Compression 1.00x	Saved 0.0 B
Access on	afp://syneto-os-0ec1027e-adi/jane-backup_20	Time machine	OFF <input checked="" type="checkbox"/> ON
		Spotlight search	OFF <input checked="" type="checkbox"/> ON
Description	[2018-12-14 16:11:05] AFP share jane-back-up_2018_12_14_16_01 created.		

A clone revert process will do the following steps:

- Delete all snapshots created after the snapshot from which the clone was created
- Replace the original share with the clone

WARNING: Once a clone is restored, all data and snapshots newer than the clone will be permanently lost.

NOTE: If you have several clones of the same share, only the clone from the most recent snapshot can be reverted.



During the revert process you will see a step-by-step progress. If something goes wrong, an appropriate message will be shown.



WARNING: On DR unit, on a received share, you can create clones but you cannot revert the clone.

After the revert process finishes, the system will update the datastore's *Description* field with the relevant information.

Description	[2018-07-16 15:13:00] AFP Share reverted from clone jane-backup_2018_07_16_12_50 [2018-07-16 15:12:38] AFP share jane-backup_2018_07_16_12_50 created.	Edit
-------------	---	----------------------

VOLUMES

Introduction

Volumes are virtual block devices shared over the iSCSI protocol. Such volumes can be accessed from VMware and from Microsoft Windows or Linux servers.

WARNING: This feature is not available on Syneto HYPER 2000 Series.


Access management

iSCSI remote initiators


Before provisioning a new volume, we need to specify at least one initiator. An iSCSI initiator is an IQN or EUI from a host that will try to connect to Syneto HYPER. You can define any number of initiators and you can organize them by groups.

Go to *Main menu* → *Shares* → *Initiators*.

This will show a dialog with the currently defined initiators.



By default, no initiators are defined. To add a new initiator click *Add initiator* and specify the group name and identifier.



Add initiator

Group: Create new group | New group name: Initiators-Group-1

Identifier: iqn.2013-02.net.syneto:02:457e299a-d509-4fe0-a882-f40360b2fca2

Cancel Add

NOTE: each identifier must be in a group named by you. This is necessary even if there is only one initiator in the group.

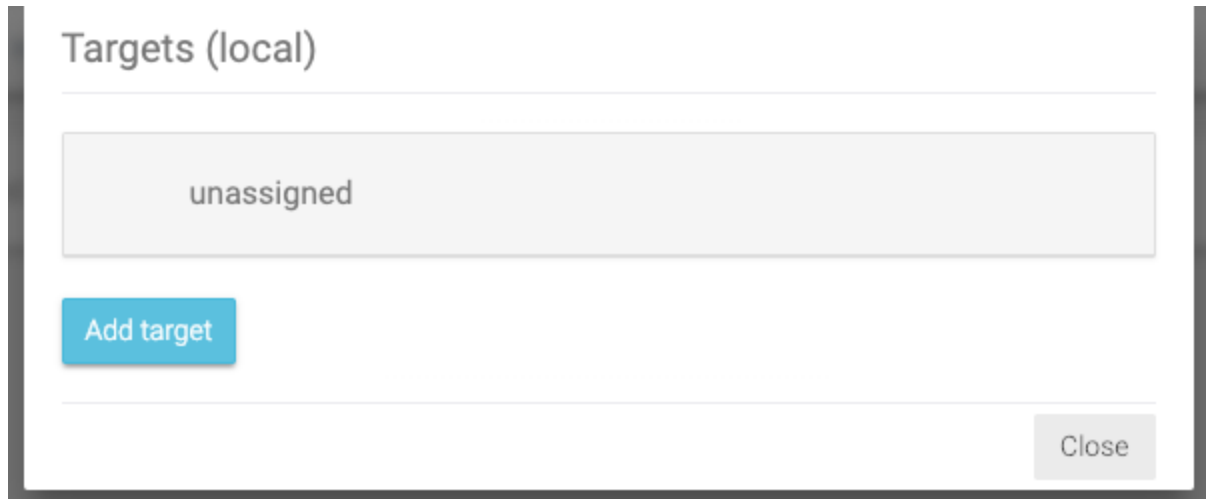
Once you are done, click *Add*. The new initiator will be shown in the list. From this list you can manage the initiators, move them from one group to another, or delete them.

iSCSI local targets

By default no local targets are defined on Syneto HYPER. Defining targets is optional. Volumes can be provisioned without specifying them.

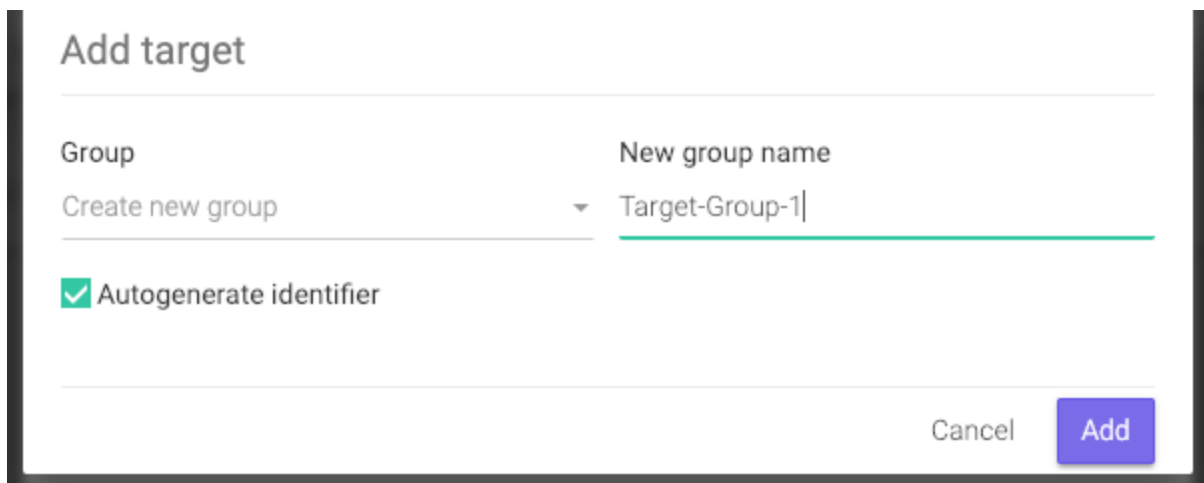
In case you wish more granular access control, you can specify local targets and groups by going to *Main menu* → *Shares* → *Targets*.

By default a list with an empty group named *unassigned* will be shown. This is a group containing targets that are not assigned to any target group. While initiators must be part of group always, targets do not need to. All these targets are part of the virtual *All* group.

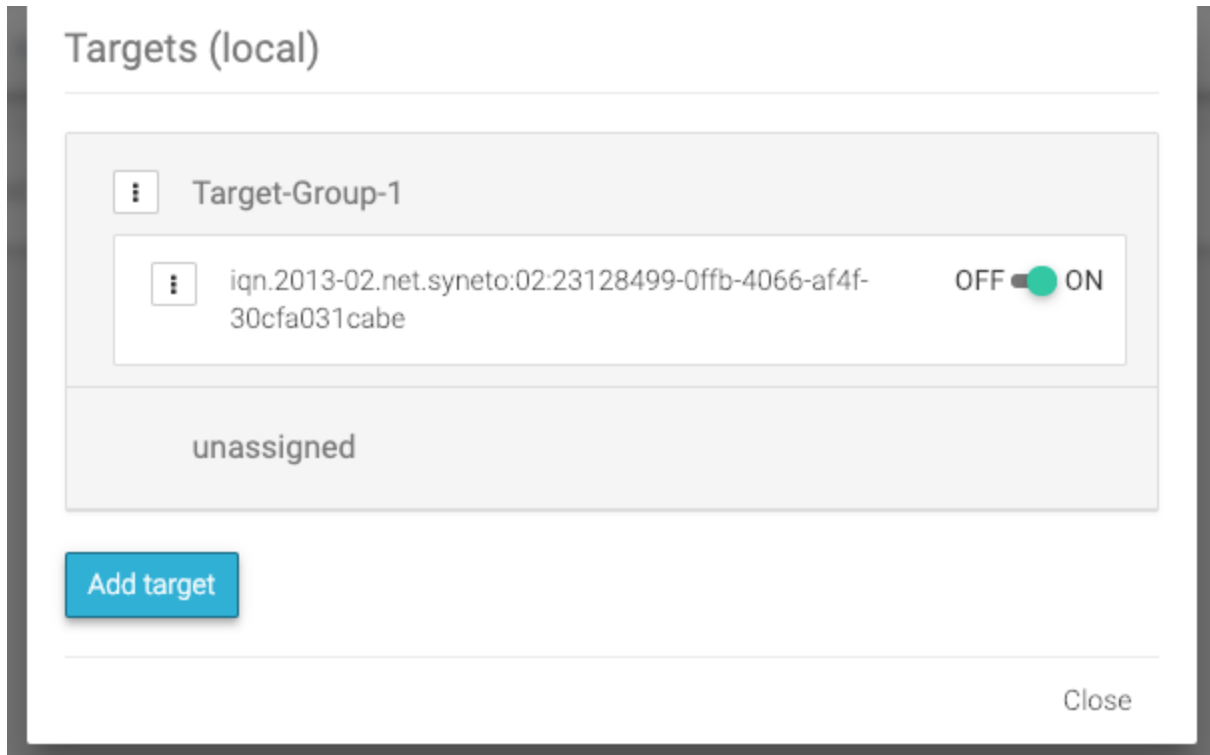


To create a new target click *Add target*. The form will ask you the group. Options are:

1. *Create new group* - creates a new group with the given name
2. *Unassigned* - adds the target to no group, accessible through the virtual *All* group
3. *No group* - same as *unassigned*



By default the *Autogenerate identifier* option will be selected. Optionally, you can uncheck it and specify an IQN yourself.



Once you are done, click *Add*. The new target will be shown in the list. From this list you can manage the targets, move them from one group to another, turn them on or off, or delete them.

Provisioning and managing

Provisioning a new volume is done on the *Main menu* → *Volumes* page. In the action bar click *New volume*.


New volume

Name
Volume1

Size
2 GB

Block size
32 KB (recommended) ▼

Provisioning Thin Thick

Enable write back cache 

iSCSI access

Initiator
Initiators-Group-1 ▼

Target
Target-Group-1 ▼

LUN
1

2.0 TB free space

Cancel **Create**

NOTE: An initiator must be defined before provisioning a new volume.

The dialog will require to fill in a set of information:

- *Name* - the name of the volume
- *Size* - the size in GB for the volume
- *Block size* - the block size for the volume. By default is 32KB as that is a common value. However, you should choose one to accommodate the type of traffic your applications are using. For example, database with many write operation may prefer an 8KB block size. A file server for video would run much better on high block size like 128KB.
- *Provisioning* → *Thin / Thick* - specifies if the whole virtual disk should be immediately reserved. Thin provisioned volumes allows you to over-provision the space you have available on your pool. Thick provisioned disks guaranties that there will be always enough underlying storage space up to the size specified by you.

WARNING: Be careful when creating snapshots of a thick provisioned volume. In order to guarantee the provisioned size, each snapshot will be as big as the provisioned space for the volume. The space will be used from the available space on the pool.

- When Provisioning is set to Thick, you will see an informative message when creating a schedule:

Adding a snapshot schedule for this volume can occupy a lot of space. [View details](#)

This volume is thick provisioned. Its snapshots are not incremental and can occupy a significant amount of space. Please monitor the pool size while the schedule is running and either reduce the retention policy or convert the volume to thin provisioning.

- *Enable write back cache* - enable an iSCSI protocol integration that will receive data to the memory first, confirm to the sender that the data is written, and flush the data to the drive at specific time intervals. This speeds up write operations a lot, but it may lead to data loss at sudden power failures.
- *iSCSI access* - requires you to specify the initiator group that can access this volume, and the local targets to which it is addressed. You must also provide a logical unit number (LUN).

After you are done, click *Create*. The Syneto HYPER will create the local volume, it will reserve the space if you requested thick provisioning, it will set up all the iSCSI configuration for it so that it will be immediately accessible from your network.

Once a new volume is created you will be able to find it on the *Main menu* → *Volumes* page.

Search 1-1 of 1

Show Originals, Clones, » Replica: ▾

Name	Used space	Provisioning	Protection
Volume1	12.0 KB	thin	NOT PROTECTED

Volume1 Delete Details Snapshots

Size: 2.0 GB Expand

Block size: 32 KB

Provisioning: THICK THIN

Write back cache: OFF ON

Protection: NOT PROTECTED Protect

Space effective used: 6.00 KB / 12.00 KB

Compression: 1.00x Saved 0.0 B

Description: [2018-12-14 17:19:50] Volume Volume1 created. Edit

iSCSI access Grant access

Initiator	Target	LUN
Initiators-Group-1	Target-Group-1	1

Volumes can be clicked in the list, and a detailed box in the bottom will appear. Here you can take a few volume management steps:

- *Delete* - deletes the volume, all the iSCSI shares related to it, all its data and snapshots.

NOTE: Volumes with clones cannot be deleted. See more about clones below.

- *Expand* - a volume's size can be increased, but not decreased. So we recommend you start with the minimum or nominal size your user will need and expand as necessary.
- *Provisioning* → *thick / thin* - you can change the provisioning type any time. Changing a volume from thin to thick provisioned will immediately "consume" the space from the pool up to the provisioned size. Switching a volume from thick to thin will release the space reserved for the the thick volume that is not used.
- *Writeback cache* - change the writeback cache type. See details above about this is.
- *iSCSI access* - this section allows you to allow access for other initiator groups by clicking *Grant access* or revoke access from an initiator group by clicking the context menu and selecting *Revoke access* in the *iSCSI access* list.
- View and edit the volume's description. This field will shows by default the time when the volume was created.

Protecting

Please refer to [Configuring snapshotting without virtual machines](#) for more details.

Monitoring

The main volumes page can be accessed by *Main menu* → *Volumes*. This page offers a list of volumes.

The screenshot shows the SynetoOS Volumes page. At the top, there is a search bar and a filter dropdown set to 'Show Originals, Clones, > Replica:'. The page shows 1-2 of 2 volumes. A table lists the volumes with columns for Name, Used space, Provisioning, and Protection. Below the table, the details for 'Volume1' are shown, including its size (2.0 GB), block size (32 KB), provisioning (THIN), write back cache (ON), and protection status. A bar chart shows 'Space effective used' with 6.00 KB of data and 12.00 KB of snapshots. The iSCSI access table shows one entry for Initiators-Group-1 connected to Target-Group-1 on LUN 1.

Name	Used space	Provisioning	Protection
Volume1	12.0 KB	thin	M W
Volume2	12.0 KB	thin	M H D W

Volume1 Delete Details Snapshots

Size: 2.0 GB Expand Space effective used: 6.00 KB (Data: 12.00 KB, Snapshots: 0.00 B)

Block size: 32 KB

Provisioning: THICK THIN

Write back cache: OFF ON

Protection: M W Protect

Description: [2018-12-14 17:19:50] Volume Volume1 created. Edit

iSCSI access Grant access

Initiator	Target	LUN
Initiators-Group-1	Target-Group-1	1

The list has two columns that can help you monitor what is happening: Used space and Provisioning

Used space

How much data is written on the volume.

Provisioning

Thick or thin, this value, combined with *Used space* can help you spot issues. For example if a user cannot write to a volume, but the user sees there is enough free space, you can check the *Provisioning* volume. Maybe your system is over-commissioned and there is no space left on the underlying storage.

When Provisioning is set to Thick, you will see an informative message when creating a schedule:

Adding a snapshot schedule for this volume can occupy a lot of space. [View details](#)


This volume is thick provisioned. Its snapshots are not incremental and can occupy a significant amount of space. Please monitor the pool size while the schedule is running and either reduce the retention policy or convert the volume to thin provisioning.

Analyzing performance

Syneto HYPER collects performance data and saves them for later analyzes. This chapter will describe what analytics can be visualized for volumes.

Dataset IOPS


Go to *Main menu* → *Analytics*, click *Add graph*, select *Dataset* → *IOPS*.



These graphics show the IOPS happening on the volume from the Syneto HYPER perspective.

Dataset bandwidth


Go to *Main menu* → *Analytics*, click *Add graph*, select *Dataset* → *Bandwidth*.



These graphics show the data transfer speed on the volume from the Syneto HYPER's perspective.

Dataset average block size

Go to *Main menu* → *Analytics*, click *Add graph*, select *Dataset* → *Average block size*.



These graphics show the average block size of the data on the volume from the Syneto HYPER's perspective.

Recovering data from a snapshot

Access a previous version of your volume

Sometimes, it happens that you lose some data. Maybe a virus corrupted something, maybe a user deleted a file.

Search Show Originals, Clones, » Replica: ▾ 1-2 of 2

Name	Used space	Provisioning	Protection
Volume1	12.0 KB	thin	M W
Volume2	12.0 KB	thin	M H D W

Volume1 Delete Details **Snapshots**

Size: 2.0 GB [Expand](#) Space effective used: 6.00 KB

Block size: 32 KB 12.00 KB

Provisioning: THICK THIN Data: 12.00 KB — Snapshots: 0.00 B

Write back cache: OFF ON Compression 1.00x Saved 0.0 B

Protection: M W [Protect](#) Description: [2018-12-14 17:19:50] Volume Volume1 created. [Edit](#)

iSCSI access [Grant access](#)

Initiator	Target	LUN
Initiators-Group-1	Target-Group-1	1

For all these, and many more cases, you can simply click on the Snapshots tab in the details view of a volume.

Volume1 Delete Details Snapshots

December 2018

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6


December 14, 2018 3 snapshots

- ▾ i 17:27 0.0 B
 - 17:27:38 Starting snapshot process ...
 - 17:27:39 ISCSI Volume name: Volume1 ...
 - 17:27:40 Frequency: every friday ...
 - 17:27:41 Host: syneto-os-0ec1027e...
 - 17:27:42 Taking ISCSI Volume snapshot ...
 - 17:27:43 Local snapshot taken.
- ▶ i 17:26 0.0 B
- ▶ i 17:25 0.0 B

Summary

Oldest snapshot on	2018/12/14
Total snapshots	3

Pick a date from the calendar, and click on the contextual menu for a snapshot.



Volume1 Delete Details Snapshots

December 2018

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6


December 14, 2018 3 snapshots

17:27	0.0 B
17:38	Starting snapshot process ...
17:39	ISCSI Volume name: Volume1 ...
17:27:40	Frequency: every friday ...
17:27:41	Host: syneto-os-0ec1027e...
17:27:42	Taking ISCSI Volume snapshot ...
17:27:43	Local snapshot taken.
17:26	0.0 B
17:25	0.0 B

Summary

Oldest snapshot on	2018/12/14
Total snapshots	3

Select *Clone*.



Clone volume

Volume1
Snapshot date: 2018/12/14 17:27

Clone volume as
Volume1_2018_12_14_17_27|

Cancel **Clone**

Click clone and wait for the process to finish.

The screenshot shows the SynetoOS interface. At the top, there is a search bar and a dropdown menu set to 'Show Originals, Clones, » Replica:'. The main area displays a table of volumes:

Name	Used space	Provisioning	Protection
Volume1	12.0 KB	thin	<input type="checkbox"/> M <input type="checkbox"/> W
<input type="checkbox"/> Volume1_2018_12_14_17_27	0.0 B	thin	NOT PROTECTED
Volume2	12.0 KB	thin	<input type="checkbox"/> M <input type="checkbox"/> H <input type="checkbox"/> D <input type="checkbox"/> W

Below the table, the details for 'Volume1_2018_12_14_17_27' are shown. It is a clone of 'Volume1' (2018/12/14 17:27). The size is 2.0 GB, and it is currently 'THICK' with a slider set to 'THIN'. The write back cache is 'OFF' with a slider set to 'ON'. The space effective used is 0.00 B. The description is '[2018-12-17 10:19:50] Volume Volume1_2018_12_14_17_27 created.' The iSCSI access is 'Volume not accessible via iSCSI'.

Now you can access a volume with the name specified at the clone step. It will contain the data as it was present at the moment of taking the snapshot.

WARNING: Do not change data on the clones. Clones cannot be protected, and after the Cleanup process (see below), all changes will be deleted.

Clean up when you are done

One of the more tedious tasks after recovering data is cleaning up after you. Deleting storage clones and volumes can be a tedious operation. And because it is time consuming and difficult, many people just let them there... forever.

With Syneto HYPER the whole cleanup process is just one click away. And we take care to clean up only what we created.

Go to *Main menu* → *Volumes* and select a share that has cloned snapshots. Click *Browse* in the snapshot section.

Search Show Originals, Clones, » Replica: ▼ 1-3 of 3

Name	Used space	Provisioning	Protection
Volume1	12.0 KB	thin	M W
📄 Volume1_2018_12_14_17_27	0.0 B	thin	NOT PROTECTED
Volume2	12.0 KB	thin	M H D W

Volume1 Delete Details Snapshots

Size	2.0 GB	Expand	Space effective used	<div style="width: 100%;"><div style="width: 6.00 KB;"></div></div> 6.00 KB
Block size	32 KB			<div style="width: 100%;"><div style="width: 12.00 KB;"></div></div> 12.00 KB
Provisioning	THICK <input checked="" type="radio"/> THIN			Compression 1.00x Saved 0.0 B
Write back cache	OFF <input checked="" type="radio"/> ON			
Protection	M W	Protect	Description	[2018-12-14 17:19:50] Volume Volume1 created. Edit
Clones			iSCSI access	Grant access
			Initiator	Target
			Initiators-Group-1	Target-Group-1
				LUN
				1

Name	Snapshot date
📄 Volume1_2018_12_14_17_27	2018/12/14 17:27

You will see that all cloned snapshots are marked with a label. Click the context menu for a cloned snapshot and select *Delete clone*.

Volume1 Delete Details Snapshots

December 2018

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6


December 14, 2018 3 snapshots

17:27	cloned	0.0 B
17:27:39	Starting snapshot process ...	
17:27:40	ISCSI Volume name: Volume1 ...	
17:27:41	Frequency: every friday ...	
17:27:41	Host: syneto-os-Dec1027e-adi ...	
17:27:42	Taking ISCSI Volume snapshot ...	
17:27:43	Local snapshot taken.	
17:26		0.0 B
17:25		0.0 B

Summary

Oldest snapshot on	2018/12/14
Total snapshots	3

The system will ask you to confirm the process.



Wait a few seconds for cleanup to finish. A green notification in the bottom right will tell you when it's done.

That's it, all that was created and shared for that snapshot clone was reverted. No more mess left behind by the recovery process.

DRIVES AND POOLS

Introduction

In the secondary menu, we have the *Hardware* section. This contains *Drives* and *Pools*. On the *Drives* page you can monitor the health of the drives and replace a faulted drive. On this page, you have two tabs that allows you to see the drives details, blink a drive, or replace a faulty one. The two tabs are *Map* and *List*.

Map tab will show you how the disks are mapped in the rack. By clicking on a disk, you will see it's details: *Name, Size, Pool, Enclosure, Position, Vendor, Model, Serial number* and *physical block size*.

The screenshot shows the 'Drives' page in SynetoOS. At the top, there are tabs for 'Map' and 'List', and a 'Rescan' button. The main content area is titled 'HYPER 3100 Series' and features a rack map with 'front' and 'rear' views. Below the map, the details for a specific drive are shown:

65535_0 ⚡ [Replace](#)

Size	111.8 GB	Vendor	ATA
Pool	OS	Model	INTEL SSDSC2BB120G6
Enclosure	HYPER 3100 Series / back	Serial	BTWA542500HQ120CGN
Position	row 1 / column 1	Physical block size	unknown

List tab will show you the drives in a table view. You can see all the details above and do all the actions like blinking and replacing drives. Here you can also *Search and Sort* by *Name, Size, Pool, Enclosure* and *Serial*.

The screenshot displays the 'Drives' management page in SynetoOS. At the top, there are tabs for 'Map' and 'List', and a 'Rescan' button. A search bar is located on the left, and '1-16 of 16' is shown on the right. The main table lists five drives. The third drive, 'c0t5000CCA255615D54d0', is highlighted. Below the table, the details for this drive are shown in a key-value format.

Name	Size	Pool	Enclosure	Serial
c0t5000CCA24D5B9984d0	5.5 TB	hybrid	Expansion Shelf 12 Bays	NCHMBNXS
c0t5000CCA24D5B9F54d0	5.5 TB	hybrid	Expansion Shelf 12 Bays	NCHMD1XS
c0t5000CCA255615D54d0	5.5 TB	hybrid	Expansion Shelf 12 Bays	K1HRJYSD
c0t5000CCA25562388Cd0	5.5 TB	hybrid	Expansion Shelf 12 Bays	K1HSOKAD
c0t500117310022358Cd0	93.2 GB	-	Expansion Shelf 12 Bays	FG008GSH

c0t5000CCA255615D54d0		Blink LED Replace	
Size	5.5 TB	Vendor	HGST
Pool	hybrid	Model	HUS726060AL5210
Enclosure	Expansion Shelf 12 Bays / front	Serial	K1HRJYSD
Position	row 1 / column 2	Physical block size	4096

Pools are used to group drives together in order to build redundancy, prevent data corruption and provide read cache and write acceleration.

On the *Pools* page you can monitor the health of the pools and do different actions related to pools. By clicking on a pool, more details will be shown.

The pool is automatically configured. It can be hybrid(HDD and SSD) or flash (only SSD). The hybrid pool on DR-Play and Hyper2000 Series uses an 8GB virtual disk for write acceleration and has no read cache.

You can graphically see how much space is being occupied on the pool. Pools with more than 80% of space usage might suffer fragmentation and slower access times.


