



SynetoOS User Guide

Version 4.2

July 2018

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](#)
- [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](#)

[Clean up when you are done](#)

[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](#)

[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](#)

[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](#)

[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](#)

[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](#)

[SYSTEM CONFIGURATION](#)

- [Networking](#)
 - [Ethernet interfaces](#)
 - [Missing or broken interfaces](#)
 - [Date/time and timezones](#)
 - [Password change](#)
 - [Changing local ESXi and SynetoOS passwords](#)
 - [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](#)

DEFAULT USERNAMES 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: syneto123

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 **4.2.0**.

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/6.5/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:

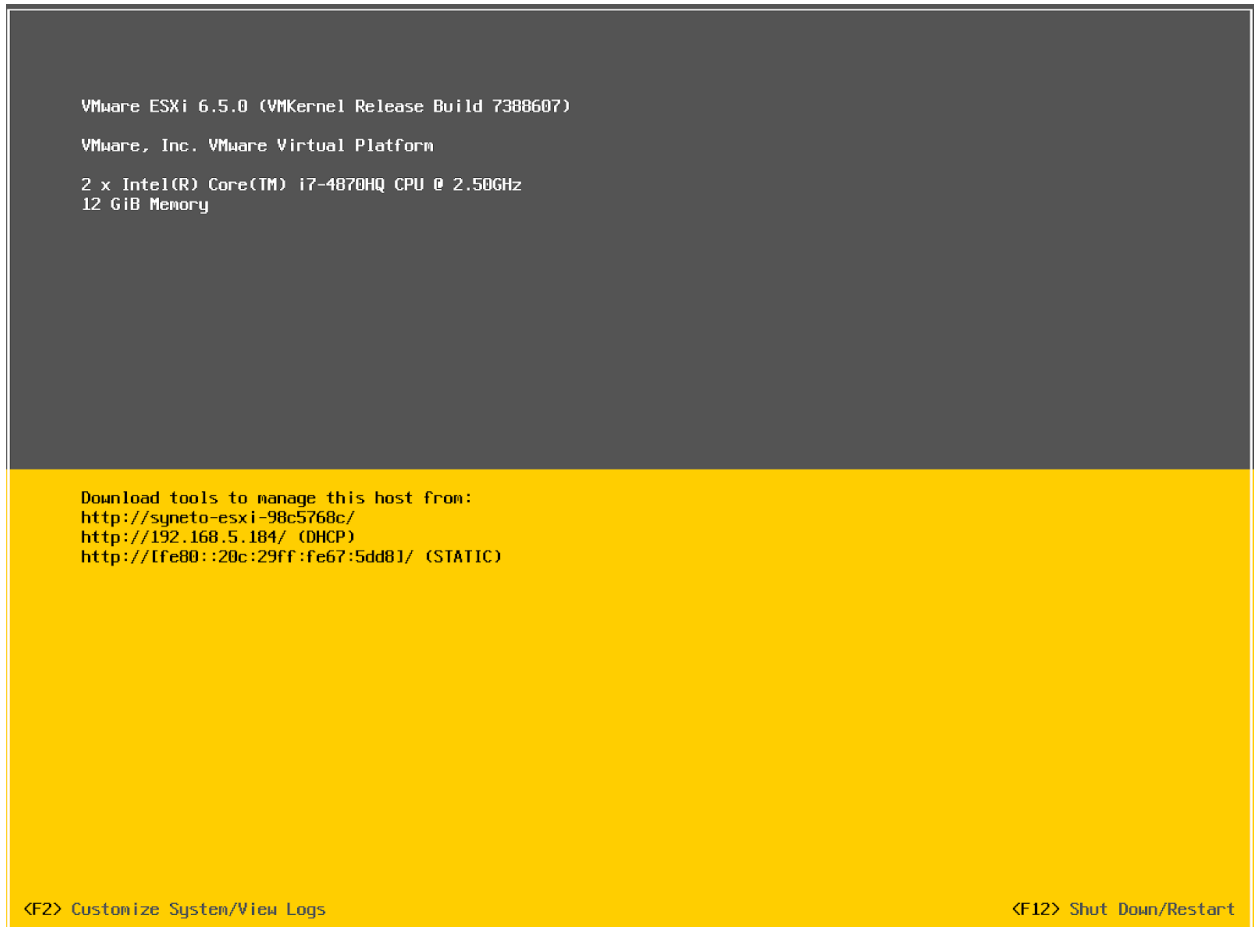
- Allocate three IPs in the **management network**, for:
 - IPMI (machine management)
 - ESXi (hypervisor management)
 - SynetoOS (controller VM)
- **DNS server** IP address - you can configure between one and three DNS servers in the management network.
- **Gateway** IP address - you will need this 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 6.5.0 build-7388607 http://www.vmware.com
Copyright (c) 2007-2017 VMware, Inc.