The screenshot displays the 'Pools' management page. At the top, there are navigation links: 'New pool', 'Import existing', 'Integrity check', and 'Rescan'. Below this is a table with the following data:

Name	Size	Used space	Scheduled check
hybrid	1.8 TB	324.0 KB	in 4 days
OS	79.5 GB	5.5 GB	disabled

Below the table, the 'OS' pool is selected, showing a detailed view. It includes a 'Data' section with two bars representing 111.8 GB and 93.2 GB. A progress bar shows '5.5 GB of 79.5 GB used'. A 'Healthy' status indicator is present. Below this, an 'Integrity check' section shows 'Last check: never checked' and 'Scheduled check: disabled', with a 'Check now' link.

Replacing a faulted drive

Go to *Secondary Menu* → *Drives*.



The *drives* page will open. The faulted drive will be highlighted

- **yellow:** it has errors, but still works
- **red:** it does not work anymore

Name	Size	Used space	Scheduled check
! hybrid	3.6 TB	364.0 KB	in 4 days
✔ OS	79.5 GB	5.5 GB	disabled

hybrid Export Destroy Clear errors

Data 1.8 TB removed removed 1.8 TB Expand	364.0 KB of 3.6 TB used	 ! Degraded
Write acceleration Manage		
Read acceleration Manage		
Spares Manage		

Integrity check Check now

Last check never checked

Scheduled check in 4 days (Fri, Jul 20, 2018 22:59)

Name	Size	Used space	Scheduled check
✖ hybrid	3.6 TB	364.0 KB	in 4 days
✔ OS	79.5 GB	5.5 GB	disabled

hybrid Export Destroy

Data 1.8 TB 1.8 TB 1.8 TB 1.8 TB Expand	364.0 KB of 3.6 TB used	 ✖ Unavailable
Write acceleration Manage		
Read acceleration Manage		
Spares Manage		

Integrity check Check now

Last check never checked

Scheduled check in 4 days (Fri, Jul 20, 2018 22:59)

Click on a drive to be redirected to the Drives page. See the disk in List view → Replace the disk.

HYPER 3100 Series

c0t5000CCA25E498AF0d0 Blink LED [Replace](#)

Size	1.8 TB	Vendor	HGST
Pool	hybrid	Model	HUS726020AL5210
Enclosure	HYPER 3100 Series / front	Serial	K5H9EUJA
Position	row 2 / column 2	Physical block size	4096

Search 1—14 of 14

Name	Size	Pool	Enclosure	Serial
c0t5000CCA25E498AF0d0	1.8 TB	hybrid	HYPER 3100 Series	K5H9EUJA
c0t5000CCA25E498B48d0	1.8 TB	hybrid	HYPER 3100 Series	K5H9EV7A
c0t5000CCA25E498F78d0	1.8 TB	hybrid	HYPER 3100 Series	K5H9G3WA
c0t5000CCA25E499FB0d0	1.8 TB	hybrid	HYPER 3100 Series	K5H9H6BA
c0t5000CCA25E49A0D4d0	1.8 TB	-	unknown	K5H9H8RA

c0t5000CCA25E498AF0d0 Blink LED [Replace](#)

Size	1.8 TB	Vendor	HGST
Pool	hybrid	Model	HUS726020AL5210
Enclosure	HYPER 3100 Series / front	Serial	K5H9EUJA
Position	row 2 / column 2	Physical block size	4096

Search 1—14 of 14

Name	Size	Pool	Enclosure	Serial
● c0t5000CCA25E498AF0d0	1.8 TB	hybrid	HYPER 3100 Series	K5H9EUJA
● c0t5000CCA25E498B48d0	1.8 TB	hybrid	HYPER 3100 Series	K5H9EV7A
● c0t5000CCA25E498F78d0	-	hybrid	-	-
● c0t5000CCA25E499FB0d0	-	hybrid	-	-
● c0t5000CCA25E49A0D4d0	1.8 TB	-	unknown	K5H9H8RA

c0t5000CCA25E498F78d0 Removed Replace

Size	-	Vendor	-
Pool	hybrid	Model	-
Enclosure	-	Serial	-
Position	-	Physical block size	-

Our machines have support for blink, so you can detect the drive easily. To activate the blink for a drive, click on *the Drive -> Drive details -> Blink*.

NOTE: OS drives do not have support for blinking.

c0t5000CCA255615D54d0 Blink LED Replace

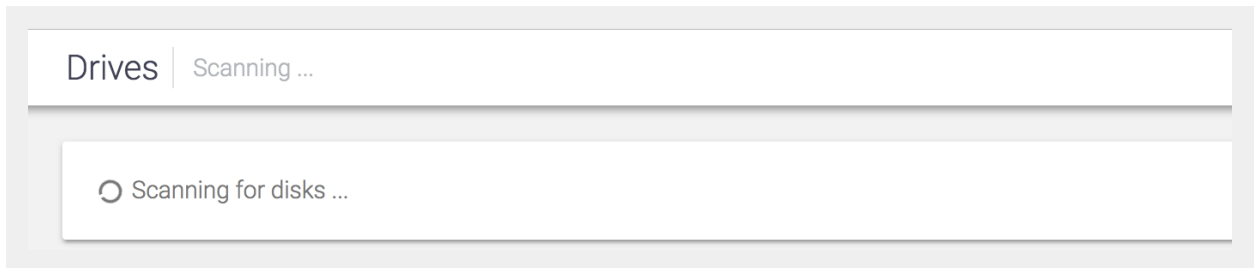
Size	5.5 TB	Vendor	HGST
Pool	hybrid	Model	HUS726060AL5210
Enclosure	Expansion Shelf 12 Bays / front	Serial	K1HRJYSD
Position	row 1 / column 2	Physical block size	4096

Go to the physical machine. Identify the faulted drive, take the bay out. Replace the drive in the bay with a new drive.

NOTE: Make sure that the replacement drive is of the same size as the one faulted.

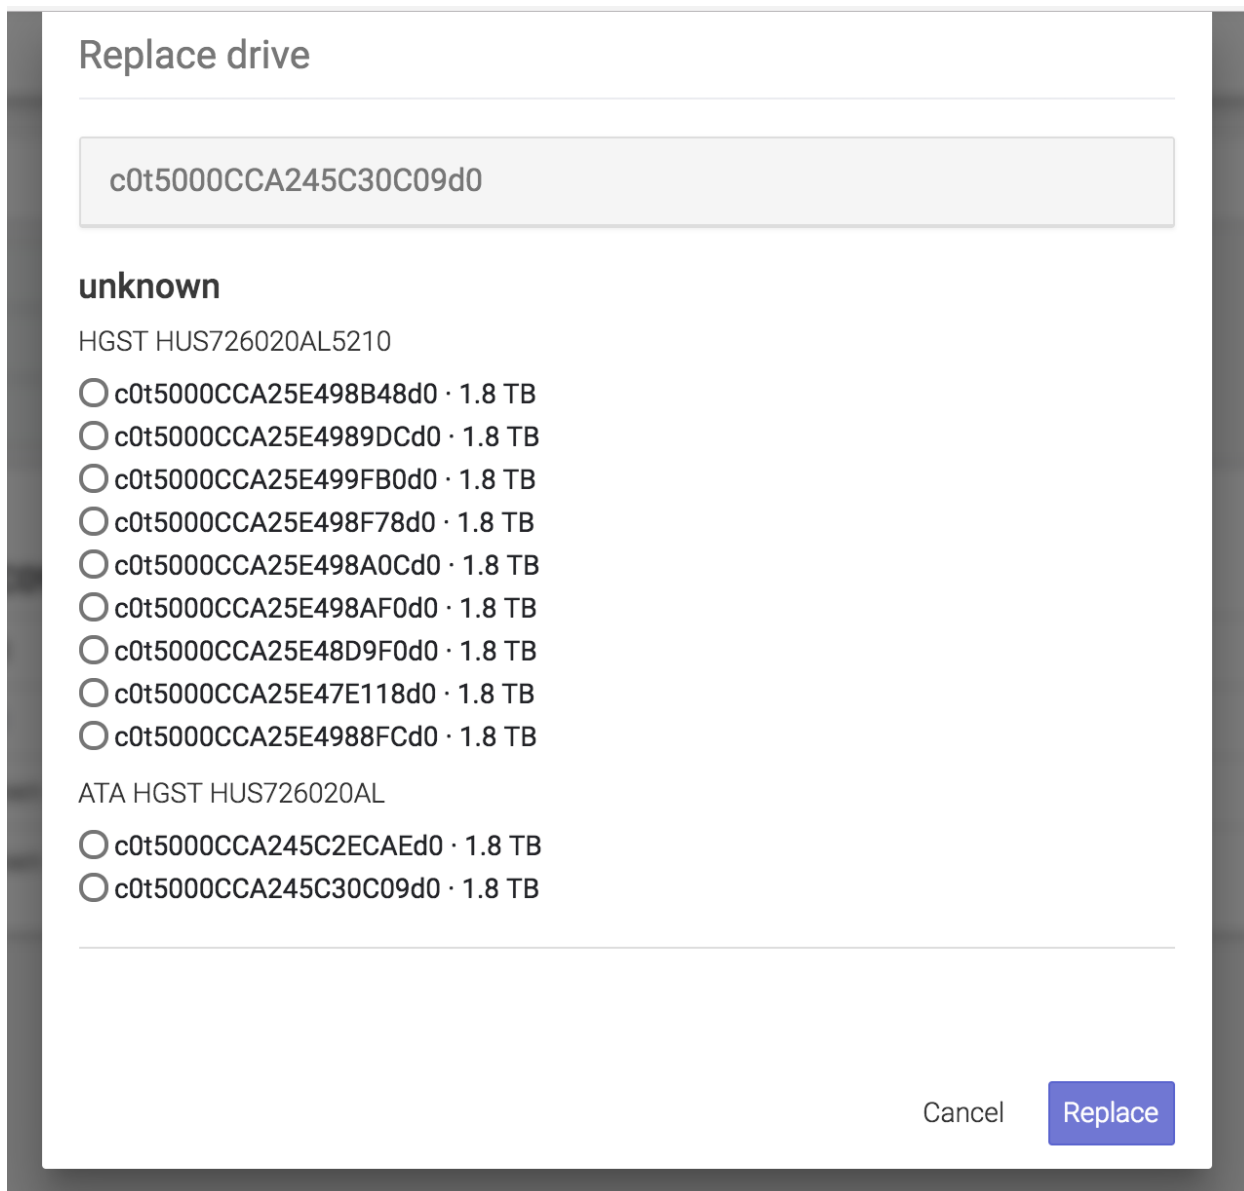
Place the bay in the machine.

Go to *Drives* page again. Click the *Rescan* button, found in the action bar.



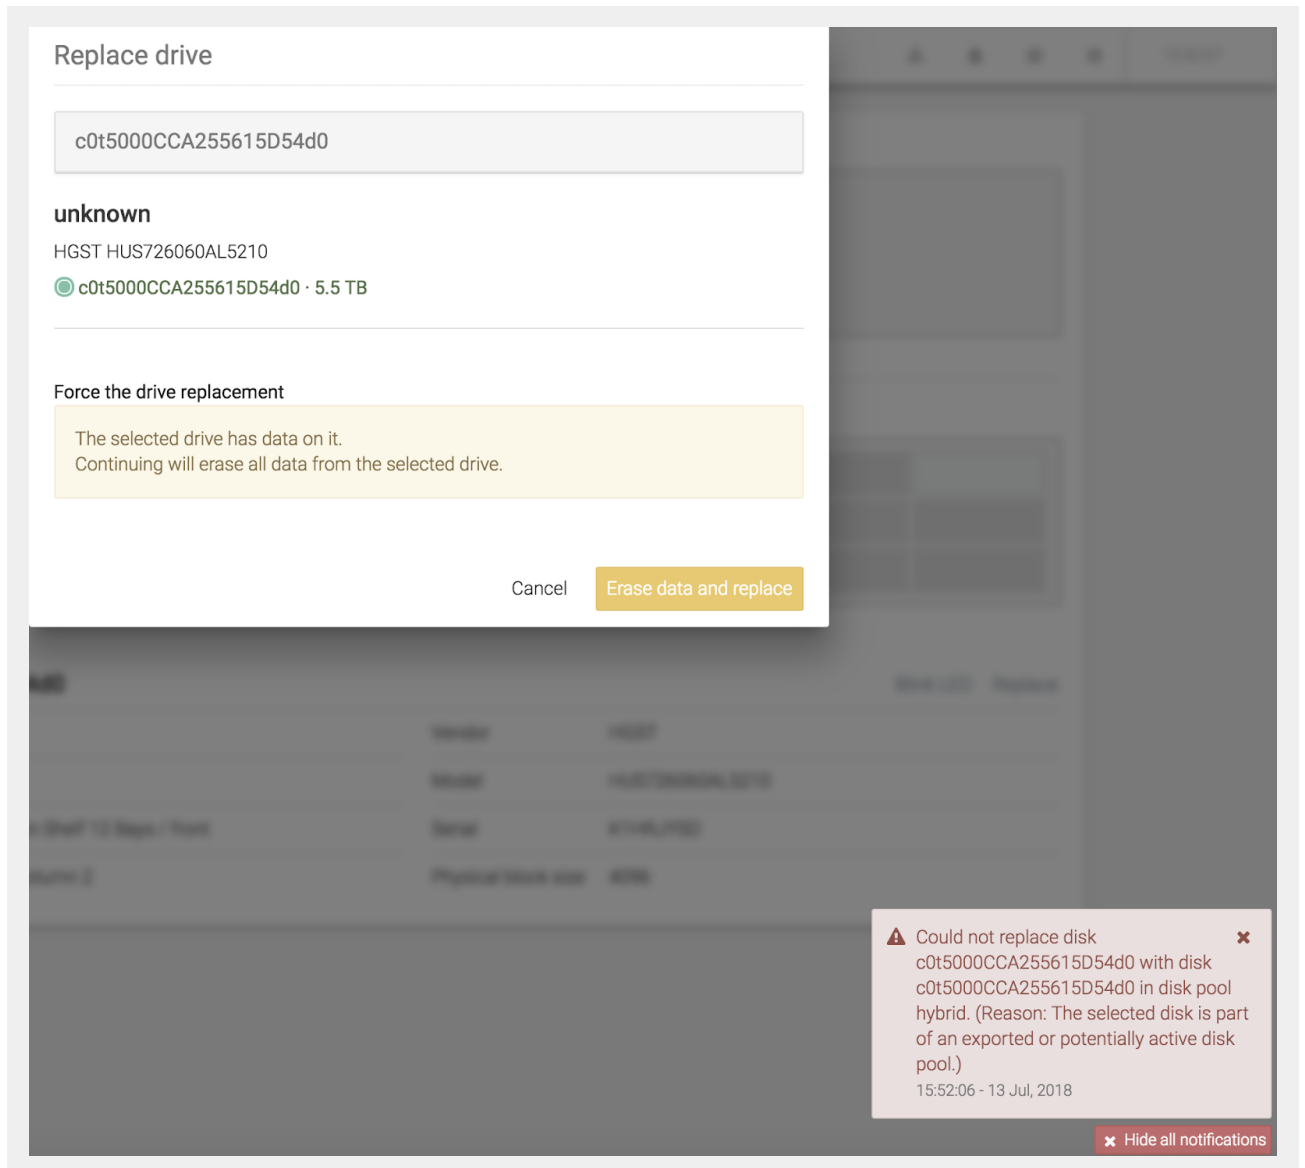
Click on *Drive details* → *Replace* .

c0t5000CCA255615D54d0		Blink LED Replace	
Size	5.5 TB	Vendor	HGST
Pool	hybrid	Model	HUS726060AL5210
Enclosure	Expansion Shelf 12 Bays / front	Serial	K1HRJYSD
Position	row 1 / column 2	Physical block size	4096



There will be a list with the available drives for replacement. Select the newly introduced one and click *Replace*. Now, the faulted drive is being replaced with the new drive. The system will automatically start restoring the data on the new drive.

If there was data on the newly introduced drive, a warning will appear.



WARNING: If you are sure there is no important data on the drive, click on Erase data and replace.

Replacing a faulted OS drive

Extract the faulted drive and replace it with a new drive.

The rest of the operation is done automatically by the hardware RAID.

Replacing a faulted cache drive

Go to *Secondary Menu* → *Pools* click on the context menu for the pool that has the faulted cache drive and click *Manage Cache*. Remove the faulted drive from cache and add the new drive that replaced it.

Creating a pool

Go to *Secondary Menu* → *Hardware* → *Pools*. The *Pools* page will open.

The screenshot displays the 'Pools' management interface. At the top, there is a navigation bar with the title 'Pools' and several action buttons: 'New pool', 'Import existing', 'Integrity check', and 'Rescan'. On the right side of the navigation bar, there are icons for user, notifications, help, and settings.

Name	Size	Used space	Scheduled check
hybrid	10.9 TB	14.4 GB	in 7 hours
OS	79.5 GB	5.2 GB	disabled

Below the table, there is a section for the selected 'OS' pool. It includes a 'Data' section with two storage indicators: '93.2 GB' and '111.8 GB'. To the right, there is a progress bar showing '5.2 GB of 79.5 GB used' and a 'Healthy' status indicator with a green checkmark. Below this, there is an 'Integrity check' section with a 'Check now' button. The 'Last check' is 'never checked' and the 'Scheduled check' is 'disabled'.

Click on the *New Pool* button in the action bar, top of the page. A dialog for creating a new pool will open.

Create pool

Name
|myData

Redundancy 0 1 2 3 Number of drives that can fail per group before data is lost.

Data drives group New data group

+

Spare drives group

+

Read cache drives group

+

Write acceleration drives group

+

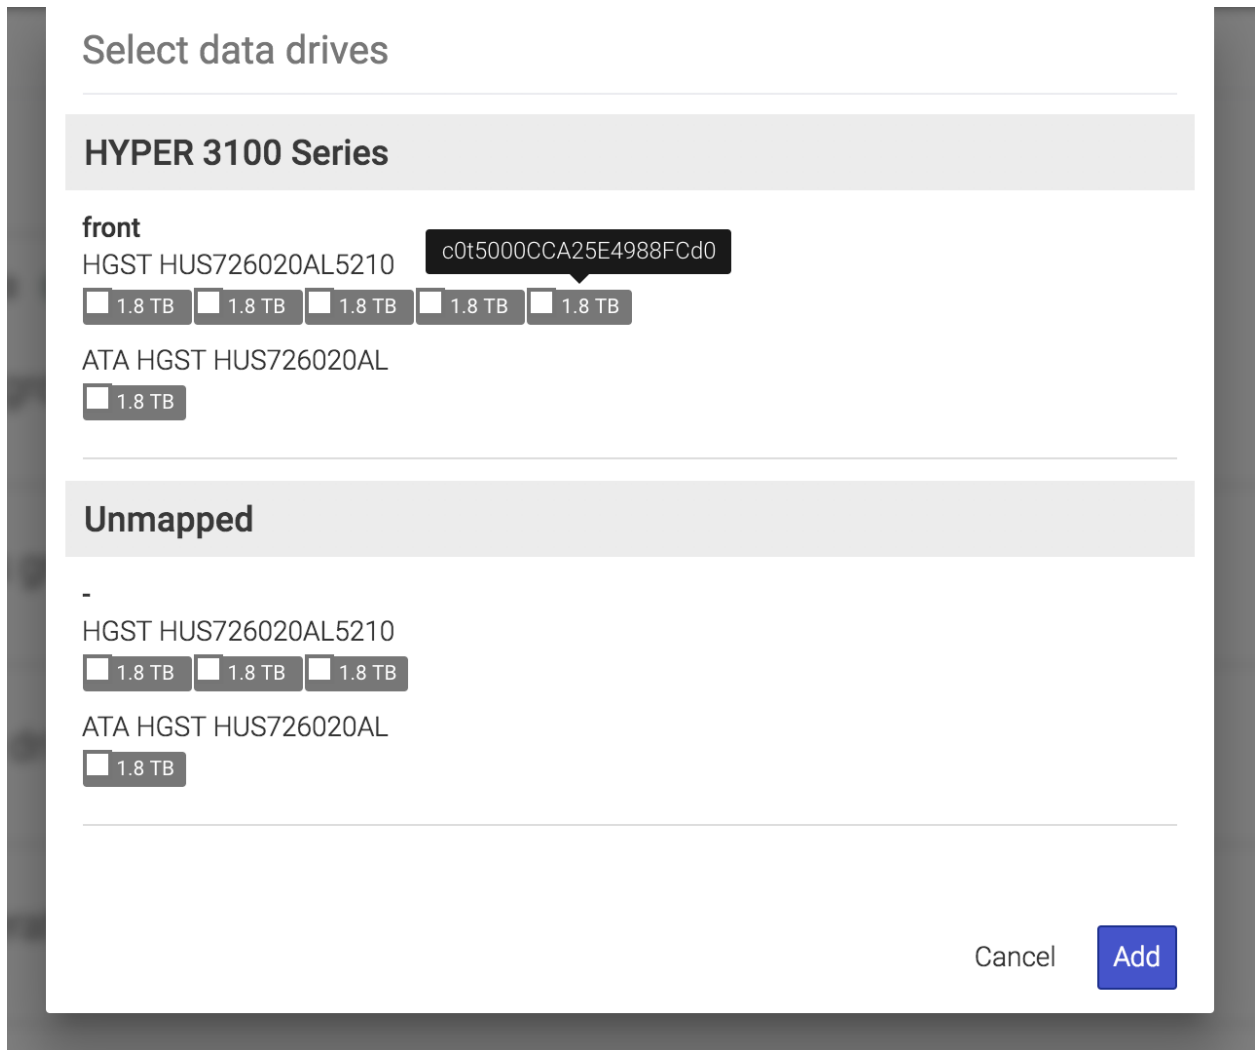
Cancel Create

You will have to give a name to the pool.

A pool is constructed of drive groups. There are several types of drive groups:

- *Data drive groups* – drives that contain the actual data. Each data drive group has a number of drives that may fail before losing data. Choosing data drive groups that will support drive failures will result in lower available space on the group – depending on the number of drives that may fail (e.g. if a data drive group of 5 drives will allow 2 drives to fail before losing data then the total available space on the drives group will be reduced by 40%).
Recommended: Select *redundancy* value 1 and two *drives* per *data group*.
- *Spare drive group* – the drives in this group are not actively used but will automatically replace a data drive that has failed. Adding spare drives is optional.
- *Read cache drive group* – this group contains drives that are actively used to speed up the data reading process. It is recommended to use high speed drives (SSD) in this group. Adding drives to this group is optional.
- *Write acceleration drive group* – this group contains drives that are actively used to speed up the data writing process. It is recommended to use only high speed drives (SSD) in this group. Adding drives to this group is optional.

To add a drive, click on the *Plus* button. A dialog for selecting a drive will open.



Create pool

Name
data

Redundancy 0 1 2 3 Number of drives that can fail per group before data is lost.

Data drives group [New data group](#)

1.8 TB 1.8 TB

Spare drives group

Read cache drives group

Write acceleration drives group

[Cancel](#) [Create](#)

Expanding a pool

Go to *Secondary Menu* → *Hardware* → *Pools*. Click *Expand* from the pool's details menu.

Name	Size	Used space	Scheduled check
✓ hybrid	1.8 TB	324.0 KB	in 4 days
✓ OS	79.5 GB	5.5 GB	disabled

hybrid Export Destroy

Data	1.8 TB 1.8 TB	Expand	<div>324.0 KB of 1.8 TB used</div> <div>Healthy</div> <div>Integrity check Check now</div> <div>Last check never checked</div> <div>Scheduled check in 4 days (Fri, Jul 20, 2018 22:59)</div>
Write acceleration		Manage	
Read acceleration		Manage	
Spares		Manage	

This will only expand the data drives.

For adding drives, click the *Plus* button



Expand pool

hybrid

Data drives group [New data group](#)

1.8 TB 1.8 TB

[+](#)

Cancel [Expand](#)

Select the ones you want by checking their boxes.

Select data drives

HYPER 3100 Series

front
HGST HUS726020AL5210
 1.8 TB 1.8 TB 1.8 TB 1.8 TB 1.8 TB

ATA HGST HUS726020AL
 1.8 TB

Unmapped

-
HGST HUS726020AL5210
 1.8 TB 1.8 TB 1.8 TB

ATA HGST HUS726020AL
 1.8 TB

Cancel

After you finished selecting, click *Add* then *Expand*.

NOTE: The new *data drive groups* must be the same size as the ones in the current pool.