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

SynetoOS management : https://192.168.5.27/

syneto-esxi-98c5768c.dev.syneto.net login:
```

vCenter and ESXi hosts credentials

HYPHER 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:

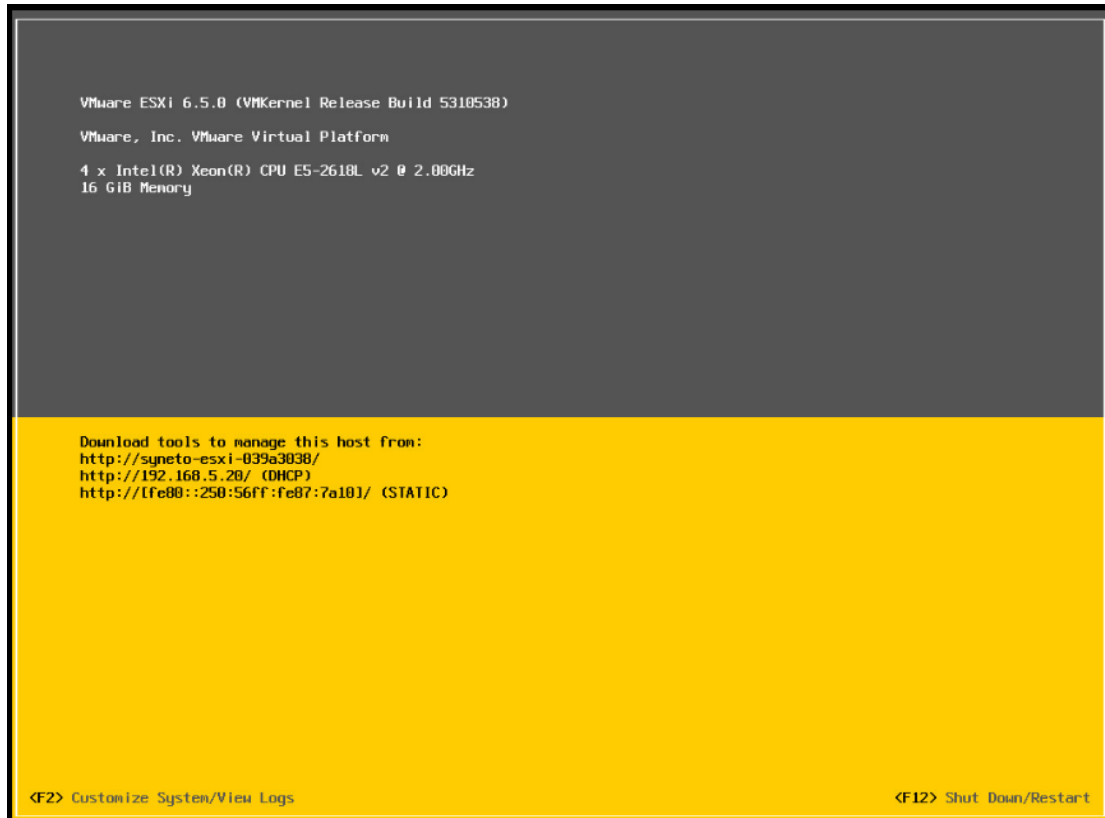
1. Get physical access to HYPER using monitor and keyboard
2. Activate the ESXi web console
3. Configure basic networking of the management VM
4. Activate machine with online Syneto Central service
5. Get access to the HYPER Management VM web console.
6. Activate IPMI remote system management interface

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: `syneto123`). 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 6.5.0 (VMKernel Release Build 5310538)
VMware, Inc. VMware Virtual Platform
4 x Intel(R) Xeon(R) CPU E5-2618L v2 @ 2.00GHz
16 GiB Memory

Download tools to manage this host from:
http://syneto-esxi-039a3038/
http://192.168.5.20/ (DHCP)
http://[fe80::250:56ff:fe87:7a10]/ (STATIC)

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

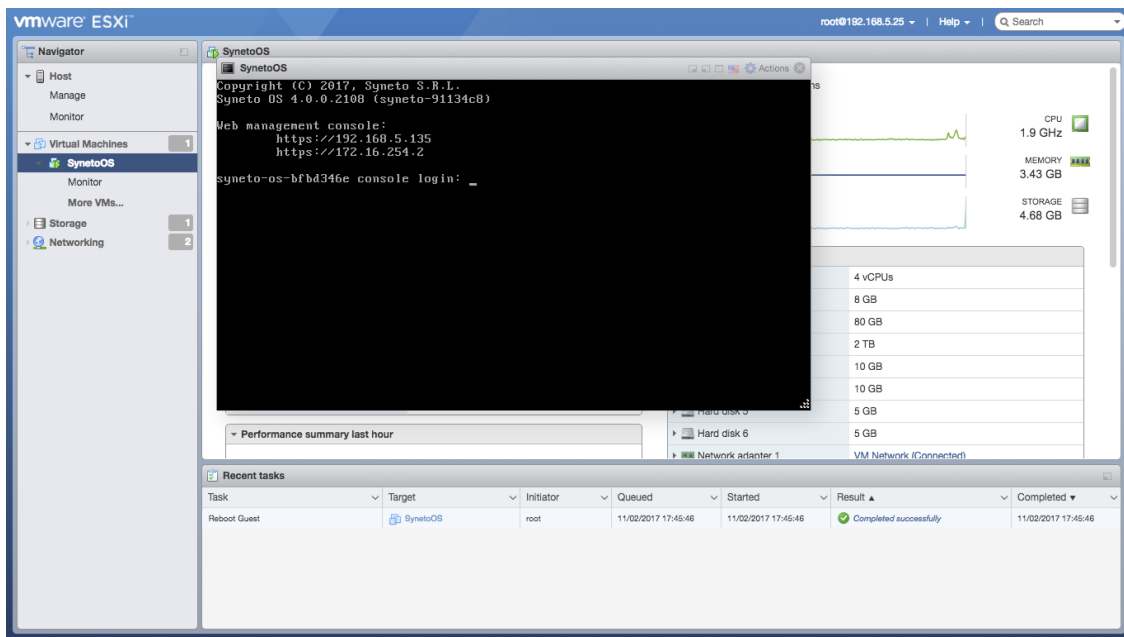
On the ESXi console, you can see the IP address to the ESXi UI. Insert the address in a browser (of your choice) and log in to the UI with the default credentials (default username: `root`; password: `syneto123`).

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).

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.

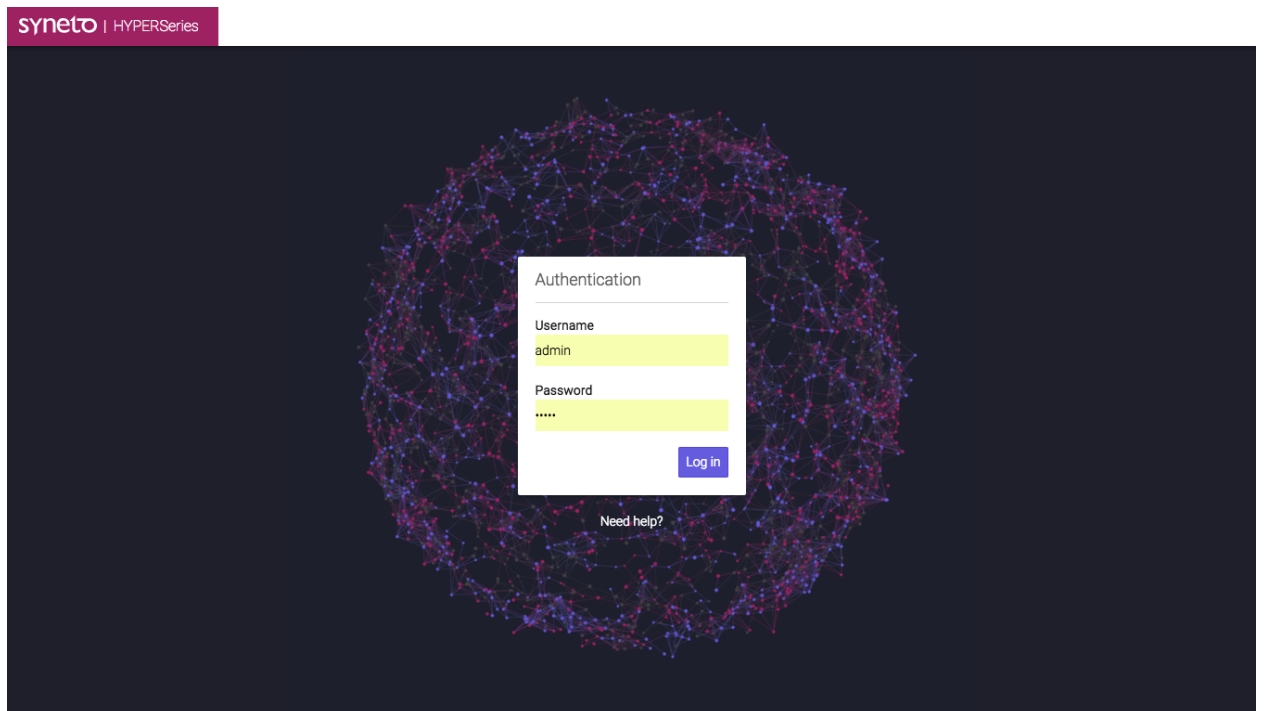
Upon receiving your Syneto HYPER, it will already be registered with Central, but in order to work on it you must create a Central account. Please read [Appendix A](#) to learn how to create a Central account.

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