To expand a pool with *spares drives*, click on *Details Menu* → *Spares* → *Manage*.


hybrid		
Data	1.8 TB 1.8 TB	Expand
Write acceleration		Manage
Read acceleration		Manage
Spares		Manage

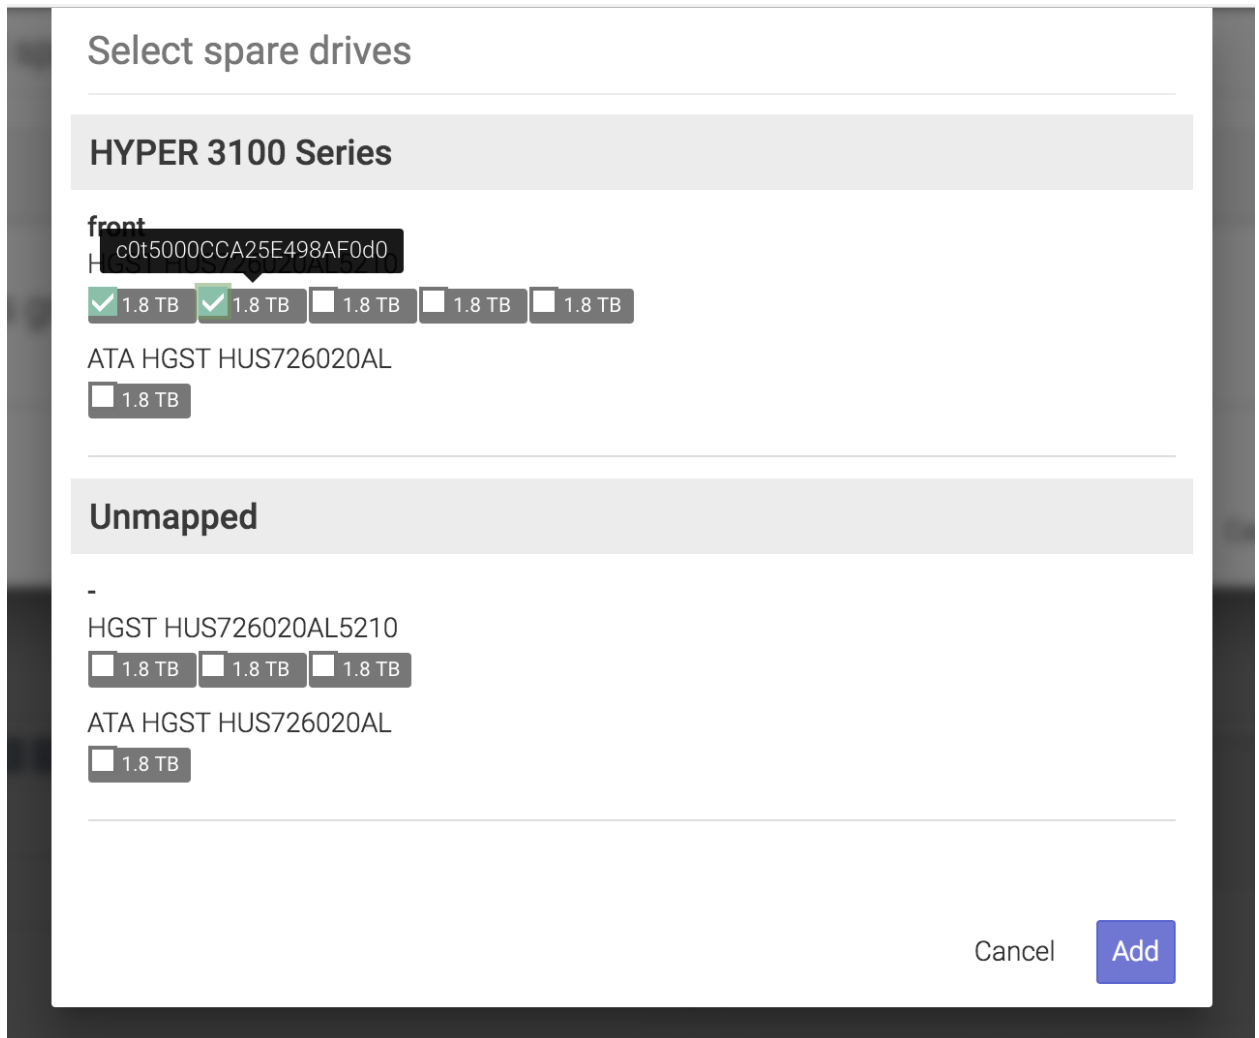
A dialog for adding *spare drives* will open.

Update pool spares

hybrid

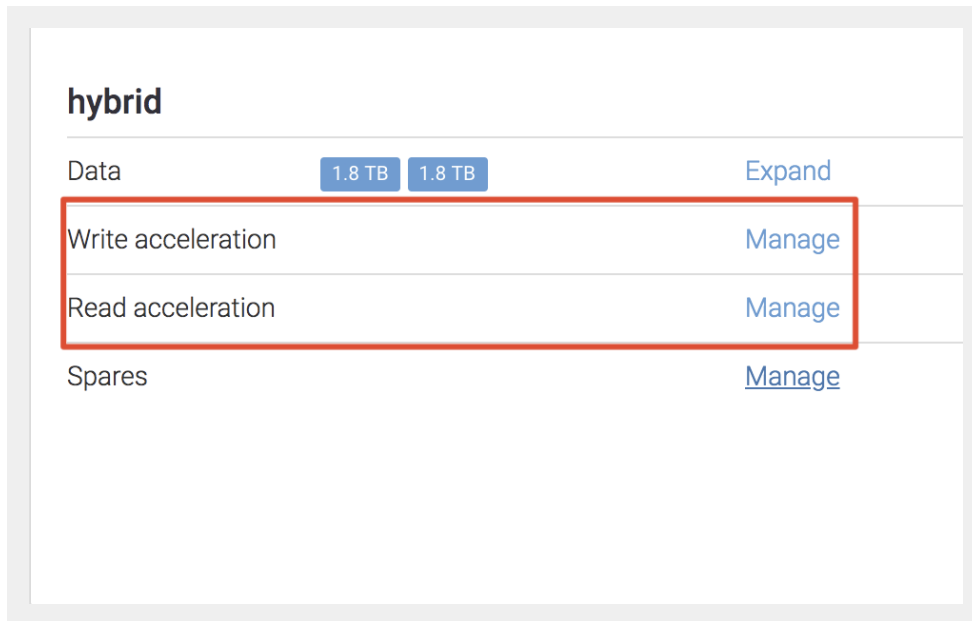
Spare drives group

Click on the *Plus* button . A dialog will open where you can select which drives to be used by checking their boxes.





After selecting the drives you want, click *Add* then *Update*.

To expand a pool with *cache drives*, click on *Details Menu* → *Read acceleration/Write acceleration* → *Manage*.



A dialog will open where you can add *read cache drives* or *write acceleration drives*.



Click on the *Plus* button . A dialog will open where you can select which drives to be used.

Select read cache drives

HYPER 3100 Series

front
HGST HUS726020AL5210
 1.8 TB 1.8 TB 1.8 TB 1.8 TB 1.8 TB

ATA HGST HUS726020AL
 1.8 TB

Unmapped

-
HGST HUS726020AL5210
 1.8 TB 1.8 TB 1.8 TB

ATA HGST HUS726020AL
 1.8 TB

Cancel


After you finish your selection, click *Add*.


Importing a disk pool

If you reinstall SynetoOS or you have introduced drives from another Syneto machine, you will want to access the data on them. This can be done by *importing the pool*.

Go to *Secondary Menu* → *Hardware* → *Pools*. The *Pools* page will open.


Click *Import existing* in the action bar. If there are drive pools that can be imported, a list will be shown.





Select the disk pool to be imported and click *Import*.

If a disk pool was previously used in another machine, a warning will appear. You will need to force import it.



Exporting a disk pool

If you reinstall SynetoOS or you want to move drives to another Syneto machine, you will want to be able to access the data. In order for the data to be saved properly, you need to *export the pool*.

Go to *Secondary Menu* → *Hardware* → *Pools*. The *Pools* page will open.

To export a pool, click on *Details Menu* → *Export*.

The screenshot shows the 'hybrid' disk pool configuration page. On the left, there are sections for 'Data', 'Write acceleration', 'Read acceleration', and 'Spares', each with '1.8 TB' indicators and 'Expand' or 'Manage' links. On the right, there is a progress bar showing '496.0 KB of 1.8 TB used' and a 'Healthy' status indicator with a green checkmark. Below this, there is an 'Integrity check' section with a 'Check now' link. The 'Export' button is highlighted with a red box in the top right corner.

NOTE: The OS pool can not be exported.

WARNING: All virtual machines and datastores related to this pool from all connected ESXi hosts will be removed.

Protecting against silent data corruption

Syneto products have a built-in mechanism for protecting against silent data corruption. Silent data corruption means that data being read is different from the data previously written to a drive. There are several ways this could happen, from controller firmware errors all the way to natural events like cosmic radiation.

Go to *Secondary Menu* → *Hardware* → *Pools*. The *Pools* page will open.

To set *integrity check*, click on *Integrity check*.

The screenshot shows the navigation bar for the 'Pools' page. It includes the following buttons: 'Pools', 'New pool', 'Import existing', 'Integrity check', and 'Rescan'. The 'Integrity check' button is highlighted with a red box.

Here you can set the speed and intervals for the integrity check of all pools.

Integrity check schedule

Start integrity check every weeks on Friday at 23:00.

This option affects the integrity checks on all the pools.
A faster speed may reduce the amount of IOPS available for read/write operations.

Set speed to slow normal fast

Cancel

Intervals section allows you to set a schedule. By default, it is enabled every week on Friday, at 23:00. You can change the week interval. First, you have to *Disable* the current schedule.

Write the new value you want in the field, between *every* and *weeks*. Click *Enable*.

Change speed section allows you set the speed of integrity check. It can be set to *slow*, *normal* and *fast*. Just select the one you want and SynetoOS will change it for future checks.

NOTE: A faster speed will consume more resources making the system perform slower.

WARNING: This option also affects the speed of rebuilding redundancy after a drive is changed.

You have an entire section, on each pool, with information about integrity checking: Last check, Schedule check and the option to *Check now*.

hybrid [Export](#) [Destroy](#)

Data	1.8 TB 1.8 TB	Expand
Write acceleration		Manage
Read acceleration		Manage
Spares		Manage

388.0 KB of 1.8 TB used

✓
Healthy

Integrity check [Check now](#)

Last check

Scheduled check in 4 days (Fri, Jul 20, 2018 22:59)

EXPANSION SHELVES

For most cases, the disks in a HYPER machine are enough to store all your data. However, there are IT infrastructures that require the archiving on large amount of historical data. For these cases, purchasing a solution with expansion shelves (JBODs) is a good idea.

If you already have a HYPER 3000 Series machine, you can purchase only the expansion shelves.

Attaching expansion shelves

NOTE: Hyper2000 series and DR-Play do not support expansion.

Connect the JBOD to the Syneto machine (for more information about physically connecting the JBOD, please consult the [JBOD Quick Start Guide](#)). Go on *Drives* page and click *Rescan*.

On the *Map* tab, you will see the enclosures for the JBOD(s) connected and their available drives. Available drives will appear in pale green. The drives which are part of a pool will be white showing a green light in their top left corner. Grey spaces suggest empty bays.

The screenshot displays the 'Drives' page in SynetoOS, showing a 'Map' view of the system's storage configuration. The interface includes a 'Rescan' button and 'Map' and 'List' tabs. The main area shows a 'HYPER 3100 Series' enclosure with a front view and a rear view. Below this, two 'Expansion Shelf 12 Bays' are shown. A specific drive is highlighted with a red box, and its details are shown in a table below.

65535_0 ⚡		Replace	
Size	111.8 GB	Vendor	ATA
Pool	OS	Model	INTEL SSDSC2BB120G6
Enclosure	HYPER 3100 Series / back	Serial	BTWA542500HQ120CGN
Position	row 1 / column 1	Physical block size	unknown


Expanding a pool with JBOD drives

Once an expansion shelf is connected, the system will suggest expanding current pools with the newly detected drives.

NOTE: You can still use the Drives and Pools page to configure a new pool in the traditional way.

When new drives are detected, a notification will appear. It will suggest you to "Use drives" and go through a wizard.

If you do not want to expand the pools, you can simply close the notification. Do not worry if you need it again, it will appear every time you land on Drives Page, until there are no more drives available for expanding a pool.



To expand the pool as the notification indicates, click on the link *Use drives*.

A new page will open, where you can see the pool which is about to be expanded and the size of the expansion.

Use new drives

2 additional drives have been detected in your system. You can increase the pool hybrid, by using them, with 2.6TB.

HYPER 3100 Series

front rear

Expansion Shelf 12 Bays **Expansion Shelf 12 Bays**

front front

Close Add

The pool will be expanded with drives of the same size as the ones already present on the data pool. The existing pool's redundancy policy will be used for all the newly added drives.

Remember that once you use the drives to expand a pool, you can not remove them unless you export or destroy the pool.

NOTE: Only data pools can be expanded.

DATA PROTECTION AND REPLICATION


Introduction

One of the main features of Syneto HYPER is the disaster recovery solution. Virtual machines, datastores, shares and volumes can all be protected and replicated to another Syneto machine. In case of a disaster your infrastructure can be up in 15 minutes.

Configuring replication targets

In case of a machine's failure, you will want to be able to recover from a backup. For this to be possible you will first need to replicate your data to the DR unit. The communication details for replication to the DR unit(s) are kept in the *replication targets* list.

Adding a replication target




The screenshot shows the 'Protect and replicate' page in SynetoOS. The 'Replication targets' tab is active, displaying a table of replication targets. The table has the following data:

Name	Type	Last snapshot	Last replication	Protection
documents	SMB	never	never	NOT PROTECTED
time_machine	AFP	never	never	NOT PROTECTED

Below the table, the 'documents' target is selected, and there are buttons for 'Progress', 'Snapshots', and 'Schedules'. A message below the buttons states: 'No snapshots have been taken yet.'


On the *Protect and replicate* page, click *Replication targets* in the action bar. A dialog will open, listing the available replication targets.



Click *Add target*. Another dialog will open where you can enter the configuration for the *new replication target*. In the field *name* write the name you want to give to the target. In the field *Hostname / IP* write the hostname or IP of the replication target. In the *Password* field write the admin's password.

The screenshot shows a dialog box titled "New replication target". It contains three input fields: "Target name" with the value "DR-target" and a help icon; "Hostname / IP" with the value "192.168.3.23"; and "Password" with five dots. At the bottom, there are "Cancel" and "Add" buttons.

After filling the form and clicking *Add*, the list will be updated with the replication target.



WARNING: It will fail when adding as replication target a HYPERseries machine with an OS below ver. 3.2.8.

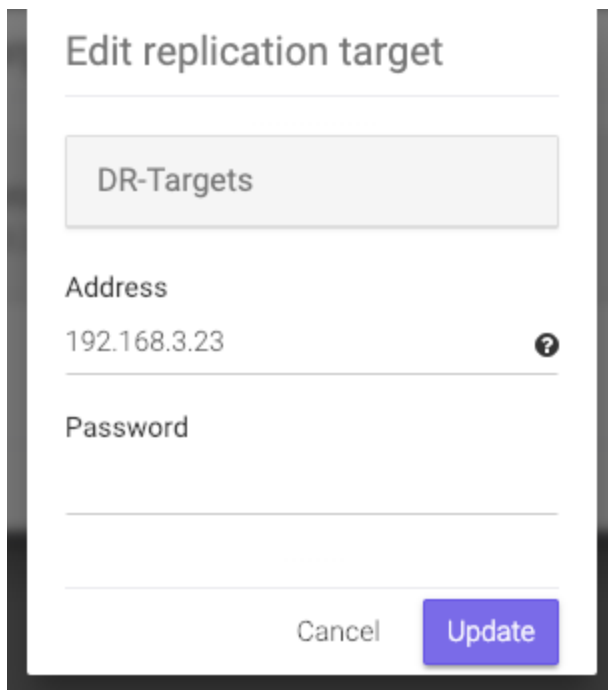
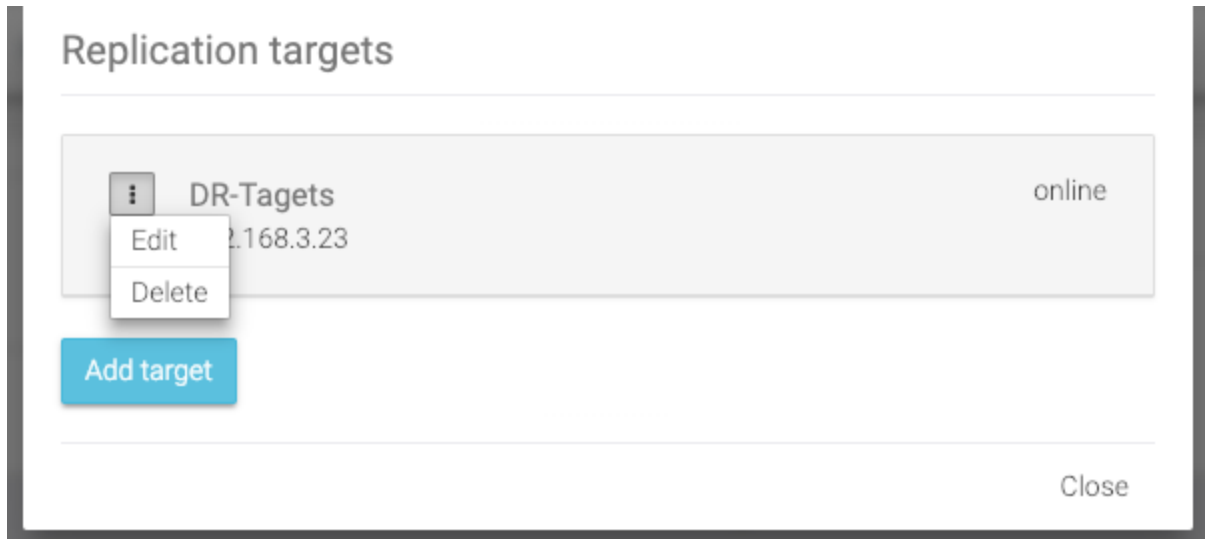
Removing a replication target

On the *Protect and replicate* page, click *Replication targets* in the action bar. A dialog will open listing the available replication targets. For the replication target you want to remove, click on *Context Menu* → *Delete*.

WARNING: Deleting a replication target, will also remove it from the snapshot schedules.

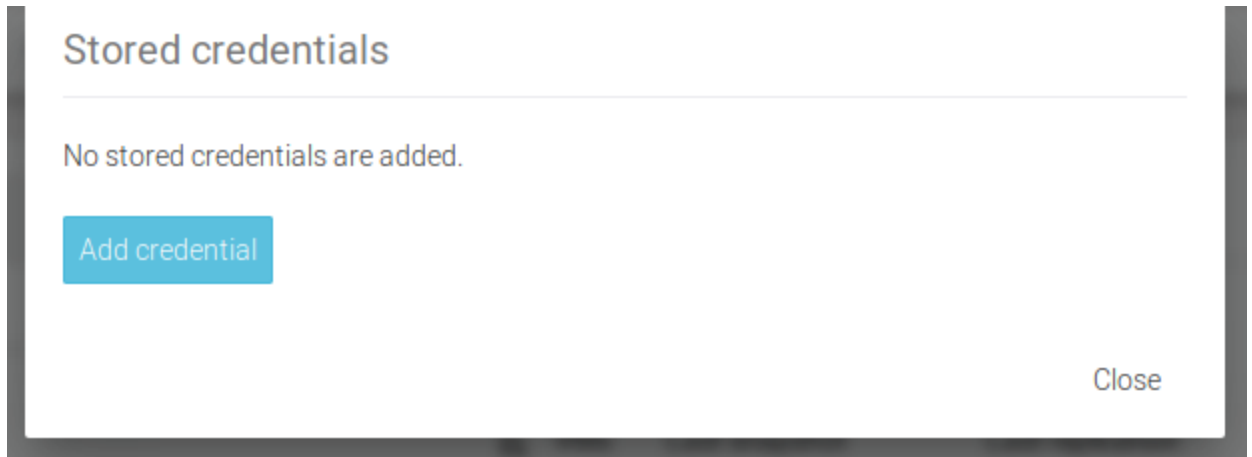
Editing a replication target

On the *Protect and replicate* page, click *Replication targets* in the action bar. A dialog will open listing the available replication targets. For the replication target you want to edit, click on *Context Menu* → *Edit*.



Configuring virtual machine credentials

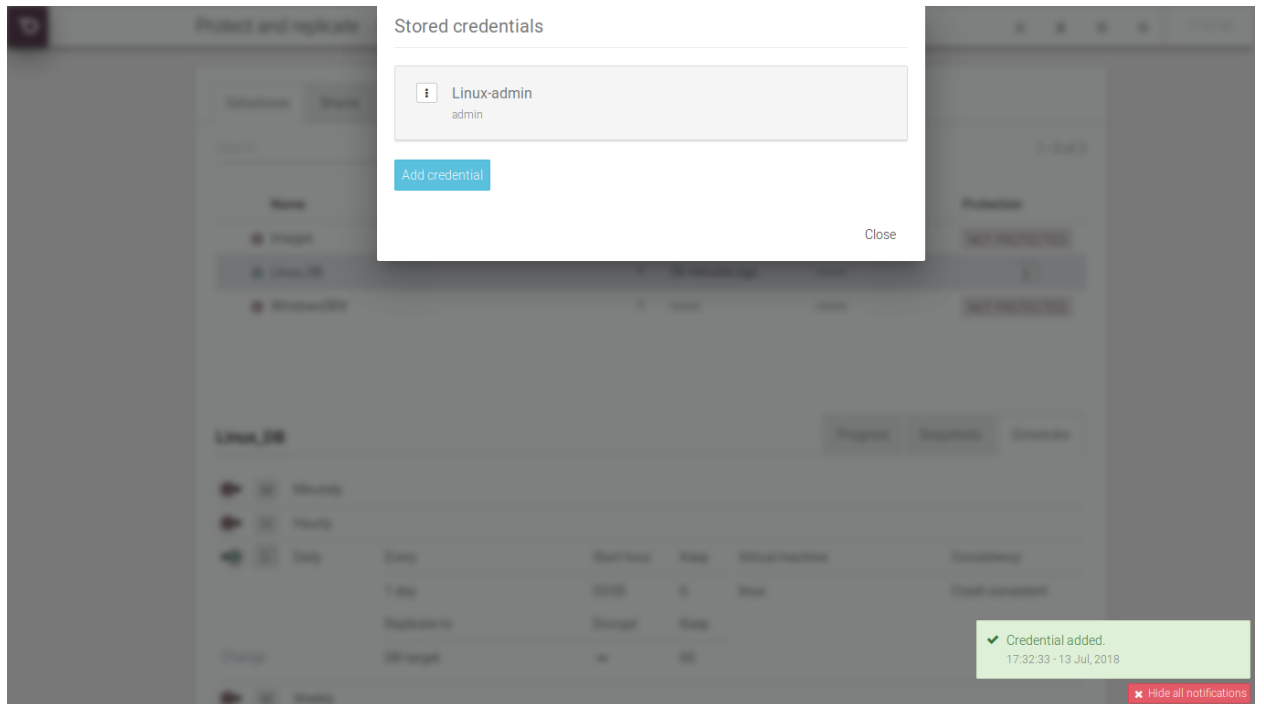
To be able to run pre-freeze and post-thaw scripts on a virtual machine during the application consistent snapshotting process, you will need to add the virtual machine's credentials. On the *Protect and replicate* page, click on *Stored credentials* in the action bar. A dialog will open listing the available credentials.



To add a credential, click on *Add credential*. A dialog will open where you can enter a name by which to identify it, and the *Username* and *Password* which will be used to log in to the virtual machine.

The screenshot shows a dialog box titled "New credential". It contains three input fields: "Credential name" with the value "Linux-admin", "Username" with the value "admin", and "Password" with five dots. Each input field has a question mark icon to its right. At the bottom, there are two buttons: "Cancel" and "Add".

Click *Add*. The credentials will appear in the list.



Protecting a datastore

Note: Cloned datastores cannot be protected because they are temporary objects.

The *Protect and replicate* page contains 3 tabs: *Datastores*, *Shares* and *Volumes*. Select *Datastores*.

Protect and replicate | [Replication targets](#) | [Stored credentials](#) | 17:51:50

Datstores | Shares | Volumes

Search Show All 1-3 of 3

Name	VMs	Last snapshot	Last replication	Protection
Images	0	never	never	NOT PROTECTED
Linux_DB	1	never	never	NOT PROTECTED
WindowsSRV	1	never	never	NOT PROTECTED

WindowsSRV | Progress | Snapshots | Schedules

No snapshots have been taken yet.

The table lists all the datastores. There is a *Search* bar, useful for searching for a specific datastore by specifying its name.

Protect and replicate | [Replication targets](#) | [Stored credentials](#) | 17:52:40

Datstores | Shares | Volumes

Win Show All 1-1 of 1 (filtered from 3 total entries)

Name	VMs	Last snapshot	Last replication	Protection
WindowsSRV	1	never	never	NOT PROTECTED


WindowsSRV | Progress | Snapshots | Schedules

No snapshots have been taken yet.

If you click on a datastore, more details about its protection will be shown.

Configuring the snapshot schedule for a datastore

In the details section, there is *Schedules tab*. Click on it. On the left are the four frequencies a user can choose: *Minutely*, *Hourly*, *Daily* and *Weekly*. You can enable one, two, three, or all four frequencies for a datastore.



The screenshot shows the SynetoOS interface for configuring snapshot schedules. The main view is titled "Protect and replicate" and includes tabs for "Replication targets" and "Stored credentials". The "Datastores" tab is active, displaying a table of datastores:

Name	VMs	Last snapshot	Last replication	Protection
Images	0	never	never	NOT PROTECTED
Linux_DB	1	never	never	NOT PROTECTED
WindowsSRV	1	never	never	NOT PROTECTED

Below the table, the "Linux_DB" datastore is selected, and the "Schedules" tab is active. The "Schedules" section shows four frequency options, each with a toggle switch:

- Minutely
- Hourly
- Daily
- Weekly

To enable one schedule type, click on the switch on the left of the desired interval. A dialog will appear, where you can configure the snapshot policy for the selected datastore and frequency.

Daily schedule

Linux_DB

Protection

Every 1 days	Start hour 03:00	Keep 6
--	----------------------------	------------------

Replicate to
 Don't replicate ▼

Consistency

Virtual machine linux	Protection level Crash consistent ▼
---------------------------------	--

Extra space used for live snapshots: **0.0 B**

Cancel Save

Here you can specify how often you want the schedule to perform by changing the *Every* field.

NOTE: We recommend spacing out snapshots taking time with at least 5 minutes. If all schedules run at the same time, and then replicate, it can lead to high network traffic and high drive IOPS on the replication target.

To the right of the dialog, there is the *Keep* field. Here you can specify how many snapshots should be kept on the Syneto HYPER. If more snapshots accumulate, older ones will be deleted.

NOTE: The retention policy applies to a type of interval only and for the selected datastore.

Replicate to section allows you to select to which replication target to replicate. By default, no replication is set. If you click on the *Target* field, a list of the defined replication targets will be shown.

Daily schedule

Linux_DB

Protection

Every	Start hour	Keep
1 days	03:00	6

Replicate to

Don't replicate

- Don't replicate
- DR-target

Consistency

Virtual machine	Protection level
linux	Crash consistent

Extra space used for live snapshots: **0.0 B**

Cancel Save

After you select the desired replication target, two more fields will appear.

Daily schedule

Linux_DB

Protection

<p>Every</p> <p>1 <input type="text"/> days</p>	<p>Start hour</p> <p>03:00</p>	<p>Keep</p> <p>6</p>
<p>Replicate to</p> <p>DR-target</p>	<p>Encrypt</p> <p><input type="checkbox"/></p>	<p>Keep</p> <p>60</p>

Consistency

<p>Virtual machine</p> <p>linux</p>	<p>Protection level</p> <p>Crash consistent</p>
-------------------------------------	---

Extra space used for live snapshots: **0.0 B**

Cancel Save

If you want the replication to be encrypted, click on the box below the *Encrypt* field.

NOTE: Non-encrypted replication is faster than encrypted one, but less secure. If the replication target is off-site, we recommend encryption.

You can specify how many replicated snapshots should be kept on the replication target under the *Keep* field.

NOTE: The retention policy applies to a type of interval only and for the selected datastore.

NOTE: We recommend a lower retention policy on the primary unit and a higher one on the replication target (DR unit).

After you finished all the configuration, click *Save*. Your schedules will appear in the details section, in *Schedules*. You can edit a schedule by clicking on *Change*.

The screenshot shows the 'Protect and replicate' interface with the 'Replication targets' tab selected. A table lists three datastores: 'Images', 'Linux_DB', and 'WindowsSRV'. The 'Linux_DB' row is highlighted, and its details are expanded below. The details section shows a 'Daily' schedule with a frequency of 'every 1 day at 03:00'. The replication target is 'DR-target, not encrypted' and the consistency is 'Crash consistent'. The retention policy is 'keep 6 locally and 60 remotely'. There are 'Change', 'Progress', 'Snapshots', and 'Schedules' buttons visible.

Name	VMs	Last snapshot	Last replication	Protection
Images	0	never	never	NOT PROTECTED
Linux_DB	1	never	never	D
WindowsSRV	1	never	never	NOT PROTECTED

Linux_DB Progress Snapshots Schedules

M Minutely
 H Hourly
 D Daily

Schedule: every 1 day at 03:00
Replication: DR-target, not encrypted
Retention policy: keep 6 locally and 60 remotely

W Weekly

A schedule can be paused by clicking on the *green switch* button. This will not pause a snapshot or replication which is currently in progress. If you want to resume it, just click on the *red switch* button.

If you want to delete a schedule, click on *Change*. A dialog will open. Click on *Delete schedule*.

Daily schedule

Linux_DB

Protection

Every	Start hour	Keep
1 days	03:00	6
Replicate to	Encrypt	Keep
DR-target	<input type="checkbox"/>	60

Consistency


Virtual machine	Protection level
linux	Crash consistent

Extra space used for live snapshots: **0.0 B**

[Delete schedule](#) Cancel Save

Checking the progress of a snapshot or replication of a datastore

If you want to check the progress of a snapshot schedule, go to *Progress* tab in details section.



Protect and replicate | [Replication targets](#) | [Stored credentials](#) | 17:05:59

Datstores | Shares | Volumes

Search Show: All 1-3 of 3


Name	VMs	Last snapshot	Last replication	Protection
Images	0	never	never	NOT PROTECTED
Linux_DB	1	never	never	D
WindowsSRV	1	never	never	NOT PROTECTED

Linux_DB | [Progress](#) | [Snapshots](#) | [Schedules](#)

D never run
Target: DR-target [Run now](#)

You will see for every enabled schedule the current state. You can click *Run now* if you want to start the schedule in this moment.

NOTE: First replication will always copy all the data found in the datastore. Subsequent replications will transfer only the differences between the last successful replication and current snapshot.



Protect and replicate | [Replication targets](#) | [Stored credentials](#) | 17:06:57

Datstores | Shares | Volumes

Search Show: All 1-3 of 3

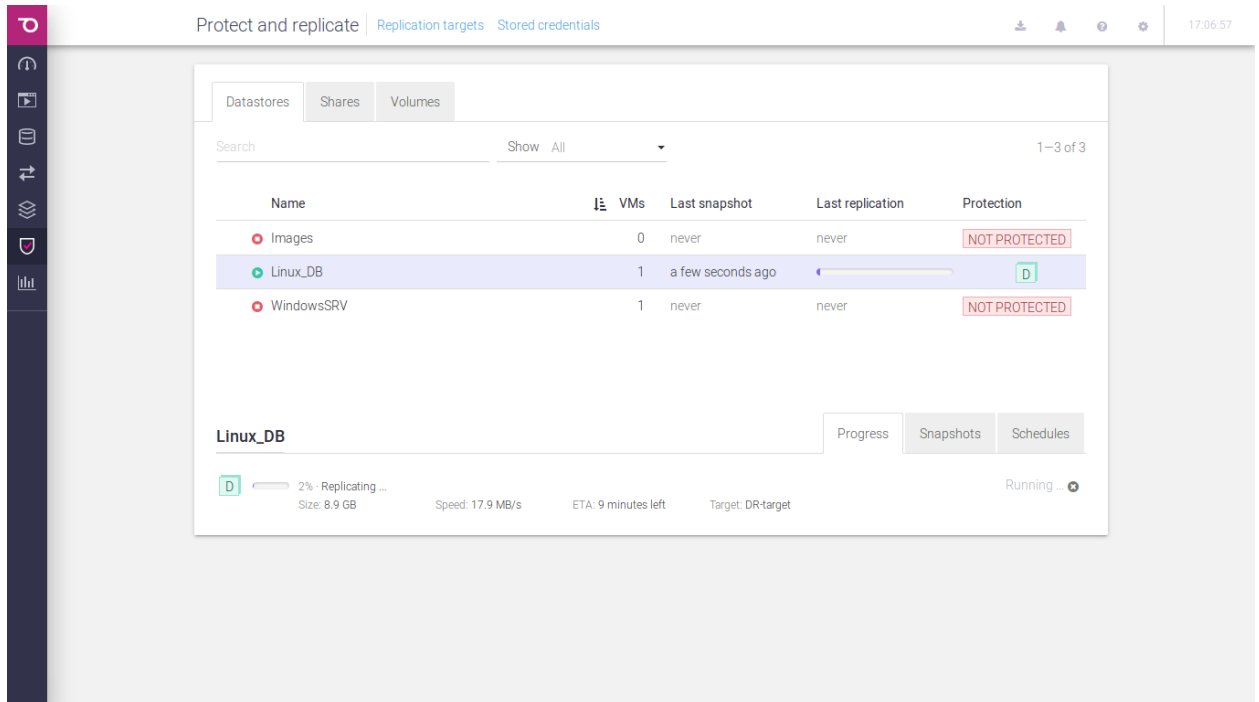
Name	VMs	Last snapshot	Last replication	Protection
Images	0	never	never	NOT PROTECTED
Linux_DB	1	a few seconds ago	<div style="width: 100px; height: 10px; background: linear-gradient(to right, #007bff, #6c757d);"></div>	D
WindowsSRV	1	never	never	NOT PROTECTED

Linux_DB | [Progress](#) | [Snapshots](#) | [Schedules](#)


D 2% Replicating ...
Size: 8.9 GB | Speed: 17.9 MB/s | ETA: 9 minutes left | Target: DR-target Running ...

Canceling a running replication of a datastore

A replication that is currently running can be canceled. As a consequence, the replica will not be created on the destination.



After you click on the cancel replication icon, you will receive a confirmation saying that the Snapshot replication was cancelled.



Monitoring the snapshots of a datastore

If you want to check the list of snapshots of a datastore, go to *Snapshots* tab in details section.

The screenshot shows the 'Protect and replicate' interface. At the top, there are tabs for 'Replication targets' and 'Stored credentials'. Below this is a table of replication targets:

Name	VMs	Last snapshot	Last replication	Protection
EmailServer-Linux	0	never	never	NOT PROTECTED
OracleServer	0	never	never	NOT PROTECTED
Windows10	1		never	M D
WindowsServer	2	3 days ago	never	M H

Below the table, there is a section for 'Windows10' with tabs for 'Progress', 'Snapshots', and 'Schedules'. The 'Snapshots' tab is active, showing a calendar for December 2018. The date December 12, 2018, is selected, and a list of snapshots is displayed:

Time	Size
15:23	140.3 MB
15:18	37.8 MB
14:33	798.6 MB

On the left, there is a calendar. On the right there is a list of snapshots from the date that is selected. You can change the day of the month by clicking on the desired day or you can go to another month by clicking the left or right arrows at the top of the calendar. Only dates with snapshots will be available for selection and the current date.

To check the history of a snapshot, click on it in the right section.

The screenshot shows the 'Protect and replicate' interface with the 'Snapshots' tab for 'Windows10' selected. The calendar shows December 12, 2018, selected. The list of snapshots is expanded to show the history of the 15:23 snapshot:

Time	Size	Details
15:23	140.3 MB	<ul style="list-style-type: none"> 15:23:41 Starting snapshot process ... 15:23:42 VMware Datastore name: Windows10 ... 15:23:42 Frequency: every 1 days ... 15:23:43 Host: syneto-os-0ec1027e-adi ... 15:23:45 Taking datastore snapshot ... 15:23:46 Virtual machines: 1 15:23:46 Windows 10 x64 (crash consistent) 15:23:47 Local snapshot taken.
15:18	37.8 MB	
14:33	798.6 MB	

Below the snapshots list, there is a 'Summary' section:

Summary

Oldest snapshot on 2018/12/12

Total snapshots 6 [Purge](#)

Configuring the type of snapshots for the virtual machines

If a datastore has virtual machines on it, in the snapshot schedule dialog there will be a list of virtual machines and their *protection level*.

Daily schedule

Linux_DB

Protection

Every 1 days Start hour 03:00 Keep 6

Replicate to DR-target Encrypt Keep 60

Consistency

Virtual machine	Protection level
linux	Crash consistent

Extra space used for live snapshots: 0.0 B

Delete schedule

There are three main types of protection levels for a virtual machine:

- *Crash consistent* - When restoring, the virtual machine will be in a state similar to a restart after a power outage. **Recommended** for virtual machines resilient to forced reboots. (ie. Linux or Solaris servers or workstations)

- *Application consistent*
 - *Windows VSS* - Before taking the snapshot a message will be sent to the operating system. This will instruct compatible applications to flush all their data to the drive. The snapshot will be taken after the flush operation finishes. Drive consistency for the application supporting the protocol will be consistent. When restoring, the virtual machine will be in a state similar with a power off.
Recommended for Microsoft servers (Active Directory, MS-SQL, etc).
 - *Hooks only* - Before taking the snapshot, the pre-freeze hook will be executed. After the snapshot of the virtual machine is done, the post-thaw hook will be executed. For information on how to configure this type of schedule go to [configuring snapshot hooks](#).
Recommended for Linux Databases and Windows databases that do not have VSS integration (Oracle, MySQL, etc).
 - *Windows VSS & Hooks* - Executes both Windows VSS and hooks.
- *Live snapshot* - Before taking a snapshot all operations on the virtual machine will be suspended and saved to the drive. This includes waiting for drive flush operations, persisting RAM memory, persisting current running state. Unfortunately, during this operation the virtual machine is not accessible. The process can take up to 5 minutes or more for large servers.
Recommended for mission critical system and only on a daily or rarer schedule, run overnight, so that work is not interrupted.

By default, the protection level of a virtual machine is set to *crash-consistent*. If you want to change it, click on the right of the virtual machine name, in *Protection level* column. A dropdown will open where you can select the protection level you want for the virtual machine.

It will contain a list with all the virtual machines contained by the selected datastore. Each virtual machine can have one of the three available protection levels: *Crash consistent*, *App consistent* and *Live snapshot*. Select your desired options and then click on *Update*.

WARNING: Live snapshot is not available on Syneto HYPER 2000 Series.

The screenshot shows the 'Protect and replicate' interface with the following data:

Name	VMs	Last snapshot	Last replication	Protection
Images	0	never	never	NOT PROTECTED
Linux_DB	1	33 minutes ago	never	D
WindowsSRV	1	never	never	NOT PROTECTED

Linux_DB details:

- M Minutely
- H Hourly
- D Daily

Every	Start hour	Keep	Virtual machine	Consistency
1 day	03:00	6	linux	Live snapshot
Replicate to	Encrypt	Keep		
DR-target	-	60		
- W Weekly

The details section will be updated with the new information. The following snapshots will take the virtual machine snapshots with the new configuration.

NOTE: For the minutely schedule smaller than 15 minutes, the protection level will always be crash consistent.

Configuring snapshot hooks for a virtual machine

If you want to make a *hooks only snapshot* of a virtual machine, select the datastore on which is mounted on. Select the schedule type on which you would like to apply *hooks only*.

Daily schedule

Linux_DB

Protection

Every	Start hour	Keep
1 days	03:00	6

Replicate to

Don't replicate

Consistency

Virtual machine	Protection level
linux	Crash consistent

Extra space used for live snapshots: **0.0 B**

Cancel Save

To select the *Protection level* of a virtual machine, click on the right of the virtual machine name in the *Protection level* column. A list will appear with the available protection levels.

NOTE: To be able to set any of the *application consistent* protection levels, you need to have *VMware Tools* installed on the virtual machine.

Daily schedule

Linux_DB

Protection

Every	Start hour	Keep
1 days	03:00	6
Replicate to	Encrypt	Keep
DR-target	<input type="checkbox"/>	60

Consistency

Virtual machine	Protection level
linux	Crash consistent

Extra space used for live snapshots: **0.0 B**

[Delete schedule](#)

- Crash consistent
- Application consistent
- Windows VSS
- Windows VSS and hooks
- Hooks only
- Live snapshot (RAM 1.0 GB)

Select *Hooks only* or *Windows VSS and hooks*.

Daily schedule

ENABLE

Protection

Every	Start hour	Keep
1 days	03:00	6

Replicate to

Don't replicate

Consistency

Virtual machine	Protection level
linux	Hooks only

Pre-freeze script

/home/user/pre-hook.sh

Post-thaw script

/home/user/post-hook.sh

Credential

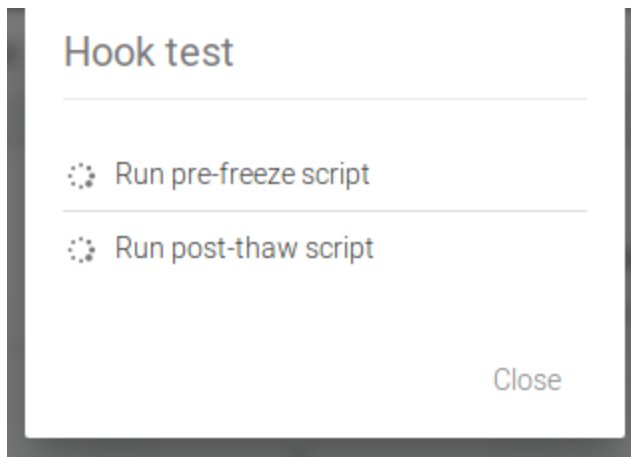
Linux-admin

Test


Extra space used for live snapshots: **0.0 B**

Delete schedule Cancel Save

You have to add the path of the *pre-freeze script* and *post-thaw script*. You also have to provide *credentials* which can connect and execute the scripts. To see how to add credentials, check [configuring virtual machine credentials](#). If you want to test the hooks configuration, click on *Test* button. This will only execute the scripts, not take a snapshot.

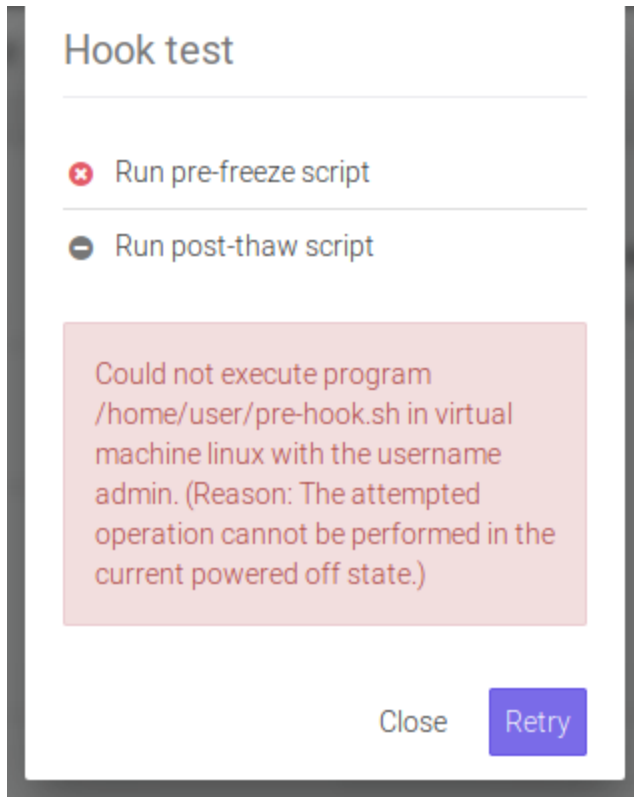


If the scripts are executed successfully, a green check will appear.



If one of the hooks fail, an error message will be displayed with the failure reason.

Note: If the hooks fail during a snapshot schedule run, the virtual machine snapshot will be crash consistent.



Note: Hooks must execute in less than 30 seconds. If they take longer than that, SynetoOS will kill the process.

Exceptions: Excluding time intervals from Schedules

The *Minutely* and *Hourly* schedules allow you to exclude time intervals - ranges of hours, throughout the days of the week. Snapshotting & replication will not happen in those time intervals.

Excluding time intervals can be useful if you want to backup your data at high granularity, but you know that one time or more during the week there is an important operation that your system needs to perform. If that operation could interfere with the scheduled backup, or is computationally expensive (e.g. intensive disk IO), it might be better to simply exclude that time interval from the backup schedule.

In order to use exceptions in the backup schedule, look for the *Exceptions* field in the *Minutely / Hourly* schedule, and click the link (initially stating *No exceptions*). A green coloured table will appear, with the days of the week as the row and hours of the day as the columns.

In order to exclude intervals from the schedule, click and drag, so as to draw rectangles in the green area: grey blocks represent hours in the week when the backup will not run, while green blocks represent hours in the week when the backups will be allowed to run.

NOTE: You can also draw rectangles in grey areas - this will cause blocks within the grey area to turn back green (and thus allow the backup to run in those hours).

NOTE: You can also drag within a single block - this will cause that block (hour in the week) to toggle from being excluded (grey) to not being excluded (green), and vice versa.


This is how you use, click and drag to draw a rectangle:

Minutely schedule

Windows10

Protection

Every 1 minutes Exceptions No exceptions Keep 6



0 · 2 · 4 · 6 · 8 · 10 · 12 · 14 · 16 · 18 · 20 · 22 · 24

Mo
Tu
We
Th
Fr
Sa
Su

Select when the schedule won't run (click + drag) Will run Won't run

Extra space used for live snapshots: **0.0 B**

Delete schedule Cancel Save

This is how the time interval looks like after you draw a rectangle:

Minutely schedule

Windows10

Protection

Every 1 minutes Exceptions: Some intervals Keep 6

Select when the schedule won't run (click + drag)

Extra space used for live snapshots: **0.0 B**

Delete schedule Cancel Save

NOTE: In the image above, the hours 06:00 -> 18:00 of the days Monday -> Friday have been excluded. That is, in those time intervals the backup schedule will not run, even though the schedule is otherwise configured to run every 1 minute.

A time interval that has all the blocks green means that there are no hours in the week excluded from the schedule:

Minutely schedule

Windows10

Protection

Every 1 minutes Exceptions No exceptions Keep 6

	0	2	4	6	8	10	12	14	16	18	20	22	24
Mo	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Tu	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
We	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Th	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Fr	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Sa	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Su	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run

Select when the schedule won't run (click + drag) ■ Will run ■ Won't run

Extra space used for live snapshots: **0.0 B**

[Delete schedule](#) Cancel [Save](#)

A time interval that has all the blocks grey means that the snapshot backups will never run:

Minutely schedule

Windows10

Protection

Every 1 minutes Exceptions: Schedule will never run Keep 6

	0	2	4	6	8	10	12	14	16	18	20	22	24
Mo													
Tu													
We													
Th													
Fr													
Sa													
Su													

Select when the schedule won't run (click + drag) Will run Won't run

Extra space used for live snapshots: **0.0 B**

[Delete schedule](#) [Cancel](#) [Save](#)

After you save the schedule, with the exceptions specified, you will be able to see the changes in the lower pane of the *Protect and replicate* page. Below, you can see that the *Minutely* schedule had been configured with *exceptions*. If you click on the *exceptions* link, you will be able to see the configured exceptions (hours in the week when the backup will not run).

The screenshot displays the 'Protect and replicate' page in SynetoOS, specifically the 'Shares' tab. The interface includes a search bar, a 'Show All' dropdown, and a list of shares. A calendar grid shows the protection status for each share across a 24-hour period. A legend indicates that green cells represent 'Will run' and grey cells represent 'Won't run'. Below the calendar, the details for a 'Windows10' share are shown, including its schedule, replication, and retention policy.

Name	VMs	Last snapshot	Last replication	Protection
EmailSer				W
Linux_DE				D
OracleSe				H
Windows				M, D
Windows				H

Windows10

- Schedule:** every 1 minute, with [exceptions](#)
- Replication:** not configured
- Retention policy:** keep 6 locally

Change

- Hourly:** Schedule: every 1 day at 03:00
- Daily:** Schedule: every 1 day at 03:00

Protecting a share

Note: Cloned shares cannot be protected because they are temporary objects.

The *Protect and replicate* page contains three tabs: *Datstores*, *Shares* and *Volumes*. Select *Shares*.

Protect and replicate | [Replication targets](#) | [Stored credentials](#) | 15:13:07

Datstores | Shares | Volumes

Search Show All 1-2 of 2

Name	Type	Last snapshot	Last replication	Protection
documents	SMB	never	never	NOT PROTECTED
time_machine	AFP	never	never	NOT PROTECTED

documents Progress Snapshots Schedules

No snapshots have been taken yet.

The table lists all the shares. There is a *Search* bar, useful for searching for a specific share by specifying its name.

Protect and replicate | [Replication targets](#) | [Stored credentials](#) | 17:56:14

Datstores | Shares | Volumes

time Show All 1-1 of 1 (filtered from 2 total entries)

Name	Type	Last snapshot	Last replication	Protection
time_machine	AFP	never	never	NOT PROTECTED


time_machine Progress Snapshots Schedules

No snapshots have been taken yet.

If you click on a share, more details about its protection will be shown.

Configuring the snapshot schedule for a share

In the details section, there is *Schedules tab*. Click on it. On the left are the four frequencies a user can choose: *Minutely*, *Hourly*, *Daily* and *Weekly*. You can enable one, two, three or all four frequencies for a share.



The screenshot displays the SynetoOS interface for configuring a share. The main window is titled "Protect and replicate" and contains tabs for "Replication targets" and "Stored credentials". The interface is divided into three sections: "Datastores", "Shares", and "Volumes". The "Shares" section is active, showing a table of shares. The table has columns for "Name", "Type", "Last snapshot", "Last replication", and "Protection". Two shares are listed: "documents" (SMB) and "time_machine" (AFP). Both shares have "never" for both "Last snapshot" and "Last replication" and are marked as "NOT PROTECTED". Below the table, the "time_machine" share is selected, and the "Schedules" tab is active. This tab shows four schedule options: "Minutely" (M), "Hourly" (H), "Daily" (D), and "Weekly" (W). Each option has a red switch to its left, indicating that none of the schedules are currently enabled.

Name	Type	Last snapshot	Last replication	Protection
documents	SMB	never	never	NOT PROTECTED
time_machine	AFP	never	never	NOT PROTECTED

time_machine Progress Snapshots Schedules

- M Minutely
- H Hourly
- D Daily
- W Weekly

To enable one schedule type, click on the switch on the left of the desired interval. A dialog will appear where you can configure the snapshot policy for the selected share and frequency.

Hourly schedule

time_machine

Protection

Every	Start minute	Exceptions	Keep
1 ▼ hours	00	No exceptions ▼	6

Replicate to

Don't replicate ▼

Cancel Save

Here you can specify how often you want the schedule to perform by changing the *Every* field.

NOTE: We recommend spacing out snapshots taking time with at least 5 minutes. If all schedules run at the same time, and then replicate, it can lead to high network traffic and high drive IOPS on the replication target.

To the right of the dialog, there is the *Keep* field. Here you can specify how many snapshots should be kept on the Syneto HYPER. If more snapshots accumulate, older ones will be deleted.

NOTE: The retention policy applies to a type of interval only and for the selected share.

Replicate to section allows you to select to which replication target to replicate. By default, no replication is set. If you click below *Replicate to*, a list of the defined replication targets will be shown.

Hourly schedule

time_machine

Protection

Every	Start minute	Exceptions	Keep
1 hours	00	No exceptions	6

Replicate to

Don't replicate

- Don't replicate
- DR-target

Cancel Save

After you select the desired replication target, two more fields will appear.

Hourly schedule

time_machine

Protection

Every	Start minute	Exceptions	Keep
1 hours	00	No exceptions	6

Replicate to

DR-target

Encrypt

Keep

48

Cancel Save

If you want the replication to be encrypted, click on the box below the *Encrypt* field.

NOTE: Non-encrypted replication is faster than encrypted one, but less secure. If the replication target is off-site, we recommend encryption.

You can specify how many replicated snapshots should be kept on the replication target under the *Keep* field.

NOTE: The retention policy applies to a type of interval only and for the selected share.