Go to any browser you want and write the management IP address using the default credentials (user: `admin`; password: `admin`). 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 HYPER and note the management IP address shown during boot. See bottom right in the image below: `BMC IP:`

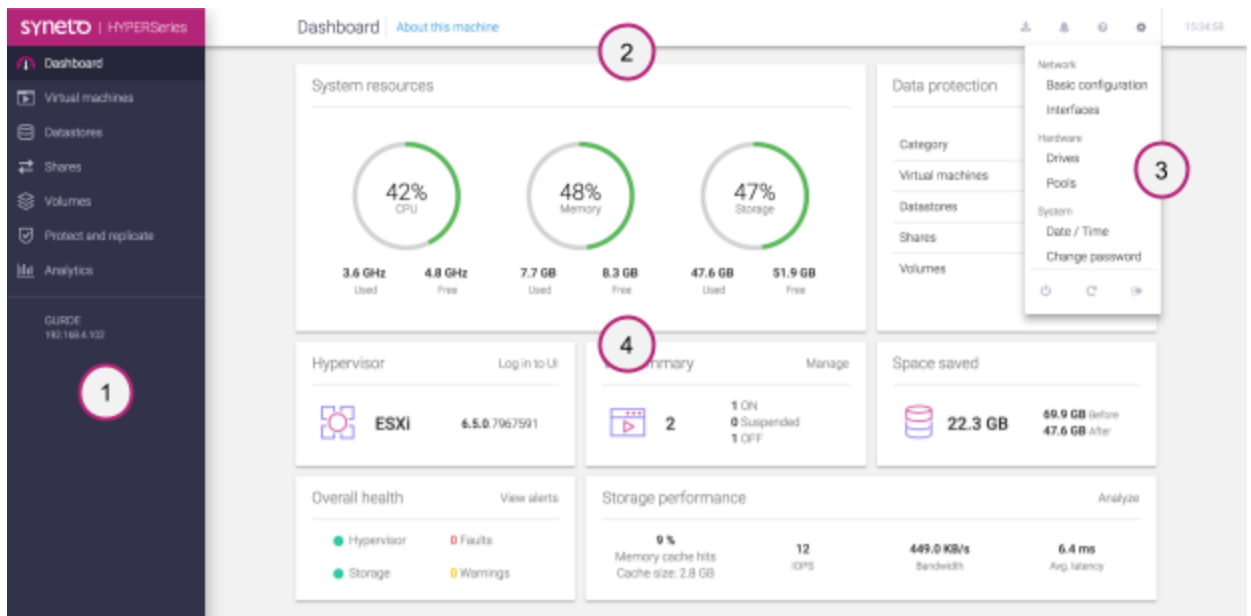


Default credentials

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

Management interface main components



These are the main components of the management interface:

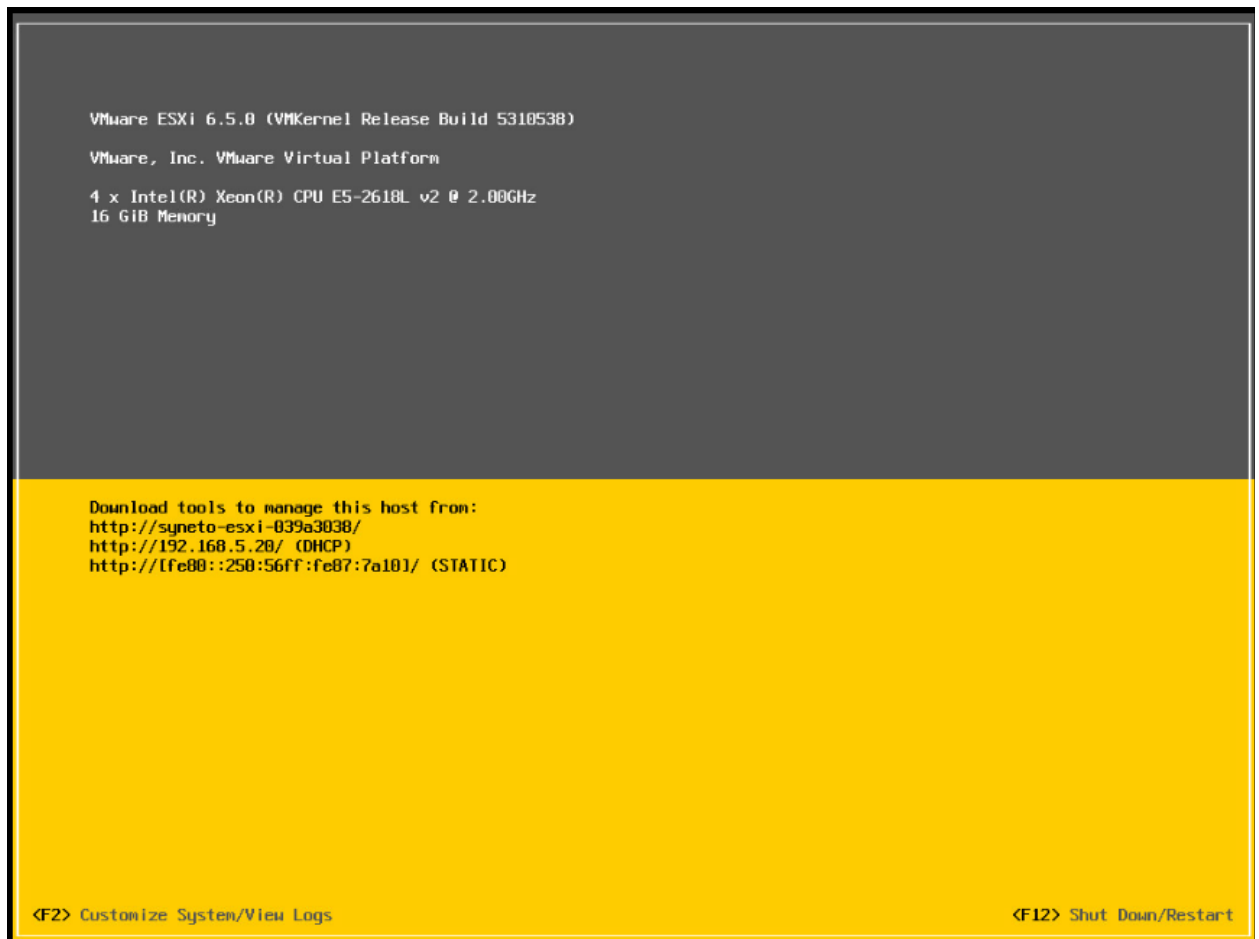
1. **Main menu** - contains the most important items related to the Syneto HYPER management. Expect to use it frequently.
2. **Action bar** - presents actions related to the current page, and a set of notifications on the right for update, alarms, expired license, missing central connection. Finally 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 6.5.0 (VMKernel Release Build 5310538)
VMware, Inc. VMware Virtual Platform
4 x Intel(R) Xeon(R) CPU E5-2618L v2 @ 2.00GHz
16 GiB Memory

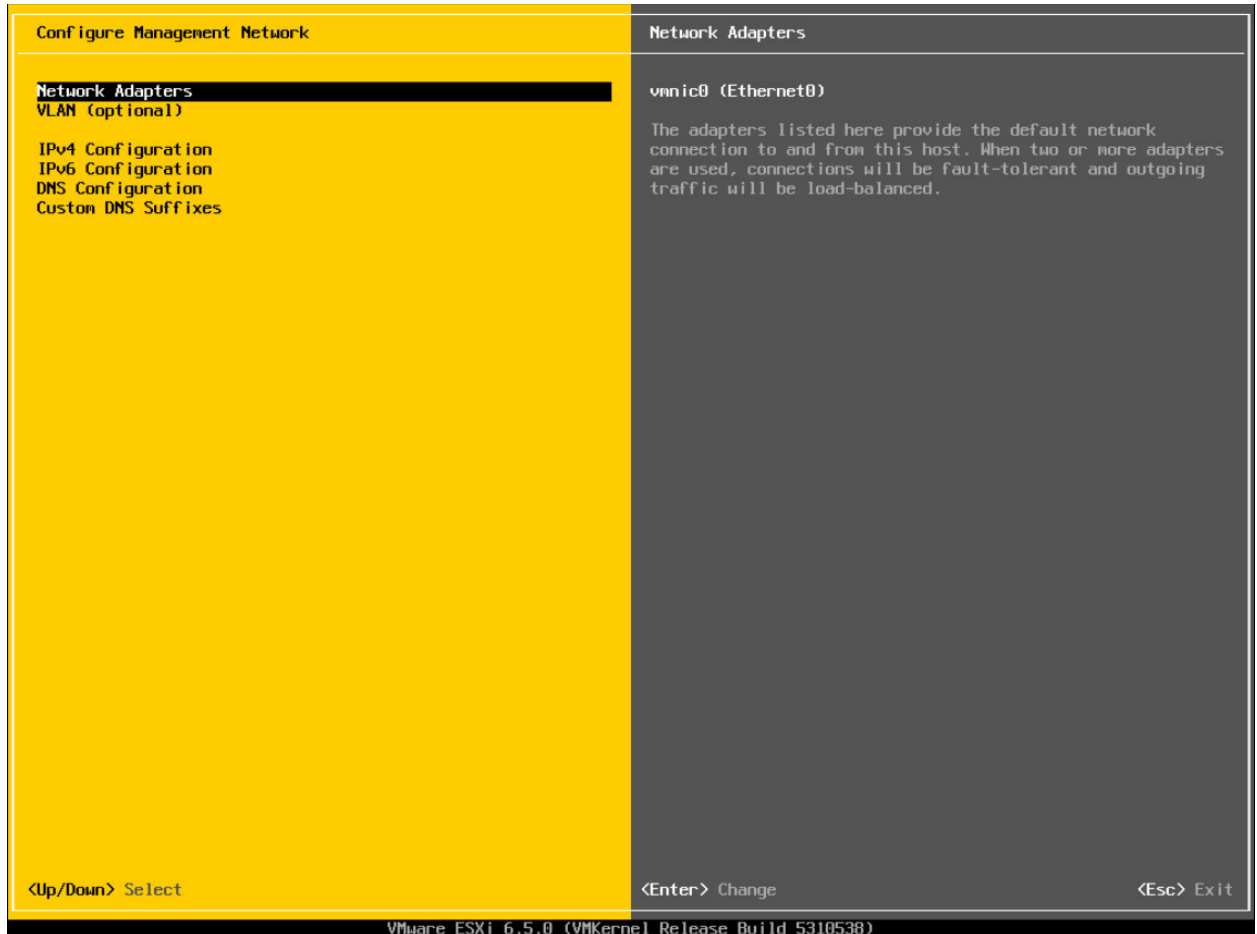
Download tools to manage this host from:
http://syneto-esxi-039a3038/
http://192.168.5.20/ (DHCP)
http://[fe80::250:56ff:fe87:7a10]/ (STATIC)

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

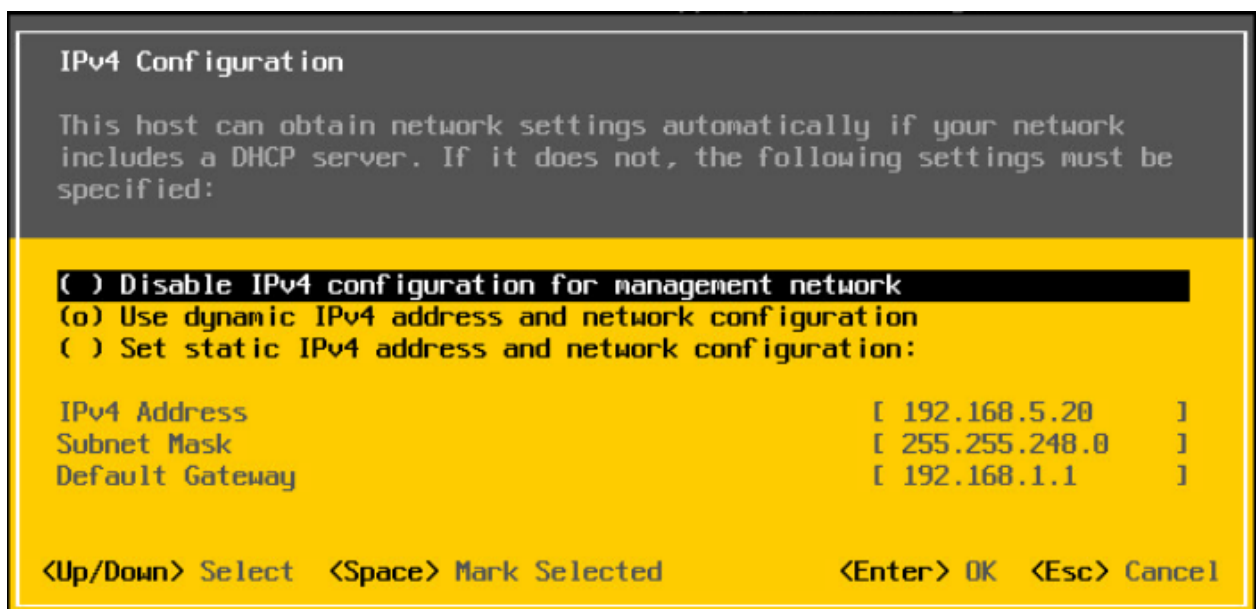
Press F2 and then enter the username and password of the ESXI. A menu with different actions will open.



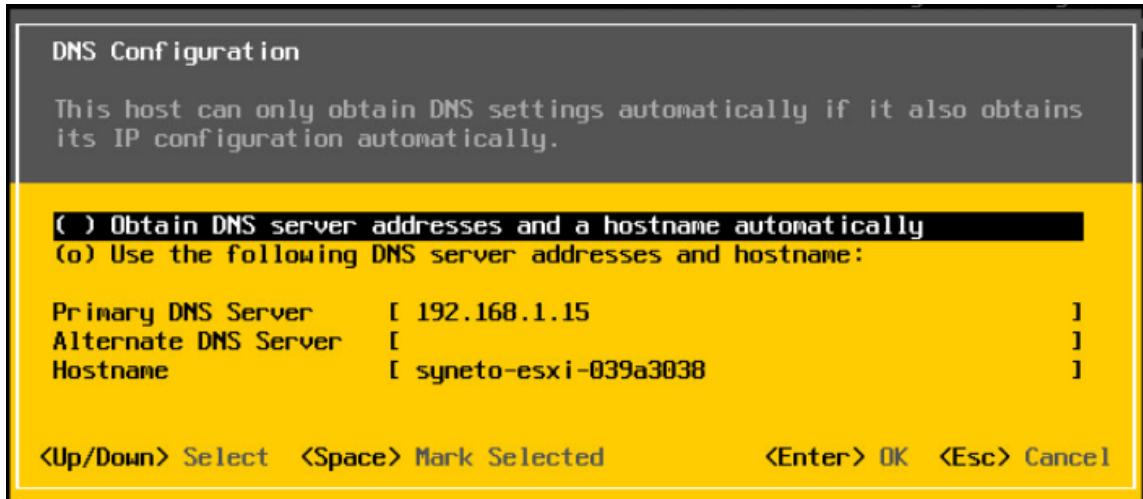
Select Configure Management Network and press Enter. A menu with network configuration options will open.



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.

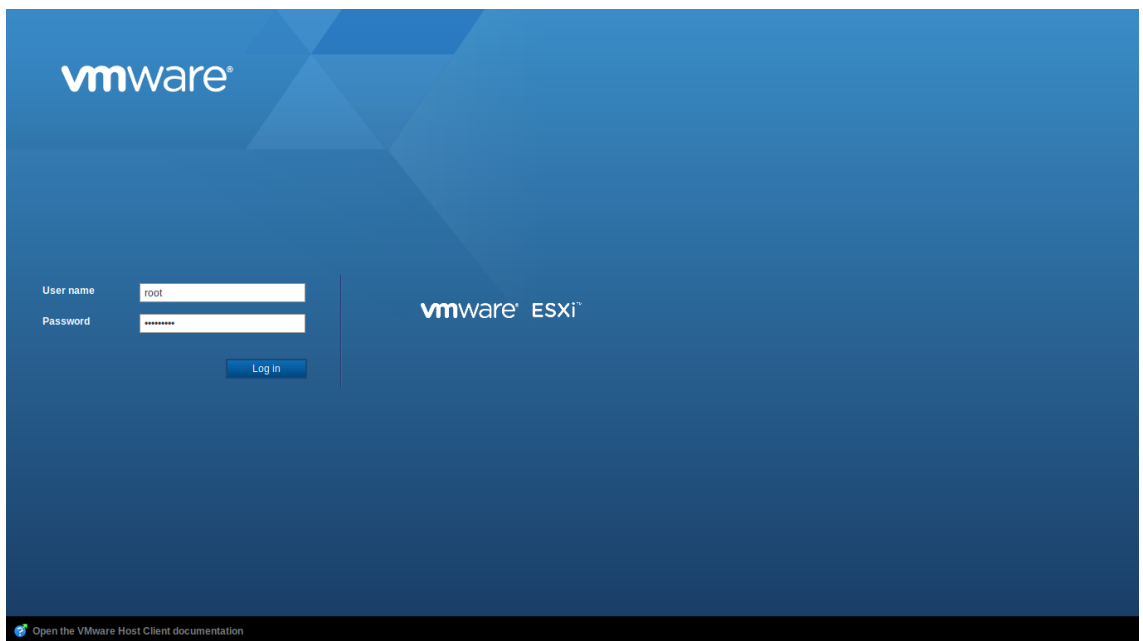


The screenshot shows a terminal window titled "DNS Configuration". The text inside reads: "This host can only obtain DNS settings automatically if it also obtains its IP configuration