NOTE: We recommend a lower retention policy on the primary unit and a higher one on the replication target (DR unit).

After you finish all the configuration, click *Save*. Your schedules will appear in the details section, in *Schedules*. You can edit by clicking on *Change*.

The screenshot displays the SynetoOS interface for managing replication targets. The main view shows a table of replication targets:

Name	Type	Last snapshot	Last replication	Protection
documents	SMB	never	never	NOT PROTECTED
time_machine	AFP	never	never	H

Below the table, the details for the 'time_machine' target are shown, including a 'Schedules' tab. The current schedule is 'Hourly' (H), which is active (green switch). The schedule details are as follows:

Interval	Frequency	Start minute	Keep
Hourly	Every	00	6
Replicate to	DR-target	Encrypt	Keep

Other available schedules include 'Minutely' (M), 'Daily' (D), and 'Weekly' (W), all of which are currently inactive (red switches). A notification at the bottom right states: 'Protection schedule saved for time_machine. 12:05:15 - 16 Jul, 2018'.

A schedule can be paused by clicking on the *green switch* button. This will not stop a snapshot or replication which is in progress. If you want to resume it, just click on the *red switch* button.

If you want to delete a schedule, click on *Change*. A dialog will open. Click on *Delete schedule*.

Hourly schedule

time_machine

Protection

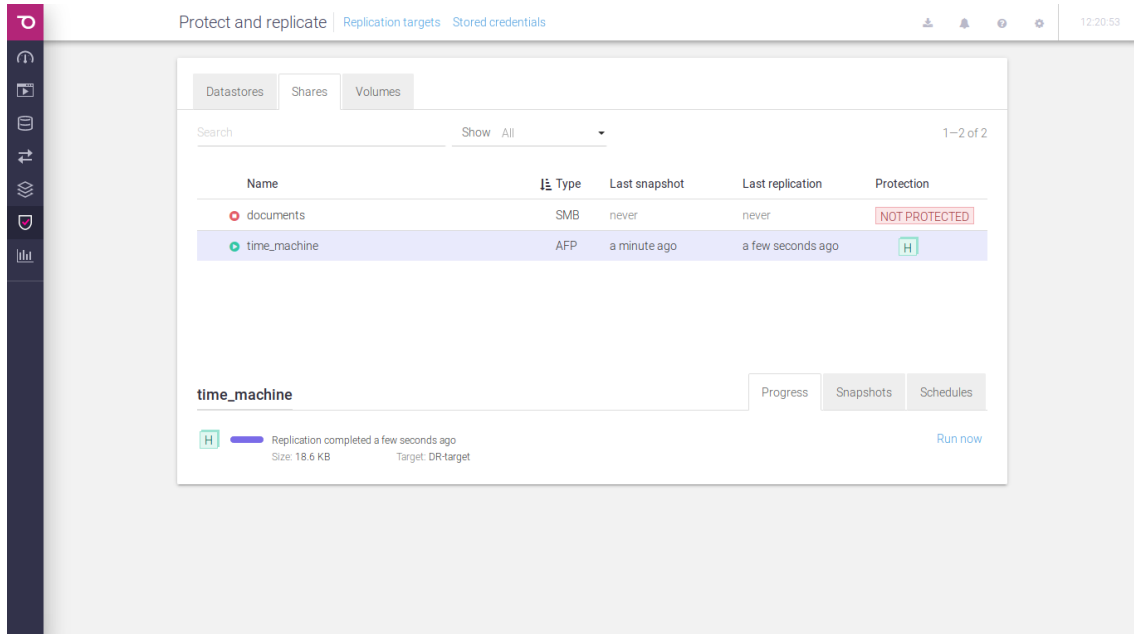
Every	Start minute	Exceptions	Keep
1 hours	00	No exceptions	6

Replicate to	Encrypt	Keep
DR-target	<input type="checkbox"/>	48

[Delete schedule](#) Cancel Save

Checking the progress of a snapshot or replication of a share

If you want to check the progress of a snapshot schedule, go to *Progress* tab in details section.



You will see for every enabled schedule the current state. You can click *Run now* if you want to start the schedule in this moment.

NOTE: First replication will always copy all the data found in the share. Subsequent replications will transfer only the differences between the last successful replication and current snapshot.

Canceling a running replication of a share

A replication that is currently running can be canceled. As a consequence, the replica will not be created on the destination. To cancel it, click on the cancel icon on the right of the *Running ...* message.

Protect and replicate | [Replication targets](#) | [Stored credentials](#) 13:08:23

Datstores Shares Volumes


Search Show All 1-2 of 2

Name	Type	Last snapshot	Last replication	Protection
documents	SMB	never	never	NOT PROTECTED
time_machine	AFP	a few seconds ago	<div style="width: 20%;"></div>	H

time_machine Progress Snapshots Schedules

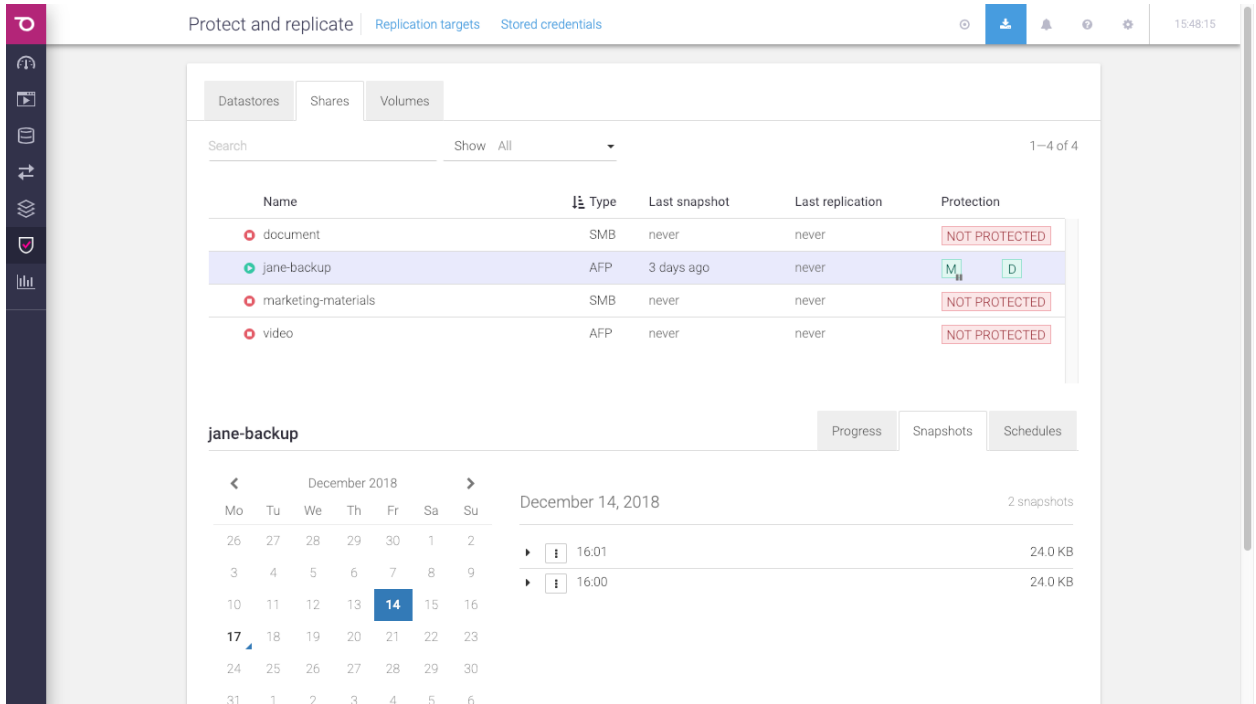
H 20% - Replicating ... Size: 1.9 GB Speed: 33.1 MB/s ETA: a minute left Target: DR-target Running ...

After you click on the cancel replication icon, you will receive a confirmation saying that the Snapshot replication was cancelled.

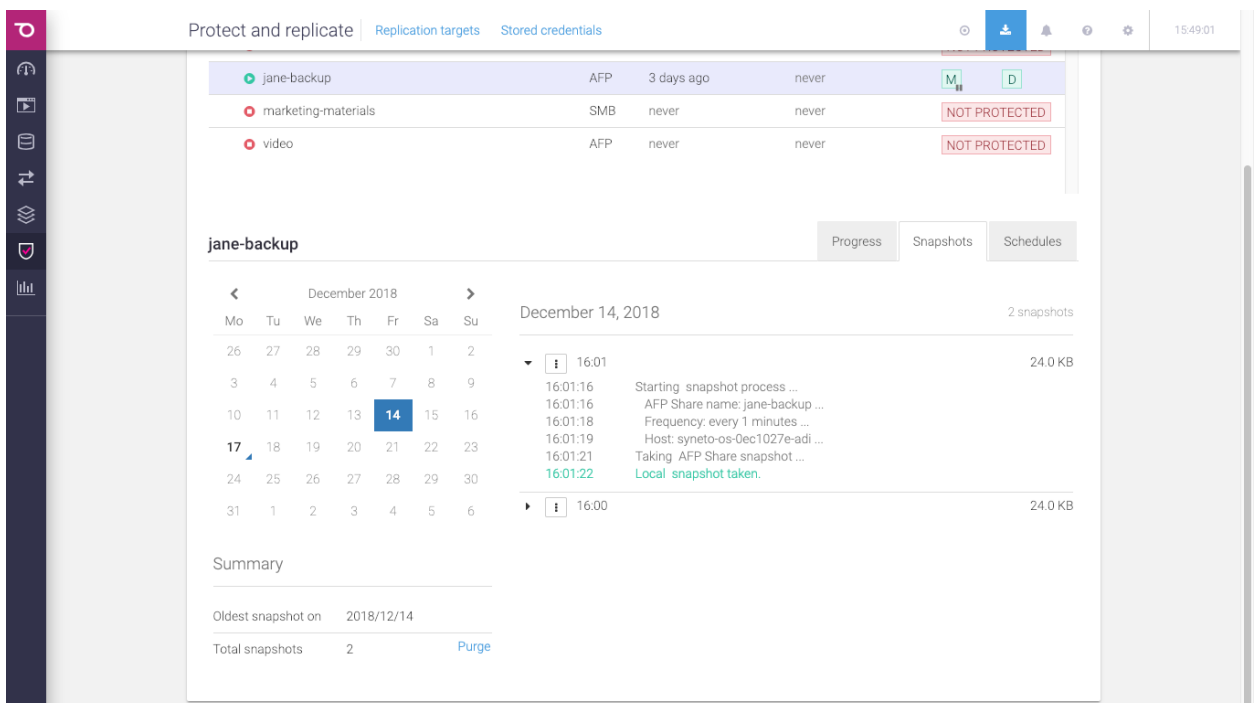


Monitoring the snapshots of a share

If you want to check the list of snapshots of a share, go to *Snapshots* tab in details section.



On the left, there is a calendar. On the right there is a list of snapshots from the date that is selected. You can change the day of the month by clicking on the desired day or you can go to another month by clicking the left or right arrows at the top of the calendar. Only dates with snapshots will be available for selection and the current date. To check the history of a snapshot, click on it in the right section.



Exceptions: Excluding time intervals from Schedules

The *Minutely* and *Hourly* schedules allow you to exclude time intervals - ranges of hours, throughout the days of the week. Snapshotting & replication will not happen in those time intervals.

Excluding time intervals can be useful if you want to backup your data at high granularity, but you know that one time or more during the week there is an important operation that your system needs to perform. If that operation could interfere with the scheduled backup, or is computationally expensive (e.g. intensive disk IO), it might be better to simply exclude that time interval from the backup schedule.

In order to use exceptions in the backup schedule, look for the *Exceptions* field in the *Minutely / Hourly* schedule, and click the link (initially stating *No exceptions*). A green coloured table will appear, with the days of the week as the row and hours of the day as the columns.

In order to exclude intervals from the schedule, click and drag, so as to draw rectangles in the green area: grey blocks represent hours in the week when the backup will not run, while green blocks represent hours in the week when the backups will be allowed to run.

NOTE: You can also draw rectangles in grey areas - this will cause blocks within the grey area to turn back green (and thus allow the backup to run in those hours).

NOTE: You can also drag within a single block - this will cause that block (hour in the week) to toggle from being excluded (grey) to not being excluded (green), and vice versa.

This is how you use, click and drag to draw a rectangle:

Minutely schedule

share-1

Protection

Every 1 minutes Exceptions: Some intervals Keep: 6

	0	2	4	6	8	10	12	14	16	18	20	22	24
Mo	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Tu	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
We	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Th	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Fr	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Sa	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Su	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run

Select when the schedule won't run (click + drag) ■ Will run ■ Won't run

This is how the time interval looks like after you draw a rectangle:

Minutely schedule

share-1

Protection

Every 1 minutes

Exceptions: Some intervals

Keep: 6

	0	2	4	6	8	10	12	14	16	18	20	22	24
Mo	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Tu	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
We	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Th	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Fr	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Sa	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Su	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run

Select when the schedule won't run (click + drag)

Will run Won't run

NOTE: In the image above, the hours 06:00 -> 18:00 of the days Monday -> Friday have been excluded. That is, in those time intervals the backup schedule will not run, even though the schedule is otherwise configured to run every 1 minute.

A time interval that has all the blocks green means that there are no hours in the week excluded from the schedule:

Minutely schedule

share-1

Protection

Every 1 minutes Exceptions Some intervals Keep 6

	0	2	4	6	8	10	12	14	16	18	20	22	24
Mo	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Tu	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
We	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Th	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Fr	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Sa	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Su	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run

Select when the schedule won't run (click + drag) Will run Won't run

A time interval that has all the blocks grey means that the snapshot backups will never run:

Minutely schedule

share-1

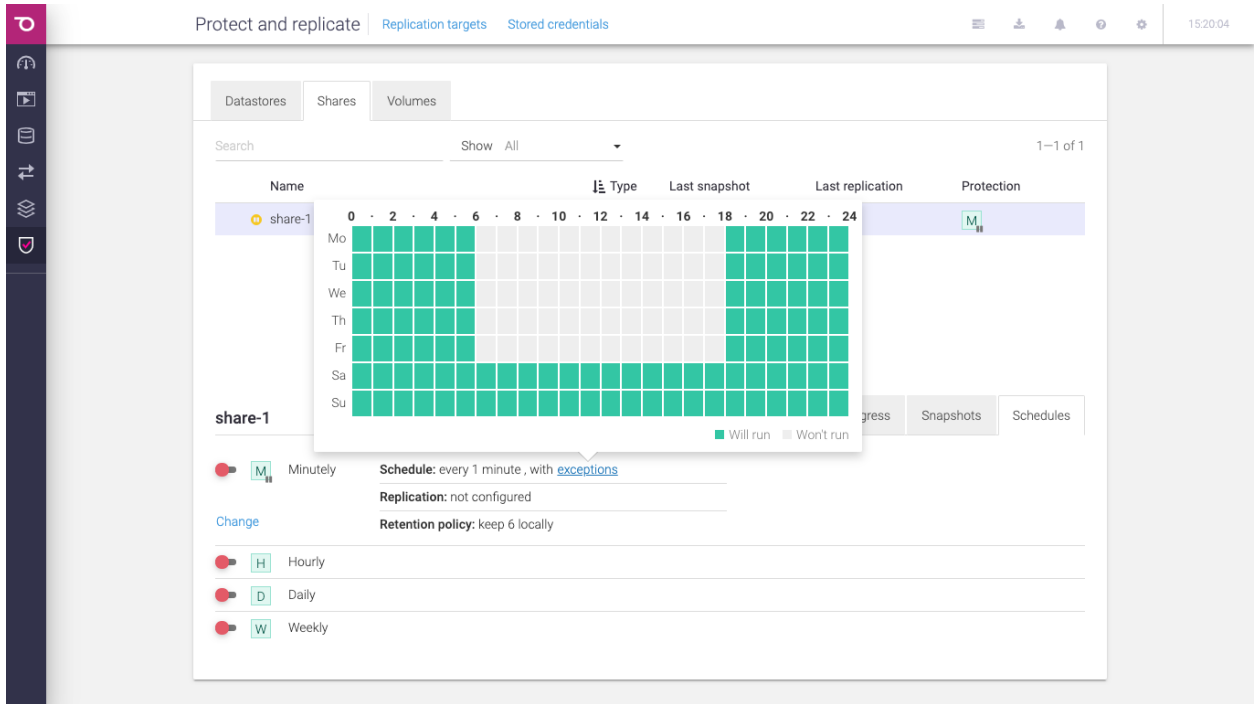
Protection

Every 1 minutes Exceptions Schedule will never run Keep 6

	0	2	4	6	8	10	12	14	16	18	20	22	24
Mo													
Tu													
We													
Th													
Fr													
Sa													
Su													

Select when the schedule won't run (click + drag) ■ Will run ■ Won't run


After you save the schedule, with the exceptions specified, you will be able to see the changes in the lower pane of the *Protect and replicate* page. Below, you can see that the *Minutely* schedule had been configured with *exceptions*. If you click on the *exceptions* link, you will be able to see the configured exceptions (hours in the week when the backup will not run).



Protecting a volume

Note: Cloned volumes cannot be protected because they are temporary objects.


The *Protect and replicate* page contains 3 sections: *Datastores*, *Shares* and *Volumes*. Select *Volumes*.



The screenshot shows the 'Protect and replicate' section of the SynetoOS interface. The 'Volumes' tab is selected, and a table lists two volumes: 'disk1' and 'disk2'. Both volumes have a status of 'NOT PROTECTED'. The table includes columns for Name, Last snapshot, Last replication, and Protection. Below the table, there are buttons for 'Progress', 'Snapshots', and 'Schedules' for the selected volume 'disk1'. A message states 'No snapshots have been taken yet.'

Name	Last snapshot	Last replication	Protection
disk1	never	never	NOT PROTECTED
disk2	never	never	NOT PROTECTED

The table lists all the volumes. There is a *Search* bar, useful for searching for a specific volume by specifying its name.



The screenshot shows the same 'Protect and replicate' section, but with the search bar containing 'disk1'. The table now only displays 'disk1', and the status is 'NOT PROTECTED'. The table includes columns for Name, Last snapshot, Last replication, and Protection. Below the table, there are buttons for 'Progress', 'Snapshots', and 'Schedules' for the selected volume 'disk1'. A message states 'No snapshots have been taken yet.'

Name	Last snapshot	Last replication	Protection
disk1	never	never	NOT PROTECTED

If you click on a volume, more details about its protection will be shown.

Configuring the snapshot schedule for a volume

In the details section, there is *Schedules tab*. Click on it. On the left are the four frequencies a user can choose: *Minutely*, *Hourly*, *Daily* and *Weekly*. You can enable one, two, three or all four frequencies for a volume.

The screenshot shows the SynetoOS interface. At the top, there are navigation links for 'Protect and replicate', 'Replication targets', and 'Stored credentials'. The main content area is divided into three tabs: 'Datastores', 'Shares', and 'Volumes'. The 'Volumes' tab is active, displaying a table with the following data:

Name	Last snapshot	Last replication	Protection
disk1	never	never	NOT PROTECTED
disk2	never	never	NOT PROTECTED

Below the table, there is a section for 'disk1' with three sub-tabs: 'Progress', 'Snapshots', and 'Schedules'. The 'Schedules' sub-tab is active, showing four frequency options, each with a toggle switch and a label:

- M Minutely
- H Hourly
- D Daily
- W Weekly

To enable one schedule type, click on the switch on the left of the desired interval. A dialog will appear where you can configure the snapshot policy for the selected volume and frequency.

Hourly schedule

disk1

Protection

Every	Start minute	Exceptions	Keep
1 ▾ hours	00	No exceptions ▾	6

Replicate to

Don't replicate ▾

Cancel Save

Here you can specify how often you want the schedule to perform by changing the *Every* field.

NOTE: We recommend spacing out snapshot taking time with at least 5 minutes. If all schedules run at the same time, and then replicate, it can lead to high network traffic and high drive IOPS on the replication target.

To the right of the dialog, there is the *Keep* field. Here you can specify how many snapshots should be kept on the Syneto HYPER. If more snapshots accumulate, older ones will be deleted.

NOTE: The retention policy applies to a type of interval only and for the selected volume.

Replicate to section allows you to select to which replication target to replicate. By default, no replication is set. If you click below *Replicate to*, a list of the defined replication targets will be shown.

Hourly schedule

disk1

Protection

Every	Start minute	Exceptions	Keep
1 ▼ hours	00	No exceptions ▼	6

Replicate to

Don't replicate ▼

- Don't replicate
- DR-target

Cancel Save

After you select the desired replication target, two more fields will appear.

Hourly schedule

disk1

Protection

Every	Start minute	Exceptions	Keep
1 ▼ hours	00	No exceptions ▼	6

Replicate to

DR-target ▼

Encrypt

Keep

48

Cancel Save

If you want the replication to be encrypted, click on the box below the *Encrypt* field.

NOTE: Non-encrypted replication is faster than encrypted one, but less secure. If the replication target is off-site, we recommend encryption.

You can specify how many replicated snapshots should be kept on the replication target under the *Keep* field.

NOTE: The retention policy applies to a type of interval only and for the selected volume.

NOTE: We recommend a lower retention policy on the primary unit and a higher one on the replication target (DR unit).

After you finish all the configuration, click *Save*. Your schedules will appear in the details section, in *Schedules*. You can edit by clicking on *Change*.

The screenshot shows the SynetoOS interface for managing replication targets. The 'Protect and replicate' section is active, with the 'Volumes' tab selected. A table lists the volumes and their protection status:

Name	Last snapshot	Last replication	Protection
disk1	never	never	H
disk2	never	never	NOT PROTECTED

Below the table, the 'disk1' details are shown, including the 'Schedules' section. The 'Hourly' schedule is selected and active, with a green switch button. The schedule details are:

- Schedule:** every 1 hour at minute 00
- Replication:** DR-target, not encrypted
- Retention policy:** keep 6 locally and 48 remotely

A notification at the bottom right states: 'Protection schedule saved for disk1. 15:02:20 - 28 Feb, 2019'. A 'Hide all notifications' button is also visible.

A schedule can be paused by clicking on the *green switch* button. This will not stop a snapshot or replication which is in progress. If you want to resume it, just click on the *red switch* button.

If you want to delete a schedule, click on *Change*. A dialog will open. Click on *Delete schedule*.

Hourly schedule

disk1

Protection

Every	Start minute	Exceptions	Keep
1 hours	00	No exceptions	6

Replicate to	Encrypt	Keep
DR-target	<input type="checkbox"/>	48

[Delete schedule](#) Cancel Save

Checking the progress of a snapshot or replication of a volume

If you want to check the progress of a snapshot schedule, go to *Progress* tab in details section.

The screenshot displays the 'Protect and replicate' interface. At the top, there are tabs for 'Datastores', 'Shares', and 'Volumes'. Below these is a search bar and a 'Show All' dropdown. A table lists replication targets:

Name	Last snapshot	Last replication	Protection
disk1	a few seconds ago	a few seconds ago	[H]
disk2	never	never	NOT PROTECTED

Below the table, a detailed view for 'disk1' is shown. It includes a 'Progress' bar, 'Snapshots', and 'Schedules' tabs. A message indicates: 'Replication completed a few seconds ago. Size: 6.6 KB. Target: DR-target. Run now'.

You will see for every enabled schedule the current state. You can click *Run now* if you want to start the schedule in this moment.

NOTE: First replication will always copy all the data found in the volume. Subsequent replications will transfer only the differences between the last successful replication and current snapshot.

Canceling a running replication of a volume


A replication that is currently running can be canceled. As a consequence, the replica will not be created on the destination. To cancel it, click on the cancel icon on the right of the *Running ...* message.

The screenshot displays the 'Protect and replicate' section of the SynetoOS interface. It features a navigation sidebar on the left and a main content area. The main content area has tabs for 'Datastores', 'Shares', and 'Volumes'. Below these tabs is a search bar and a 'Show All' dropdown. A table lists two replication targets:

Name	Last snapshot	Last replication	Protection
disk1	a few seconds ago		
disk2	never	never	NOT PROTECTED

Below the table, there is a detailed view for 'disk1' with tabs for 'Progress', 'Snapshots', and 'Schedules'. The 'Progress' tab is active, showing a progress bar at 67% with the text 'Replicating ...'. Other details include 'Size: 6.6 KB', 'Speed: 6.5 MB/s', 'ETA: a few seconds left', and 'Target: DR-target'. A 'Running' status with a refresh icon is also visible.

After you click on the cancel replication icon, you will receive a confirmation saying that the Snapshot replication was cancelled.



Monitoring the snapshots of a volume

If you want to check the list of snapshots of a datastore, go to *Snapshots* tab in details section.

Protect and replicate | [Replication targets](#) | [Stored credentials](#)

Datastores | Shares | Volumes

Search Show: All 1—2 of 2

Name	Last snapshot	Last replication	Protection
Volume1	3 days ago	never	M ₁₁ D ₁₁ W ₁₁
Volume2	3 days ago	never	M ₁₁ H ₁₁ D ₁₁ W ₁₁

Volume1 Progress | Snapshots | Schedules

< December 2018 >

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

December 14, 2018 3 snapshots

- 17:27 0.0 B
- 17:26 0.0 B
- 17:25 0.0 B

On the left, there is a calendar. On the right there is a list of snapshots from the date that is selected. You can change the day of the month by clicking on the desired day or you can go to another month by clicking the left or right arrows at the top of the calendar. Only dates with snapshots will be available for selection and the current date. To check the history of a snapshot, click on it in the right section.