automatically." Below this, there are two options: "() Obtain DNS server addresses and a hostname automatically" and "(o) Use the following DNS server addresses and hostname:". The second option is selected. Below the options, there are three fields: "Primary DNS Server" with the value "192.168.1.15", "Alternate DNS Server" which is empty, and "Hostname" with the value "syneto-esxi-039a3038". At the bottom, there are navigation instructions: "<Up/Down> Select", "<Space> Mark Selected", "<Enter> OK", and "<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.




The screenshot shows the VMware ESXi login interface. The VMware logo is in the top left. On the left side, there are two input fields: "User name" with the value "root" and "Password" with a masked password. Below these fields is a "Log in" button. On the right side, the text "vmware ESXi" is displayed. At the bottom left, there is a small icon and the text "Open the VMware Host Client documentation".



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



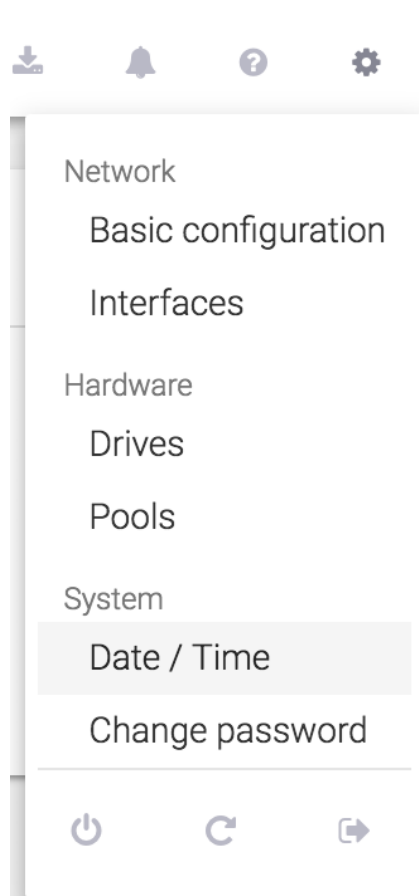
Select Licensing tab → Assign License.



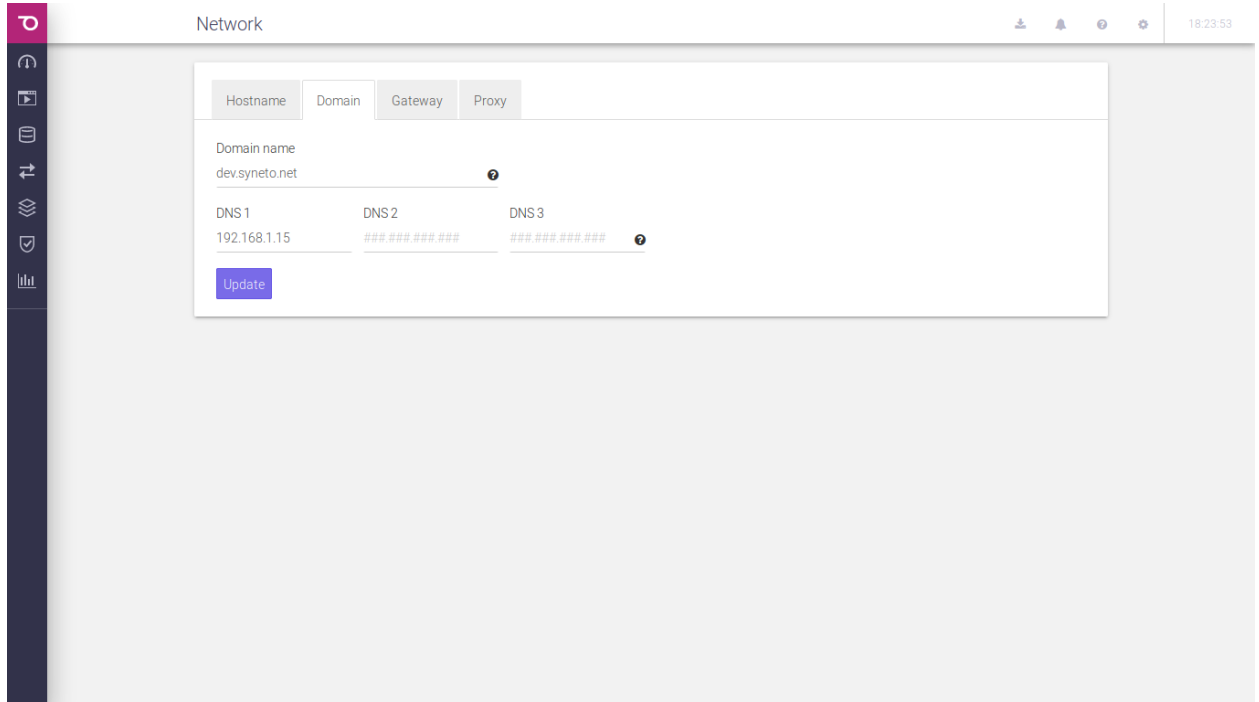
Type or paste the ESXi license key. Click Check License.

Configure Management VM networking

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



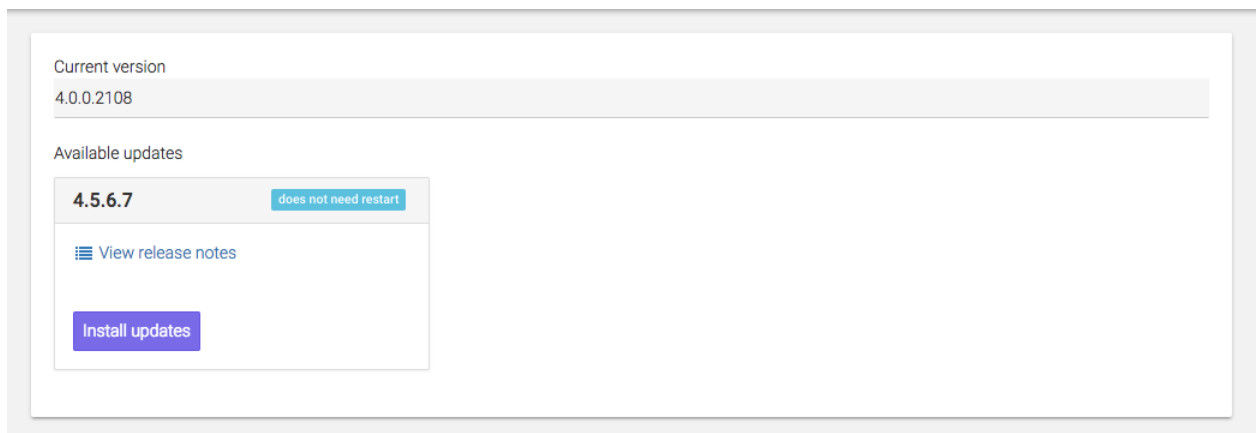
The Network page will open. Here you can set the hostname, domain name, DNS, gateway and proxy server.



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.

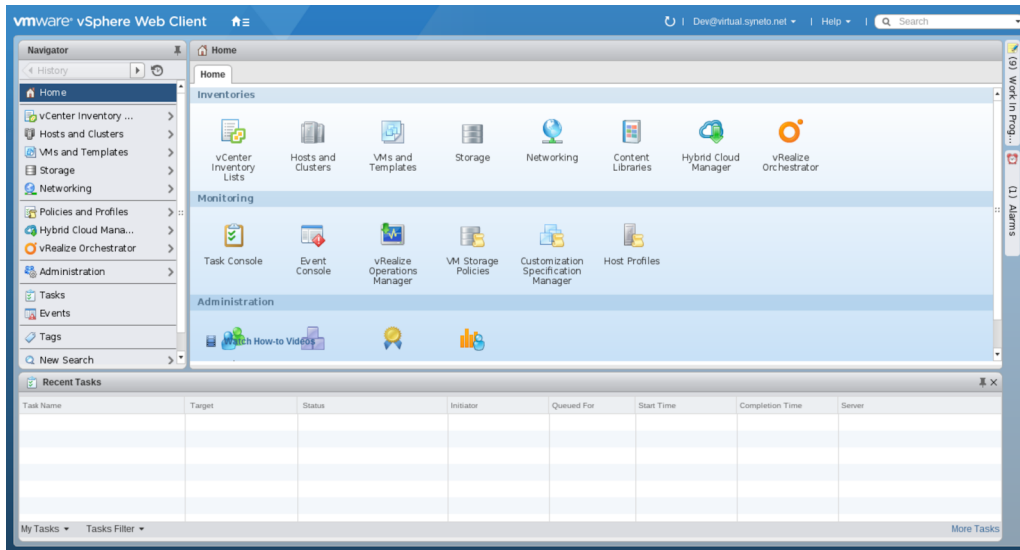


For more information, check [the software update section](#).

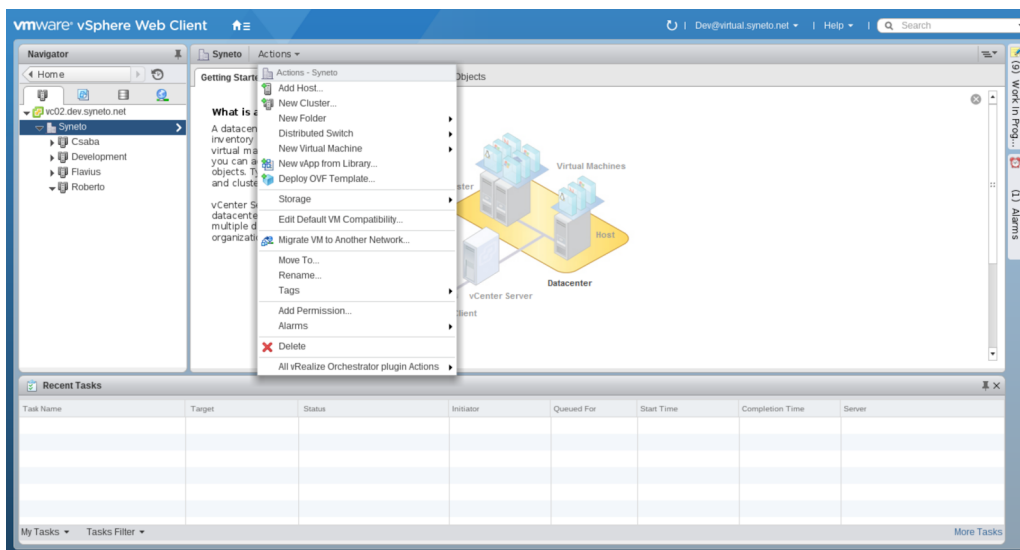
Join vCenter

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

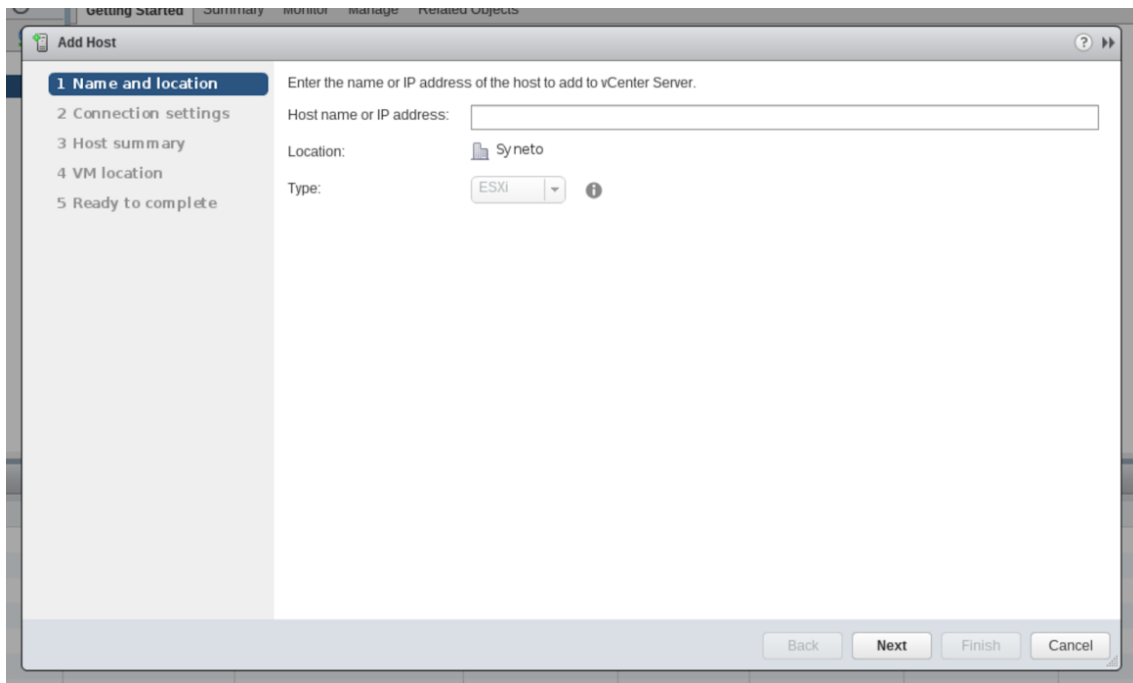