Protect and replicate | [Replication targets](#) | [Stored credentials](#)

Volume2 3 days ago never M₁₁ H₁₁ D₁₁ W₁₁

Volume1 Progress | Snapshots | Schedules

< December 2018 >

Mo	Tu	We	Th	Fr	Sa	Su
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

December 14, 2018 3 snapshots

- 17:27 0.0 B
 - 17:27:38 Starting snapshot process ...
 - 17:27:39 ISCSI Volume name: Volume1 ...
 - 17:27:40 Frequency: every friday ...
 - 17:27:41 Host: syneto-os-0ec1027e-adi ...
 - 17:27:42 Taking ISCSI Volume snapshot ...
 - 17:27:43 Local snapshot taken.
- 17:26 0.0 B
- 17:25 0.0 B

Summary

Oldest snapshot on 2018/12/14

Total snapshots 3 [Purge](#)

Exceptions: Excluding time intervals from Schedules

The *Minutely* and *Hourly* schedules allow you to exclude time intervals - ranges of hours, throughout the days of the week. Snapshotting & replication will not happen in those time intervals.

Excluding time intervals can be useful if you want to backup your data at high granularity, but you know that one time or more during the week there is an important operation that your system needs to perform. If that operation could interfere with the scheduled backup, or is computationally expensive (e.g. intensive disk IO), it might be better to simply exclude that time interval from the backup schedule.

In order to use exceptions in the backup schedule, look for the *Exceptions* field in the *Minutely / Hourly* schedule, and click the link (initially stating *No exceptions*). A green coloured table will appear, with the days of the week as the row and hours of the day as the columns.

In order to exclude intervals from the schedule, click and drag, so as to draw rectangles in the green area: grey blocks represent hours in the week when the backup will not run, while green blocks represent hours in the week when the backups will be allowed to run.

NOTE: You can also draw rectangles in grey areas - this will cause blocks within the grey area to turn back green (and thus allow the backup to run in those hours).

NOTE: You can also drag within a single block - this will cause that block (hour in the week) to toggle from being excluded (grey) to not being excluded (green), and vice versa.

This is how you use click and drag to draw a rectangle:

Minutely schedule

volume-1

Protection

Every 1 minutes Exceptions Some intervals Keep 6

	0	2	4	6	8	10	12	14	16	18	20	22	24
Mo	Will run	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Tu	Will run	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
We	Will run	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Th	Will run	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Fr	Will run	Will run	Will run	Will run	Will run	Won't run	Won't run	Won't run	Won't run	Won't run	Will run	Will run	Will run
Sa	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Su	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run

Select when the schedule won't run (click + drag) Will run Won't run

This is how the time interval looks like after you draw a rectangle:

Minutely schedule

volume-1

Protection

Every 1 minutes Exceptions Some intervals Keep 6

Select when the schedule won't run (click + drag) ■ Will run ■ Won't run

NOTE: In the image above, the hours 06:00 -> 18:00 of the days Monday -> Friday have been excluded. That is, in those time intervals the backup schedule will not run, even though the schedule is otherwise configured to run every 1 minute.

A time interval that has all the blocks green means that there are no hours in the week excluded from the schedule:

Minutely schedule

volume-1

Protection

Every 1 minutes Exceptions No exceptions Keep 6

	0	2	4	6	8	10	12	14	16	18	20	22	24
Mo	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Tu	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
We	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Th	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Fr	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Sa	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run
Su	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run	Will run

Select when the schedule won't run (click + drag) ■ Will run ■ Won't run

A time interval that has all the blocks grey means that the snapshot backups will never run:

Minutely schedule

volume-1

Protection

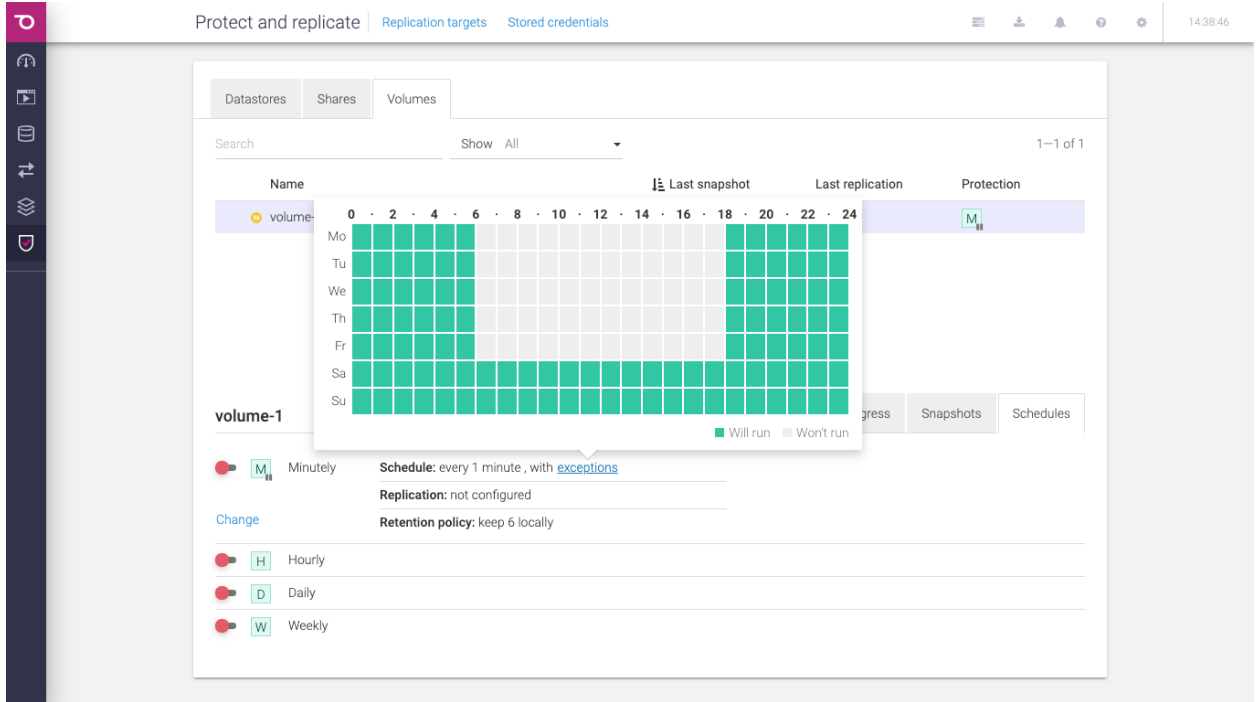
Every 1 minutes Exceptions Schedule will never run Keep 6

	0	2	4	6	8	10	12	14	16	18	20	22	24
Mo													
Tu													
We													
Th													
Fr													
Sa													
Su													

Select when the schedule won't run (click + drag) ■ Will run ■ Won't run

After you save the schedule, with the exceptions specified, you will be able to see the changes in the lower pane of the *Protect and replicate* page. Below, you can see that the *Minutely* schedule had been configured with *exceptions*. If you click on the *exceptions* link, you will be able to see the configured exceptions (hours in the week when the backup will not run).

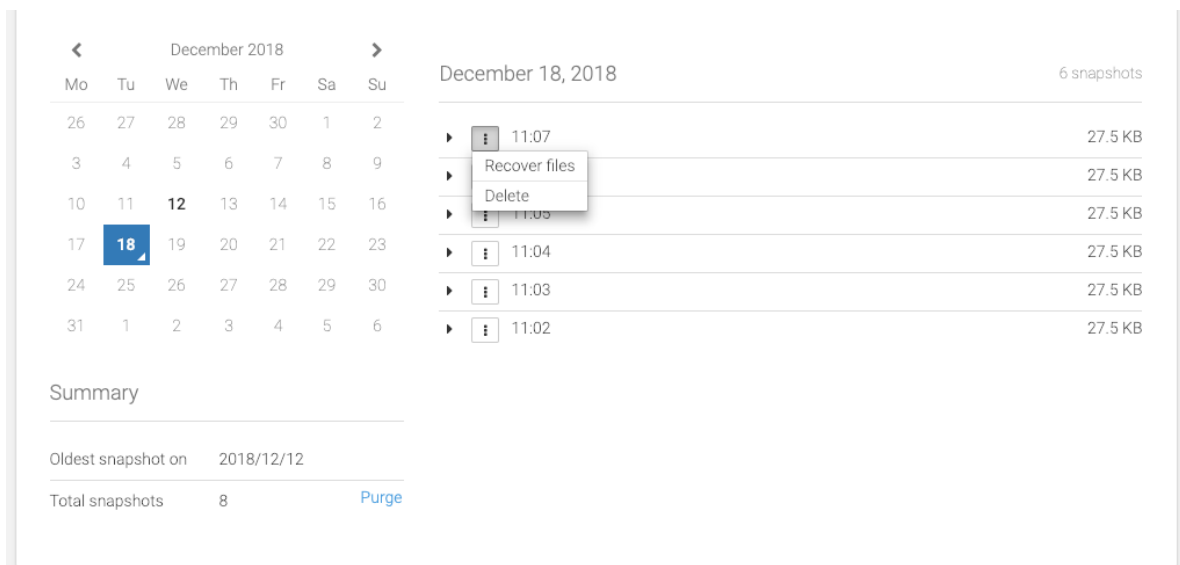
NOTE: Only the *Minutely* and the *Hourly* schedules support *exceptions*.



Recovering files from a virtual machine

To recover files from a virtual machine you can go to the Snapshots tab on the Virtual machines, Datastores or Protect and replicate page.

To recover files click *Recover files* on the desired snapshot context menu.



Browse file recover

Recover files from a virtual machine

Files to recover


Virtual machine

Windows10_x64

Partition 0 (System Reserved, 500.0 MB)

Partition 1 (31.5 GB)


Recovery method

 Please select folders and files to continue the recovery process.

Cancel

After selecting the folders and files choose the recovery method you want to use, *Download* or *Attach to VM*.

Regardless of the recovery method you choose, when starting the recovery process, SynetoOS will create a new datastore on the ESXi for recovery. You will not be able to see this datastore from the SynetoOS UI, but it will be visible from within the ESXi UI.



The screenshot shows the SynetoOS Storage interface for the host 'syneto-esxi-ba87c8c4.dev.syneto.net'. The 'Datastores' tab is active. The interface includes a search bar and several action buttons: 'New datastore', 'Increase capacity', 'Register a VM', 'Datastore browser', 'Refresh', and 'Actions'. A table lists the following datastores:

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provisio...	Access
datastore1	Non-SSD	32.5 GB	11.52 GB	20.98 GB	VMFS6	Supported	Single
OS Images	Unknown	5.94 TB	523.6 GB	5.43 TB	NFS	Supported	Single
syneto-file-recovery-VMware-42073c3225544723-...	Unknown	2.24 GB	100.09 MB	2.14 GB	NFS	Supported	Single
virtual-machines	Unknown	18.88 GB	789.89 MB	18.11 GB	NFS	Supported	Single

4 items


In the image above, the generated datastore is:

syneto-file-recovery-VMware-42073c3225544723-ef00f567b108ce9c

NOTE: If you happen to delete the recovery datastore (by mistake), there is no need to worry: SynetoOS will generate it again if it does not exist. You should, however, not delete this datastore yourself - especially if there is a recovery of files in progress.


Download files

Click the *Download* tab → *Create archive*.



This will compress the selected folders and files to a zip archive and make it available for download in file recovery progresses.

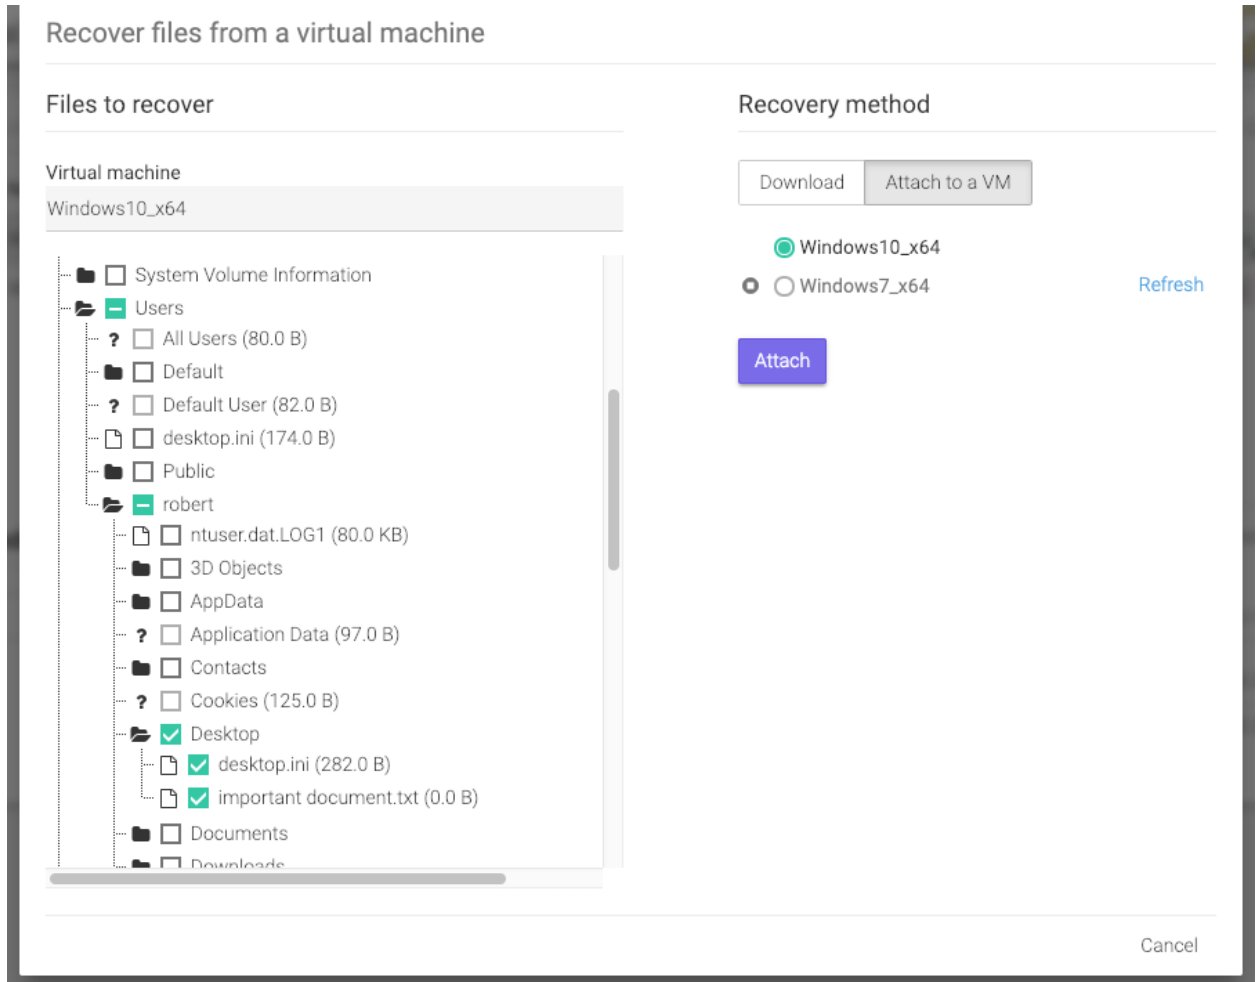
After the archive was created , click → *Download*




Attach files to a VM

Click the *Attach to a VM* tab → select in which VM to attach the files → *Attach*.

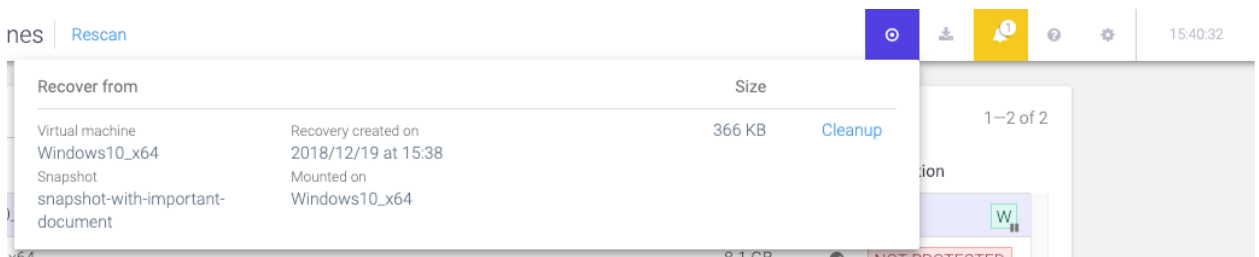
NOTE: Virtual machines without a CD-ROM or powered off can not be selected.



WARNING: The selected files and folders will be attached to the selected virtual machine's CD-ROM. Current content of the CD-ROM will be ejected.




During the restore process the recovered files will be bundled into an ISO. This will use the system resources to archive files and mount it to the selected virtual machine.




NOTE: Creating the ISO may take a long time if the selected files and folders have a big size.

Open the CD-ROM to view the selected files and folders.



You will be able to copy the files and folders.




NOTE: After 24 hours, the ISO will be automatically ejected and deleted.

SYSTEM CONFIGURATION

Networking

Select *Secondary menu* → *Basic configuration*, in the *Network* section.



The *Basic network configuration* page allows you to configure the following:

- *Hostname* - the name of the SynetoOS machine

WARNING: This will require a reboot.

A screenshot of the 'Basic network configuration' page. The page has a header with four tabs: 'Hostname', 'Domain', 'Gateway', and 'Proxy'. The 'Hostname' tab is active. Below the tabs, there is a yellow warning box that says: 'After you click update, the storage device will need to reboot in order to apply the new hostname.' Below the warning box, there is a text input field labeled 'Hostname' with the value 'syneto-os' and a help icon. At the bottom left, there is a blue 'Update' button.

- *Domain* - a domain name to append to hostnames during DNS name resolution and up to three domain name servers to be used when resolving internet names

The screenshot shows a configuration panel with four tabs: Hostname, Domain, Gateway, and Proxy. The 'Domain' tab is selected. It contains a 'Domain name' field with the value 'company.domain' and a help icon. Below it are three 'DNS' fields: 'DNS 1' with the value '192.168.1.15', 'DNS 2' with a masked value '#####', and 'DNS 3' with a masked value '#####' and a help icon. An 'Update' button is located at the bottom left.

- **Gateway** - a router used by the storage to connect to the internet

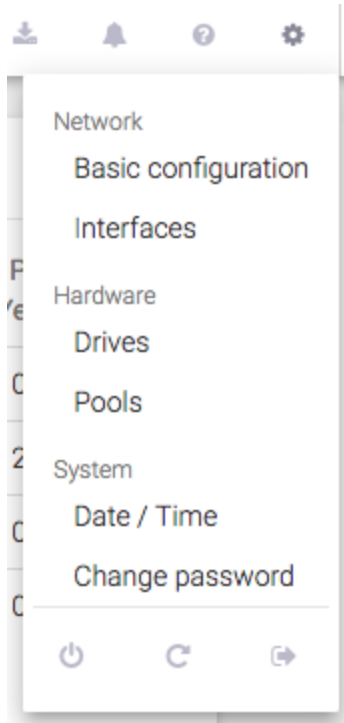
The screenshot shows a configuration panel with four tabs: Hostname, Domain, Gateway, and Proxy. The 'Gateway' tab is selected. It contains a 'Default gateway' field with the value '192.168.1.1' and a help icon. An 'Update' button is located at the bottom left.

- **Proxy** - the host and port of the HTTP proxy server


The screenshot shows a configuration panel with four tabs: Hostname, Domain, Gateway, and Proxy. The 'Proxy' tab is selected. It contains a checkbox labeled 'Use a HTTP proxy?' which is currently unchecked. An 'Update' button is located at the bottom left.

Ethernet interfaces

Select *Secondary menu* → *Interfaces*, in the *Network* section.




Syneto HYPER comes preconfigured with two network interfaces. The first one is the management and the second one is used for interaction with the ESXi host. These are not physical interfaces, they are virtualized interfaces provided by the local ESX.



The *Network interfaces* page allows you to view the state and configuration of all the ethernet interfaces. It also allows you to configure the ethernet interface of any interface other than the one for communication with the local hypervisor.

Click the *Edit* option in the context menu.




A dialog will open where you can configure the network interface. You can have a *dynamic* or *static* IP. Also, you can configure a limit for the data being transmitted through that interface in the *MTU* field.

Edit interface

`vmxnet3s0`

Address Static DHCP


IP / netmask
192.168.5.14/21

MTU
1500 

Changing the network interface will remount all datastores mounted through this network. This may take a few minutes. Virtual machines will be unregistered, you will have to register them manually afterwards.

Cancel

NOTE: SynetoOS 4 does not offer support for aggregates or VLANs.

The IP address of the `vmxnet3s1` interface cannot be modified. Hovering the cursor over the  button associated to this interface (which cannot be clicked, for editing) will bring out a notice mentioning this.



Missing or broken interfaces


In case one or more of the ethernet interfaces are faulted, the system will display the message: *missing: removed or broken!*

If **the interface was intentionally removed**, click on delete.

If **the interface was replaced with another type of virtualized network**, the system will detect the new interface. The faulted interface will still be shown and can be removed by clicking on *Delete*, in the *Context menu*.

Date/time and timezones

Select *Secondary menu* → *Date / Time*, in the *System* section.



The time page allows you to configure the following:

- *Date / Time*: The date and time can be automatically synchronized with a NTP server or set manually.

Recommended: Enable *automatic time synchronization*.

Date / Time
Time zone

Automatic time synchronization OFF ON

NTP server
 ?

[Update & synchronize](#)

Date / Time
Time zone

Automatic time synchronization OFF ON

Date / Time
 📅

[Set date & time](#)

- *Timezone*

Date / Time
Time zone

Changing the time zone will restart SynetoOS.

Time zone
 ▼

[Change time zone](#)

WARNING: This will require a reboot.

Password change

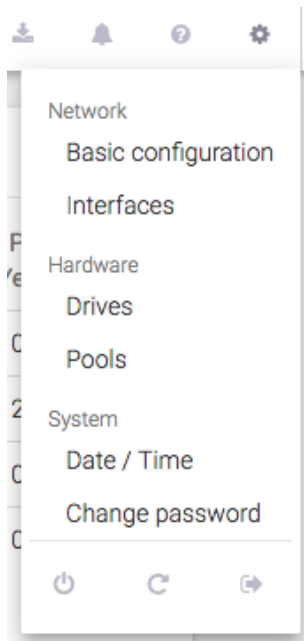
There are two sets of users and passwords that you need to manage regarding Syneto HYPER: local ESXi password, and SynetoOS management interface password.

NOTE: We highly recommend changing default passwords for obvious security reasons.

Changing local ESXi and SynetoOS passwords

The passwords can be changed by going to the *Secondary (Cog) Menu* and clicking *Change password*.

WARNING: If a wrong current ESXi password is provided, after 5 failed attempts the account will be locked for 60 seconds.



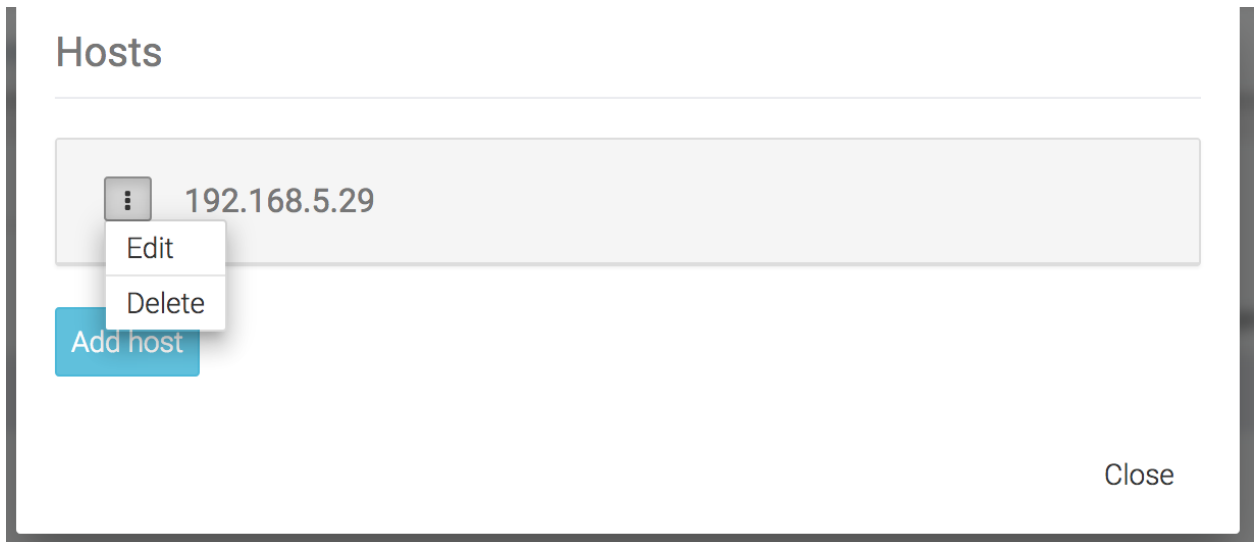
Change passwords

Management VM (Syneto OS)	Hypervisor (ESXi)
Current password <input type="password"/>	Current password <input type="password"/>
New password <input type="password"/>	New password <input type="password"/>
Re-type new password <input type="password"/>	Re-type new password <input type="password"/>
<input type="button" value="Change"/>	<input type="button" value="Change"/>


Close

Changing External ESXi Host Password


If an external ESXi password is changed from the vSphere management interface the password must be updated in SynetoOS too. In SynetoOS go to *Datastores* → *External hosts* and select *Edit* from contextual menu.



Then a dialog will open, where the new password can be introduced.



After hypervisor password is updated successfully, a confirmation message is displayed.



Alerts and notifications

Syneto wants to keep you well informed about what goes around with your machine. Alerts and email notifications are very useful in managing the few errors that may appear. We feel it is important that you receive feedback for every action you take, at all times. Also, you are the first to see when a new software version is available.

Alerts Page


Alerts are important because they let you know when something went wrong. There can be *Warnings (Orange Alerts)* or *Errors (Red Alerts)*. You will also be shown the number of issues on your machine.

The screenshot displays the Alerts page interface. At the top, there are navigation links for 'Alerts', 'Email recipients', and 'Email sender service'. A notification bell icon with a '2' is visible in the top right corner. The main content area is titled 'System errors' and contains two error entries:

- A system service failed** (e438b50c-c2c6-6593-e969-b7952bf54a43) - 2018/07/13, 14:01:40. Description: Failed replicating datastore 411-ds1 to 192.168.5.20. While running hourly schedule for datastore 411-ds1, an error occurred. The system will try replicating the snapshot again when the next schedule runs.
- A system service failed** (98bb38cf-b674-e11d-c932-a448de7b502a) - 2018/07/13, 12:39:29. Description: Failed taking snapshot of datastore 411-ds1. While running minute-by-minute schedule for datastore 411-ds1, an error occurred. The system will try replicating the snapshot again when the next schedule runs.

Below the error list, there are several health status cards:

- CPU health**: Get help. Everything works fine.
- Memory health**: Get help. Everything works fine.
- Drives health**: See more. Everything works fine.
- Storage health**: See more. Everything works fine.
- Virtual machines health**: See more. Everything works fine.

To see the *Alerts* page, go to action bar and select the *bell* .

Here, you will find information about the *System errors*, *CPU health*, *Memory health*, *Drives health*, *Storage health*, *Virtual machines health* and *Sensors health*.

The screenshot displays a health dashboard with the following components:

- System errors:** A single card at the top showing "Everything works fine".
- CPU health:** A card with "Everything works fine" and a "Get help" link.
- Memory health:** A card with "Everything works fine" and a "Get help" link.
- Drives health:** A card with "Everything works fine" and a "See more" link.
- Storage health:** A card with "Everything works fine" and a "See more" link.
- Virtual machines health:** A card with "Everything works fine" and a "See more" link.
- Sensors health:** A wide card at the bottom with "Everything works fine" and a "Get help" link.

You can access support by following the *Get help* links,, or you can go directly to the advised page by following the *See more* links.

Some issues don't need support help. You can clear them by going on the *alert* → *more* → *Mark as repaired*.

NOTE: *Mark as repaired* is not recommended unless you know and understand what the problem was, and you cleared it first.

System errors

Expand all / Collapse all

A system service failed (51447a2e-f22f-6d8a-a90d-bbd9b2764f32) 2018/07/16, 15:49:34

Mark as repaired

CPU health [Get help](#)

Memory health [Get help](#)

Drives health [See more](#)

Everything works fine

Everything works fine

Everything works fine

Storage health [See more](#)

Virtual machines health [See more](#)

Everything works fine

Everything works fine

Sensors health [Get help](#)

Everything works fine

Confirm

A fault should be marked as repaired only after the repaired procedure was completed.

Cancel [Mark as repaired](#)

NOTE: It is highly recommended that you introduce at least one email address in order to receive fast and accurate emails in case there are any issues on your machine.




To add an email address, go to *Email recipients* → *Add recipient*. A form will open, where you can add as many valid email addresses as you need.

Add recipient

Email address
name.surname@company.eu

Close **Add**


Email recipients

 name.surname@company.com
 name@company.eu
 name2017@company.it

Add recipient

Close

In order to check that your email can be reached, go to *Email sender service* → *Send test email*.



Email sender service

Outgoing email server Local service SMTP

Sender email
no-reply@syneto

Close Send test email Update

You will receive an email on all added addresses.

no-reply@syneto-os-7919f26c.dev.syneto.net

to me

Event notification test mail


If you received this email, the event notification mail configuration works.


[Click here to login](#)

This is an automated email generated by Syneto OS ([syneto-os-7919f26c.dev.syneto.net](#)) on 16 Jul 2018, 07:58:05.
Please do not respond to this message, as your reply will be discarded.

New Software Update

When there is a new software version available, you will see, besides receiving an email,

that the *Software update* icon in the *Action Bar* turns blue . After update, the color will change back to grey.



You can read more information about New Software Update in [Software updates](#) Chapter.


Notification boxes

With every action you take, there will be a feedback, a confirmation box in the bottom right corner of the screen.


This box lets you know what has changed with your action, if the operation was successful or not, if there is a process about to start (for example: replications, snapshots).

There are four types of notification boxes in Syneto HYPER:


- *Green boxes* - when an action is successful.



- *Red boxes* - when an action is not successful.



- *Yellow boxes* - when there is a connectivity issue or when the session expired, and the user will be redirected to the authentication page.



- *Blue boxes* - when an action is received and underway.



After several seconds, the notification boxes (except errors) will disappear, but you can also remove them from the screen manually. You can remove the notification boxes either one by one, by clicking the *Close* button on each box, or you can remove all boxes at once by clicking *Hide all notifications* button under all notifications on a page.

Notification emails

Notification emails will be sent when:

- Snapshot or replication fails
- Maintenance expires or becomes invalid
- New updates are available
- Network IP state changes
- Disk pool and datastores usage is above defined limits
- Operating system faults occur

ANALYTICS

Introduction

Analyzing the performance and the correct behavior of your Syneto HYPER is an essential task. There are many scenarios when you want to look at performance data. For example you may want to check to graphs to confirm your device is working within specified parameters. In other cases there may be some issues that are not obvious and you need to take a look at the analytics to pinpoint the issue.

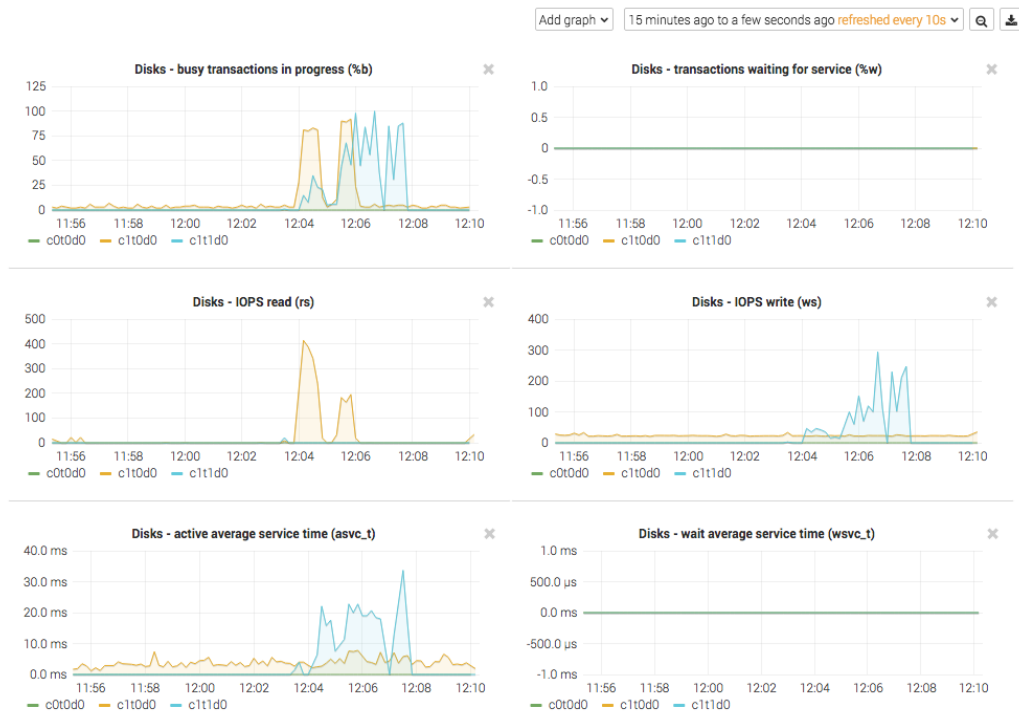
Low level analytics

The lowest level of analytics concern physical HDDs or SSDs.

Most of the time when drives fail, they do so immediately and fatally. In these cases it is easy to figure out what went wrong, and the drive is simply replaced.

However, in some cases, problems are more subtle. The drive may appear healthy and in working order, but it works at a much slower speed as it should. Even more, it may behave erratically only from time to time.

In order to analyze these cases, go to *Main menu* → *Analytics* → *Add graph* and select *Drives* → *Average service time*. Then add the *Drives* → *IOPS* graphics. Then add the *Drives* → *Transactions* graphics.



The most important indicator is the *Drives - active average service time (asvc_t)*. This represents the time required to process active transactions. But looking at it in isolation is not enough. There are legitimate cases for this value to be high. That is why we need other graphs as well. We need context.

In the images above we are interested in the blue lines on the graphs. We need to correlate *asvc_t* with *Drives - IOPS write (ws)* and with *Drives - busy transactions in progress (%b)*.

Usually a high *asvc_t* will determine a high *%b* as well. In other words, it takes a long time to process active transactions so the drive is busy.

The key indicator for a problem is the correlation with *ws*. A sudden burst of write will normally determine a higher *%b* and *asvc_t*. The images above depict a normal situation.

When a drive misbehaves, you will see high *asvc_t* and *%b* but unchanged or low *IOPS*. Most of the times another indicator is the rise of *Drives - wait average service time (wsvc_t)* alongside with *asvc_t* but no high *IOPS*.

In a simplified way, if you see increased time to process transactions, and the drive is busy, but it is not doing a large amount of reads or writes, something seems to be wrong. The drive may need to be replaced.

Mid level analytics

Drives are organized in pools on Syneto HYPER. The overall performance of a pool depends on the type of data it writes or reads, and its structure.












Similar to drives, you can analyze pool performance with several graphics and correlate between them.

The most important indicators are similar to drives: service time, IOPS, and transaction.

But pools are more complex than simple drives, so additional indicators must be taken into account for particular cases: *Pool* → *ZIL operations*, *System* → *ARC size and hit ratio* are two sets of graphics that can help correlate observed behavior with recorded data.

In this document we will present some typical cases only, for detailed performance analysis services contact Syneto customer support.

Add graph 15 minutes ago to a few seconds ago refreshed every 10s



Case 1 - I read a lot of data, but I don't see any read activity in analytics

Syneto HYPER employs read cache (called ARC in the graphs) all the time. By default all our products come with ARC placed in RAM, and L2ARC on an SSD drive.

NOTE: HYPER DR Play and HYPER 2000 Series products don't have L2ARC caching.

When you access recently or frequently used data, you will not see any read IOPS on the pool itself. You will see on the ARC or L2ARC graphics.

Case 2 - I write continuously to my Syneto HYPER, but pool graphics show write IOPS in spikes instead of a continuous line

All our products come with write acceleration (ZIL) preconfigured. The drive assigned to write acceleration, a fast SSD, takes over all the write operations and data from the clients. Every few seconds the system takes the data from the ZIL, optimizes it for writing, and writes it to the pool.

The pool write IOPS may show spikes instead of continuous writing.

Using a ZIL has significant performance benefits even with all flash pools.

Case 3 - I want to make sure my pool can read data fast enough from the drives

If you have an infrastructure that need to read very different data from a large pool, the read caches may not be effective for you. In these cases you need a the pool to be able to provide a certain read IOPS for you.

To run such a test, generate a large quantity of random data. Reboot the Syneto HYPER to clear any caches. Read the data generated before the reboot, for example by copying it over the network to another device.

Then follow the *Pools* → *IOPS read (rs)* graphics to confirm the pool is capable retrieving data as fast as the client can consume it over the network.

Case 4 - I have several disk pools and one is very slow

Correlate the pools IOPS graphs with *asvc_t*. If there is high traffic on the pool, see High level analytics. If there is normal traffic on the pool, but you have high *asvc_t* go one step down and check Low level analytics. Probably a drive is misbehaving or broken.

High level analytics

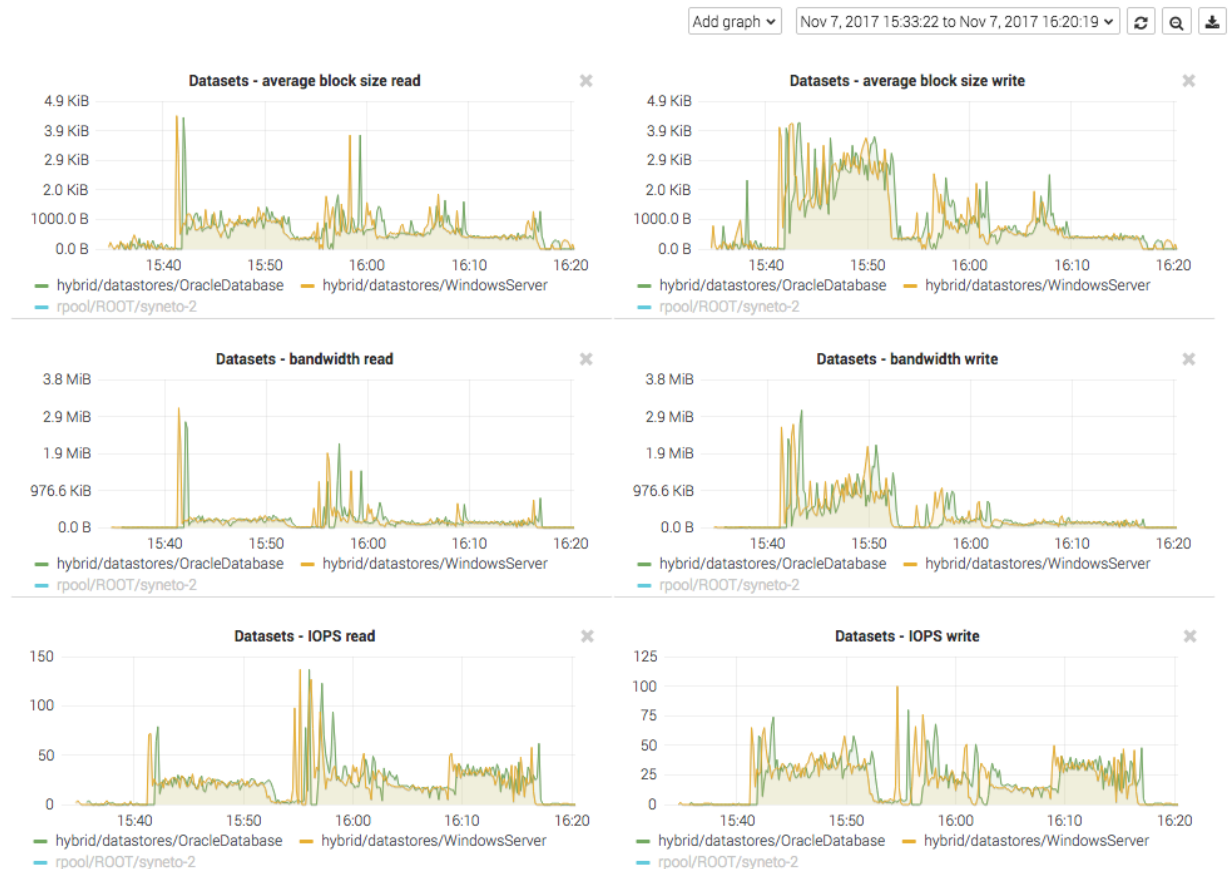
Until now we were talking mostly about hardware issues with the Syneto HYPER. However most issues are not related to the Syneto HYPERS functionality or performance. Rather they are problems with misconfiguration or misbehaving of some infrastructure components.

As these systems are very complex, interaction between virtual machines, or a wrong sharing of resources between virtualized components can lead to unexpected behavior.


For example one virtual machine can consume all available resources and make other virtual machines starve for resources as well, leading to general slowness of the whole infrastructure.

NOTE: For easy analyzes of such situations, we recommend creating one virtual machine per datastore.

From a storage perspective, a datastore is a dataset. And the analytics page has a set of tools to monitor dataset performance. The *Dataset* → *IOPS*, *Dataset* → *Bandwidth*, and *Dataset* → *Average block size* graphs will help you spot virtual machines monopolizing drive activity.



In the image you can see two virtual machines competing for dataset IOPS and bandwidth. If one of those virtual machine would take all IOPS, the other would have no resources left to operate properly.



Additionally you can add networking graphs and correlate virtual machine activity with network activity and ARC size and hits graphs.

In these images you can see that at about the same time the two virtual machines started to read data, the ARC cache size started to increase. It just happens that these two virtual machine were reading the same thing, and you can observe a slight increase in ARC hits for the same time period.

Miscellaneous analytics

Finally we will mention another set of graphics that can help pinpoint software issues with the SynetoOS operating system running on a Syneto HYPER product.

These are in the *Add graph* → *System* category. The most important graphs are *Load*, *CPU and memory*, *CPUs idle and wait*, and *CPUs system and user*.

You can check these to ensure that storage part of the Syneto HYPER is not overloaded.

SOFTWARE UPDATES

SynetoOS checks for updates automatically once every night and will notify the administrator via email when a new update is available.

You can see the current software version either by going mouse-over the Syneto | HYPERSeries tag, or by going to *Dashboard* → *About this machine* → *Software version*.

About this machine

Central

Linked with account `name.surname@company.eu` [Unlink](#)

Hardware


ID	7a5a3d15390e1ce03347193bb6f6027a
Serial number	SYN0001
Processor	4x GenuineIntel Intel(R) Xeon(R) CPU E5-1620 v4 @ 3.50GHz
Memory	8.0 GB
Data disks	5
Data storage (raw)	7.3 TB
Network	2 x 10 Gb/s

Software

Software version	4.0.0.2108
Last system boot	2017-11-02 13:05
Uptime	16:11pm up 1 day 3:06, 0 users, load average: 0.27, 0.24, 0.18

Close


Internet access is required to check for updates. If the storage needs a *HTTP proxy* to access the internet, configure the proxy on the *Network* → *Basic configuration* page.



Hostname Domain Gateway Proxy

Use a HTTP proxy?

Update




Hostname Domain Gateway Proxy


Use a HTTP proxy?

Host
110.77.159.25 ? The IP address or hostname of the HTTP Proxy server

Port
51225 ?

Update

When updates are available, the icon in the top right menu will be highlighted . Click it in order to start the update process. If a reboot is required, you will be told so by a tag next to the update option.



Current version
4.0.0.2108

Available updates

4.5.6.7 does not need restart


[View release notes](#)

Install updates

Update without restart

There are two types of updates. If the changes in the new version do not apply to the core of the operating system, the update will be installed without needing to reboot the machine.


1. Click *View release notes* to read the list of changes provided by the new version.
2. Click *Install updates* and confirm your action.




3. The update to the new version starts. Based on your internet speed, downloading the new version can take up to tens of minutes.



4. Wait for the update to finish.



5. If the software update was not successfully finished, you will be notified. This could happen if there are some issues with the connectivity. Check your internet connection and try again. Please contact support if you need any assistance.



Update with restart


New versions that contain changes which touch the core of the operating system will require a reboot. Versions requiring reboot are marked with a red warning label. We recommend the scheduling of this type of update outside working hours.

1. Click *View release notes* to read the list of changes provided by the new version.
2. Click *Install updates* and in the next dialog confirm you would like to continue.
3. The update to the new version starts. Based on your internet speed, downloading the new version can take up to several tens of minutes.
4. For the update to finish please *Restart*.
5. Wait for Syneto HYPER to restart.

TROUBLESHOOTING

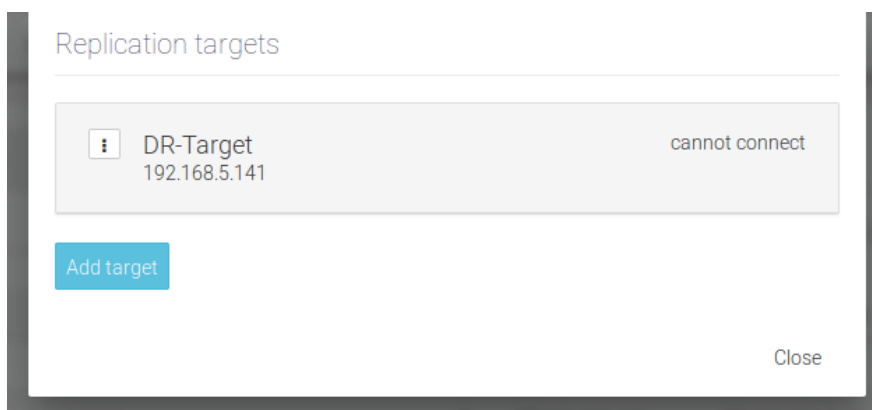
A snapshot replication fails

If you have received the following email or you have seen the *Last Replication* field highlighted with red on the *Protect and Replicate* page, it means a replication has failed.



First of all, check that the replication target is running. There are various reasons like power outage, hardware failure, system error that can prevent the proper functioning of the machine.

If the replication target is running, check that the machine which replicates can connect to it. Go to the replication targets list and you will see on the right the status, online or offline. If the status is offline, but the machine is up it means there is a network connectivity problem between the two machines.



If it can connect to it, check the available space on the machines (both the primary and disaster-recovery). This can be done by opening the *Pools* page.

Name	Size	Used space	Scheduled check
✓ hybrid	99.5 GB	47.6 GB	in 25 days
✓ OS	79.5 GB	4.3 GB	in 25 days

hybrid Export Destroy

Data	100.0 GB	Expand	<div style="border: 1px solid #ccc; padding: 5px;"> <p>47.6 GB of 99.5 GB used</p> <div style="width: 47.6%; height: 10px; background-color: #4a7ebb; margin-bottom: 5px;"></div> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 5px;"> <p style="text-align: center;">Healthy</p> </div>
Write acceleration		Manage	
Read acceleration		Manage	
Spares		Manage	

Integrity check Check now

Last check	3 days ago — All fine
Scheduled check	in 25 days (Fri, Aug 10, 2018 22:59)

If it has enough space, remove the replication target and try adding it again. Wait for the next scheduled replication or click on *Run* for the desired schedule to start now.

If it still fails, please contact Syneto customer support.

Cannot enable customer support

⚠ Could not enable remote support ✕

Host support.syneto.eu:22 is not reachable. (Reason: connection failed)

12:34:24 - 13 Nov, 2017


✕ Hide all notifications

If you can't enable remote access for Syneto customer support, it means there is a network connectivity problem. First of all, check the internet connection. If you are on the same network as the machine, you can simply try to connect to a website.

If the internet connection seems to be working, check the network configuration of the SynetoOS. Verify if you have the correct domain name, DNS server and gateway set. Also, if your network uses proxy be sure you have specified the right address of the proxy server and port. Check the [network configuration section](#).

If network configuration checks out, contact your network administrator and ask them to verify the network firewall.

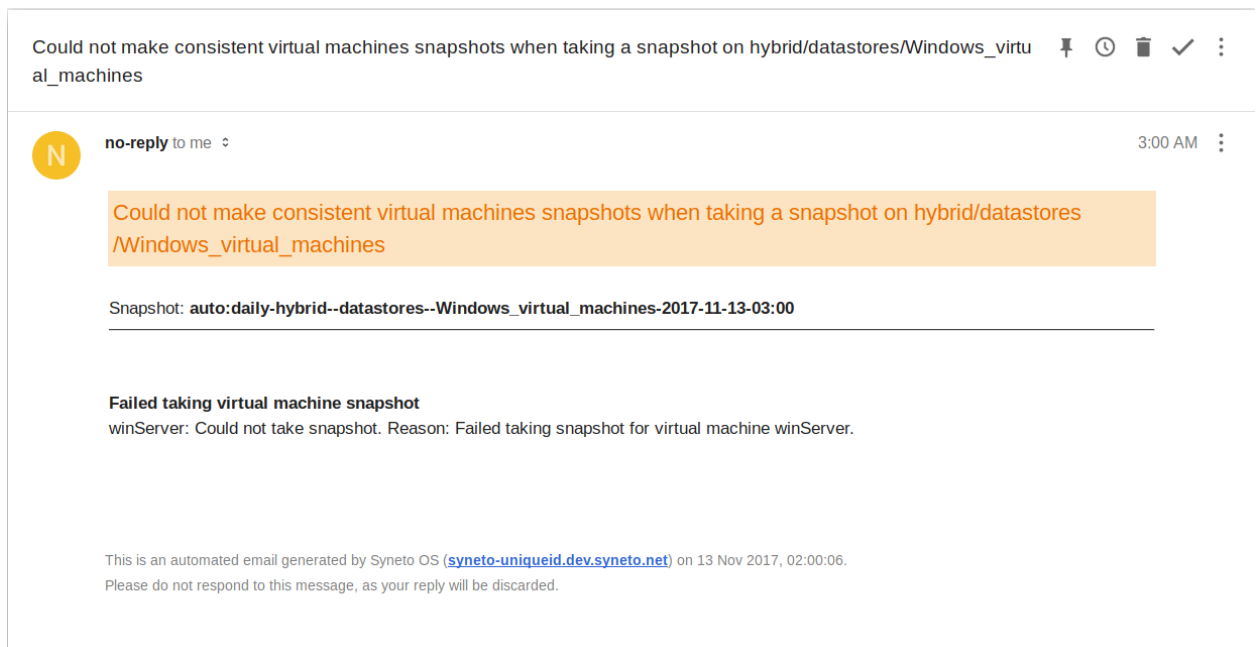
Can't connect to Active Directory



If a notification error is displayed when trying to connect to Active Directory, or if the SynetoOS lost the connection ensure the following network configurations are set in the right way:

- **The DNS must be configured** to a DNS server part of the Windows domain. In most of the cases the **AD server also acts as a DNS server**
- **The Domain Name** of the SynetoOS must be identical with the AD domain name.

Can't create app-consistent/live snapshots because existing VMware snapshot




There are two situations that can cause this problem.

First one, if the user creates VMware snapshots intentionally. In this case, if the user wants to manage the virtual machine snapshots himself, SynetoOS will not interfere. To not receive emails with the warning, select crash-consistent state for the virtual machine snapshots from *SynetoOS Management Interface, Protect and Replicate* page.

Second one, if a VMware snapshot fails it will remain attached to the virtual machine. In this case, it will have to be deleted manually. Enter the *vSphere Management Interface*, select the virtual machine and right click on it. A menu will appear with the *Snapshots* sub-menu. Place the mouse over (don't click on it) and another menu will appear. Click on *Manage snapshots*.

A dialog will open listing the virtual machine's snapshots. On the upper section of the dialog, there are a few buttons. Click on *Delete all*. This will remove the snapshots and SynetoOS will be able to make app-consistent/live snapshots of the virtual machine.



After reinstall, data pool cannot be imported

This affects **Syneto HYPER 2000 Series** and **Syneto HYPER DR Play Series**. The write acceleration drive is a virtual one, and after a reinstall it will be recreated. As a consequence, the pool cannot find the old drive.

Connect to CLI, switch to unsupported shell by running this command:

```
un s
```

Import the pool without write acceleration:

```
zpool import -m hybrid
```

By default the pool is named *hybrid*. If you created your own pool, replace *hybrid* with the name of the pool.

Go to *Secondary menu* → *Pools* and select the data pool (hybrid). Go to Write acceleration on Details view and click *Manage*. From the *Write acceleration drive group* remove all drives. Click *Update*. Go to *Manage caches* again and select the drive that has 8GB (usually c1t1d0) and then click *Update*.

Update pool cache

hybrid

Read cache drives group

+ 1.8 TB

Write acceleration drives group

+ 8 GB

Cancel

Update

Error when creating a datastore, share, or volume

When you try to create a datastore, share, or volume, you may see an error similar to this:

A screenshot of a red error message box. The text inside reads: "Could not create filesystem with path /shares/smb/smb-share-test." followed by a red 'x' icon. Below this, the timestamp "17:45:56 - 6 Nov, 2017" is displayed.

Make sure there is a data pool imported. Go to *Secondary menu* → *Pools* and confirm you can see data pools (other than the OS pools). If no data pool is present, please create a new one or import an existing one.

If the problem persists, contact Syneto customer support.

Datastores are not accessible on vSphere management interface after Syneto HYPER IP change

If this happens, go to *Main menu* → *Datastores*, unmount and mount again the affected datastores.

Datastores are not accessible on vSphere management interface after ESXi host IP change

If this happen, go to *Main menu* → *Datastores*, unmount and mount again the affected datastores.

I don't know if Syneto Central is accessible

My Syneto HYPER was activated before, but I want to know the Syneto Central connectivity status.

Go to *Main menu* → *Dashboard* → *About this machine*. The first section in the dialog is called *Central*. If something is wrong it will be shown there.

About this machine

Central


No connection with Central
Linked with account **user@my.domain**

[Unlink](#)

Dialog appears asking for the new local ESXi password

If a user changed the ESXi password from the vSphere management interface, in SynetoOS a dialog will ask to introduce the new ESXi password. If a wrong password is provided, after 5 failed attempts the account will be locked for 60 seconds.

Please note that the recommended way to change ESXi passwords is using the SynetoOS web interface.



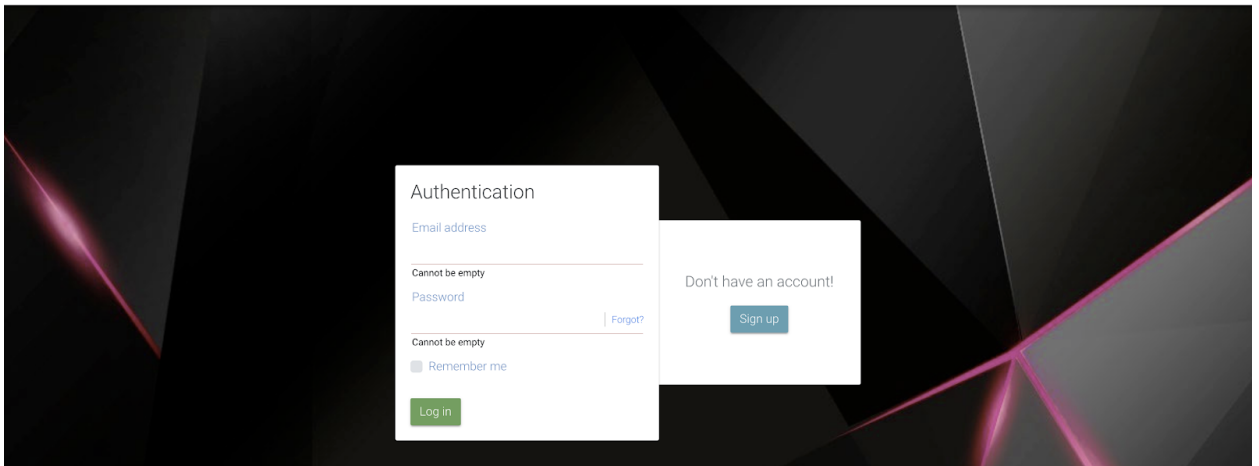
The screenshot shows a dialog box titled "Update hypervisor password". Inside the dialog, there is a yellow message box that reads: "It looks like you have changed the local hypervisor password. Please update the hypervisor password in Syneto OS too." Below the message box, there is a label "Password" followed by a horizontal input field. At the bottom right of the dialog, there is a blue button labeled "Update".

APPENDIX A: CREATING A CENTRAL ACCOUNT

In order to create a Central account, you must follow the link to the Central Portal (<http://central.syneto.eu/>).

On the page that opens, go to *Don't have an account* → *Sign up*. A form will open, where you have to configure your profile.

syneto CENTRAL



Type in the email address of the person who will be using the product and choose a password. Also insert a backup email address. Complete the name and phone number for the person in charge with this account, but also complete the information fields for your company.

WARNING: Please make sure that you completed all fields correctly. Failure to do so might prevent you from obtaining full support on your product.

Welcome,
your account is **almost ready**.

We just need a few details about yourself and your company, to properly activate your product. To obtain full support, please make sure all data is correct.

1. Account	Email address	Company email address
	name.surname@company.eu	company_name@company.eu
	Password	Re-type password

2. Contact details	First name	Last name
	Name	Surname
	Phone number	
	00303203556179	
3. Company	Name	Market sector
	Company	Education
	City	Country
	Bologna	Italy
	Address	Number of employees
	Mirafiori 25	1-20 employees

SIGN UP


After all the required fields are completed, press *Sign up*. Your Central account is now ready. You can start activating your products.

For more information about activating a product with Central, please go to [Activate using a Central account](#) chapter.



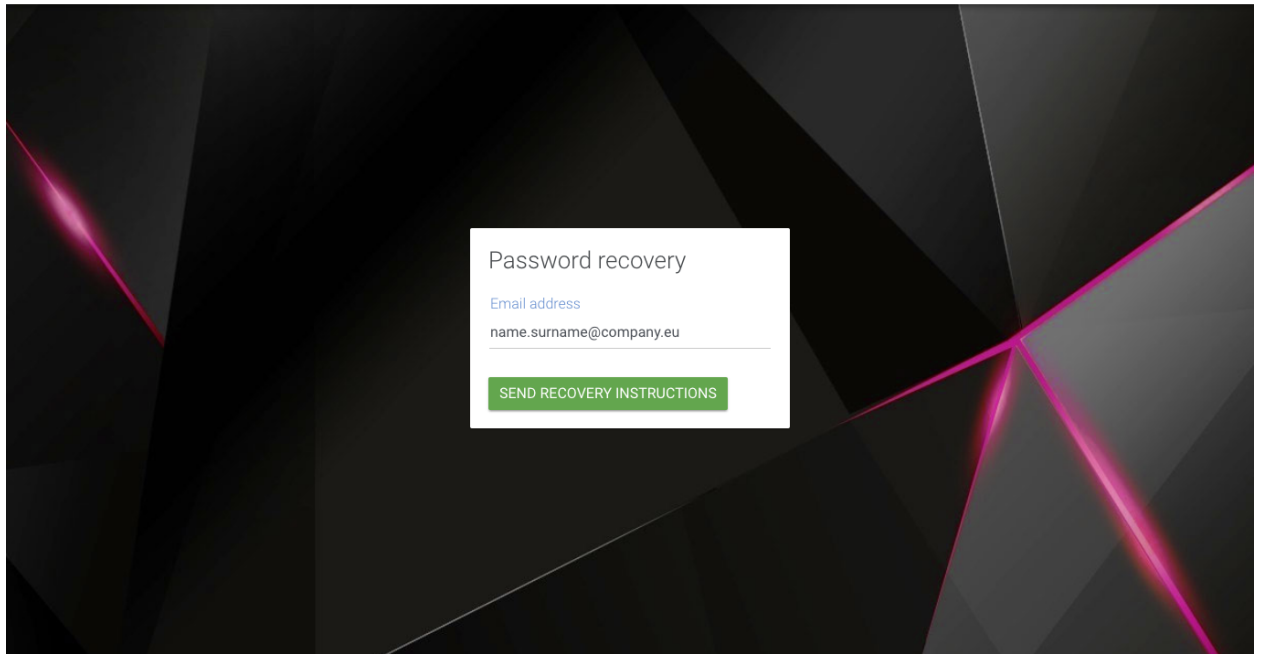
There are no machines in your list yet. Go on Hyper and activate your first one.

After you start activating products, they will all show up on the page with information about their *Name*, *Support type*, *Expiration date*, *Hardware ID* and *Model* .

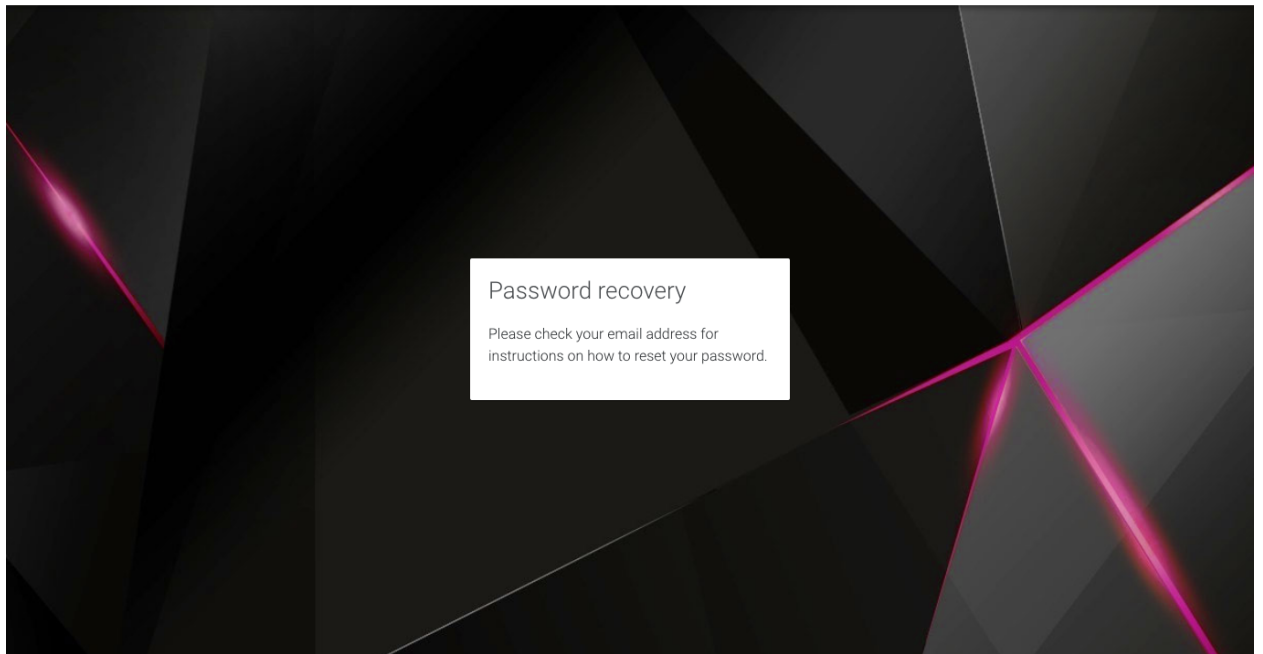


If you forgot your password to the Central account, from the login page, go to *Forgot? Link* . A new page will open , where you insert your email address and click on **SEND RECOVERY INSTRUCTIONS**.

syneto | CENTRAL



syneto | CENTRAL



An email with password recovery instructions will be sent to the provided email address.

no-reply@syneto.eu
Central: password reset request
To: luciana.baboniu@syneto.net



You are receiving this because you (or someone else) have requested the reset of the password for your account.
To reset the password [click here](#).

If you did not request this, please ignore this email and your password will remain unchanged.

Follow the *click here* link in the email, and you will be redirected to a new page, here you can insert a new password. Make sure you confirm it by inserting it twice and then click on CHANGE PASSWORD. You can now log in to your Central account with the same username and the newly changed password.

syneto | CENTRAL


Password reset

New password


Re-type new password

CHANGE PASSWORD

APPENDIX B: ACCESSING CUSTOMER SUPPORT

With Syneto HYPER it is very easy to get in touch with support. There is a  *Help* icon in the *Actions* bar that you can follow.

Here, you can *View support page*, *Open a ticket*, or *Enable/Disable remote access*. From here, you can also *Generate support bundle*.



The screenshot shows a 'Help' interface with a top navigation bar containing icons for menu, download, notifications, and help. Below the bar is a grid of six support options:

- Maintenance plan**: Type: Advanced ⓘ, Active period: 2018/05/11 - 2023/05/10. Button: Refresh status.
- Open a ticket**: Solve any situations fast by opening a support ticket. Button: Open.
- Remote support**: Enabling remote access will allow Syneto's technical support team to connect and troubleshoot any existing issues. Please make sure you have a working internet connection before enabling remote access. Button: Enable remote access.
- Support bundle**: Please download and provide the support bundle archive to Syneto's technical support team. Button: Generate support bundle.
- Support**: Learn more about how Syneto support works. Button: View support page ↗.
- Serenity**: Serenity service is not activated for this machine. To activate Serenity please contact your supplier. Serenity frees your inbox of the machine related issues and reports them directly to Syneto Support for further investigation. Button: Learn more ↗.

Because your product is registered with Central, when opening a ticket, the data will be already filled with the information provided in your *Central profile*. All you have to do now, is *Describe your request*.

New ticket

Please make sure you have a internet connection before opening the ticket. Opening a ticket will submit anonymously basic diagnostics, analytics and usage information. This information will help our engineers to diagnose and solve any issue faster.

Name
Luciana Baboniu

Email address
luciana.baboniu@syneto.eu

Company name
SYNETO

Summary
Two errors popped out!

Describe your request
My alerts icon is red, and two errors showed up...

Cancel Send

Enabling remote access, will allow our support engineers to connect and troubleshoot any existing issues. Make sure you have a working internet connection before enabling remote access. You can disable the remote access anytime you want.

NOTE: Since SynetoOS 4.4, the remote port numbers are no longer provided by enabling remote support access. You no longer need to copy support numbers to clipboard and paste them to your open conversation with the support engineers.

Remote support

Remote support access is enabled.

You can also access support when you see *Alerts*, by simply following the *Get help* links you can find on the *Alerts* page.

NOTE: Please note that the Support services differ with the Maintenance type you have.

Contacting Syneto Customer Support

Via the web	https://syneto.eu/support/
Via e-mail	support@syneto.eu
Via phone	+39 0809 080 522
Via customer support portal	https://helpdesk.syneto.eu

NOTE: In order to have support access, you need a Maintenance plan and a Central activated product.

Support for VMware ESXi hypervisor issues

Syneto Hyper Series products use VMware vSphere as the hypervisor. Support for issues related to the ESXi hypervisor, will be escalated by Syneto's support team directly to VMware, and Syneto's support team will manage the issue through the joint TSANet support platform. Please be aware that Syneto's support team will require a valid and current VMware support contract in order to fully support you on issues related to VMware products.

APPENDIX C: SERENITY SERVICE


There is a new service, for proactive support: Serenity.

Serenity relieves the customer of the burden and worries associated with receiving emails in case of failures and warnings, and contacting Customer Support. When Serenity is enabled, our support team will automatically be notified of the events (failures and warnings) on your system. We will connect on the machine via a remote support tunnel to investigate and fix the problems.

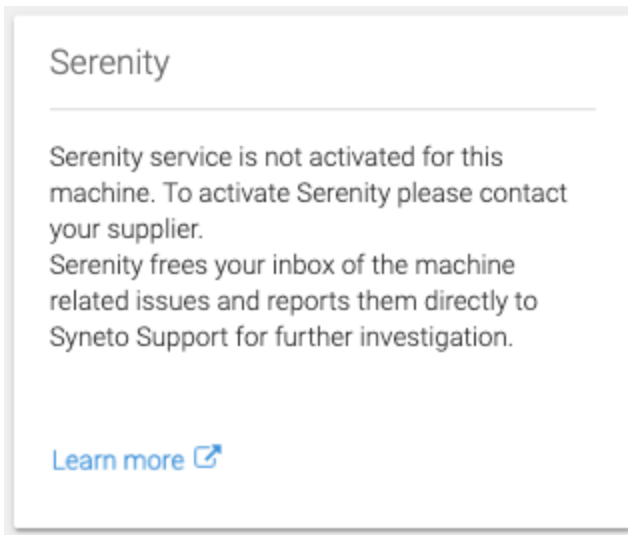
You can check the availability of the Serenity service from two locations: The Alerts page and the Help page.

If your support plan does not include the Serenity service, then you will find that it is inactive.

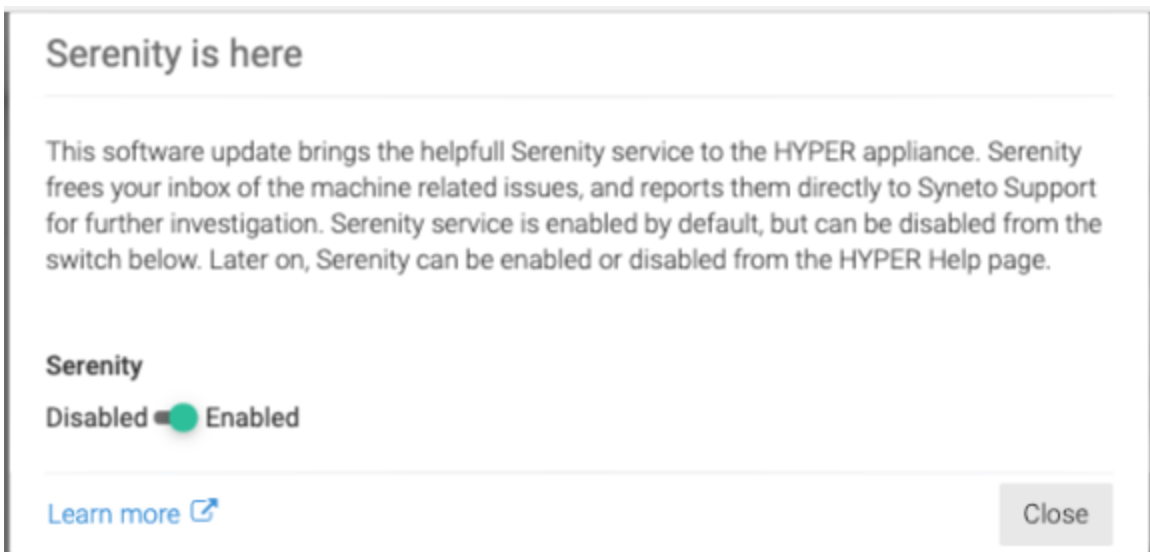
In the Alerts page:



In the Help page:




After you purchase a support plan that includes Serenity, a popup notification will appear, asking you whether you want to enable the service or not. By default, it will be enabled.




If you have a support plan that includes the Serenity service, but you wish to be notified via email for errors and warnings, you can opt out of using it. If you do so, you will be able to see in the Alerts page and in the Help page that the service is available, but is disabled.


In the Alerts page:



In the Help page:



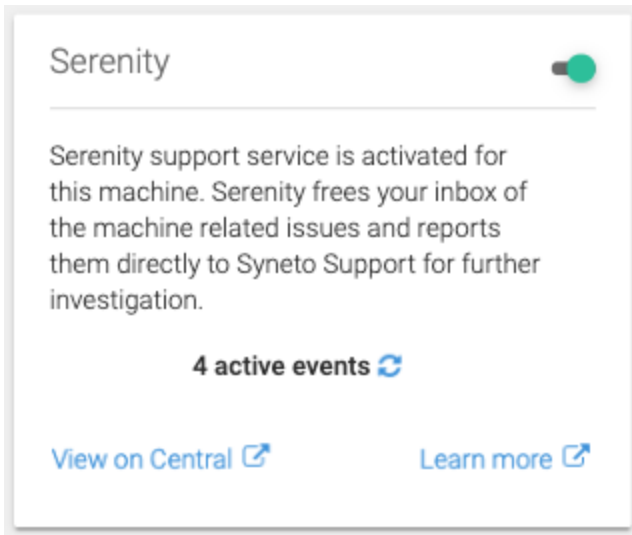
Disabling Serenity will also disable Remote support:



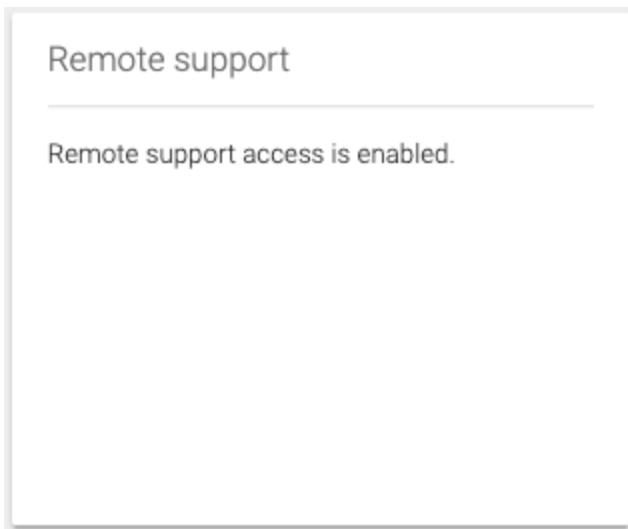
You can use the toggle to enable Serenity. Once you do so, you will see the Alerts page changed:




And you will see on the Help page as well that Serenity is now enabled:



By enabling Serenity, the Remote Support access will be automatically enabled:



If you choose to disable Serenity, you will be asked to confirm.



After your support plan expires, if you had purchased Serenity, the errors and warnings that will occur will cease to be sent to Serenity for proactive support. Instead, you will receive emails.

APPENDIX D: GRANULAR FILE RECOVERY

The File recovery feature allows you to recover individual files from snapshots.

Below are a few technical details that might be helpful to you to be aware of, if you choose to use the feature:

- In order to use the file recovery functionality, you need to keep (i.e. not delete) the *SynetoFileRecovery* virtual machine from your ESXi.
- All the operations, from browsing the snapshots (listing filesystems, listing files inside a directory) to generating the recovery artifact (zip, ISO), are being performed in the background, by the *SynetoFileRecovery* virtual machine.
- You should expect the browsing of a snapshot to take several seconds. As for the recovery process, it may take minutes or more - depending on the number of files chosen for recovery and their size.
- If you choose to recover a single file, it will be only that single file put in the zip archive / ISO. If you choose to recover a single (non-empty) directory, it will be the contents of that directory stored in the zip / ISO. If you choose multiple files from multiple directories, then the parent directories will be stripped in the zip / ISO as much as possible. E.g. If you choose to recover the files:
 - /home/admin/a.txt

- /home/admin/subdir/b.txt
- /home/admin/subdir/c.txt

Then in the zip / ISO you will have the following files & directories stored:

- a.txt
- subdir/b.txt
- subdir/c.txt
- When you recover files for the first time, SynetoOS will create a datastore on the ESXi (which will be visible from ESXi UI, but not from the UI of SynetoOS). A quota will be configured on this datastore, of maximum 10 GB (could be less, depending on the available space). This means that:
 - For the first time you choose to recover files, you will have to make sure that the total size of the files does not exceed 10 GB.
 - After you successfully recover files, you should click *Cleanup* in the task view so as to free space on the recovery datastores. Otherwise, the recovery datastore will slowly fill up and you may no longer be able to recover files due to insufficient space.
- Daily, at 02:12 am, SynetoOS will perform a cleanup operation, which will automatically remove old artifacts from the recovery datastores; it will also remove task entries from the task list.
- When choosing a recovering method, there are two options available:
 - Attach ISO: useful when you want to restore an older version of a file (or directory) back onto the same virtual machine. E.g. If you have a Windows virtual machine, and you have a file on desktop (say, "important.docx") which had become corrupted (or removed), you could browse back onto a previous snapshot of the virtual machine, and pick that file for recovery. You would attach the ISO onto the Windows virtual machine, from where you would copy back the older (good) version of the file.
 - Download (zip archive): useful when you want to retrieve data from an old version of a file (or directory), and you would use that data from your work computer (e.g. your laptop). For example, you would recover a file "email-contacts.txt" by downloading it as an archive, if you need to email a certain person. As another example, you might want to download an archive with pictures to view on your laptop.
- When browsing files from a snapshot, SynetoOS will display maximum 250 items (files & directories) for each directory. In other words, if you have a directory in your snapshot with 3000 files, when browsing, you will only be able to see 250 of them. However, if you choose to recover the directory that contains them, all 3000 files will be recovered. This limit (of 250 files) is not directly configurable (i.e. you cannot change it from the SynetoOS UI) - it is a discretionary value meant to:
 - Simplify the viewing of files - it might be difficult for the user to look for individual files in a directory of thousands or dozens of thousands of files.
 - Limit the memory usage on the web UI
 - Reduce wait time - i.e. it is more rapid to read 250 entries from a directory than to read thousands or dozens of thousands, especially since the directory is being accessed via a network share.

- The File recovery functionality can browse and recover files from the following filesystems: EXT2, EXT3, EXT4, XFS, NTFS, FAT32.
- Only regular files (and directories) will be recovered. The following will not be recovered:
 - Symbolic links
 - Block devices
 - Character devices