If you have a vCenter you can join our solution to it for easier management.
Log in to vCenter Management Interface. In the left menu, select Hosts and Clusters.



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 left side. The progress indicator shows five steps: 1 Name and location (selected), 2 Connection settings, 3 Host summary, 4 VM location, and 5 Ready to complete. The main area of the dialog contains the following text and fields:

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

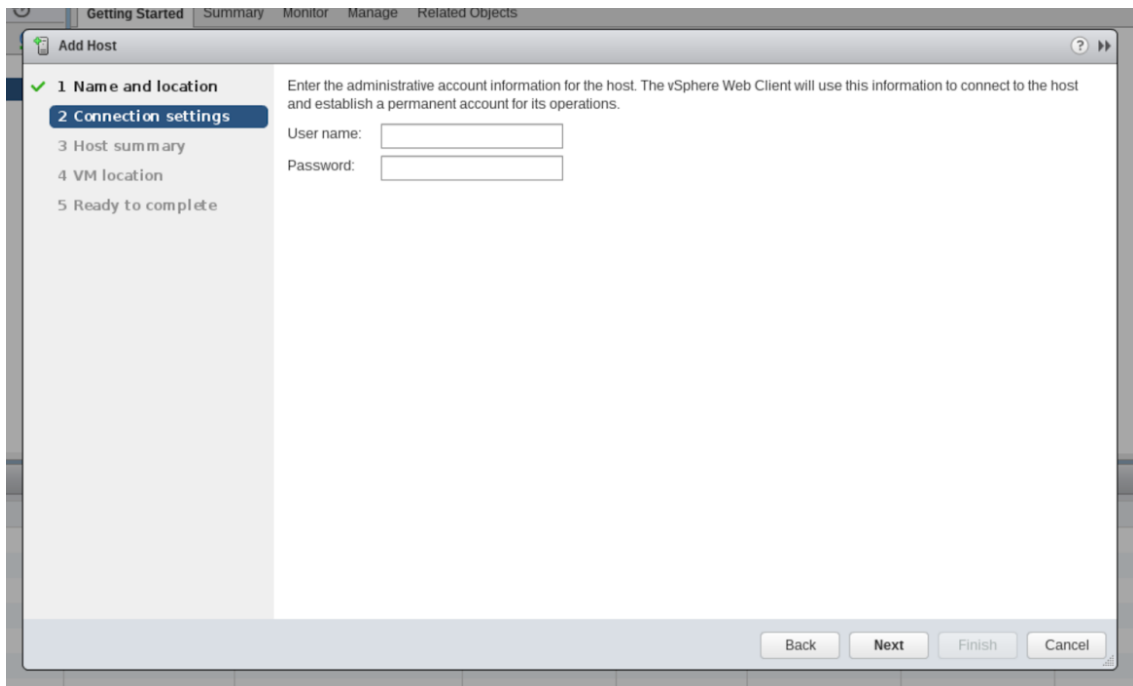
Host name or IP address:

Location:

Type: ⓘ

At the bottom of the dialog, 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 progress indicator on the left shows Step 1 completed with a green checkmark and Step 2 selected. The main area of the dialog 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 of the dialog, 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.

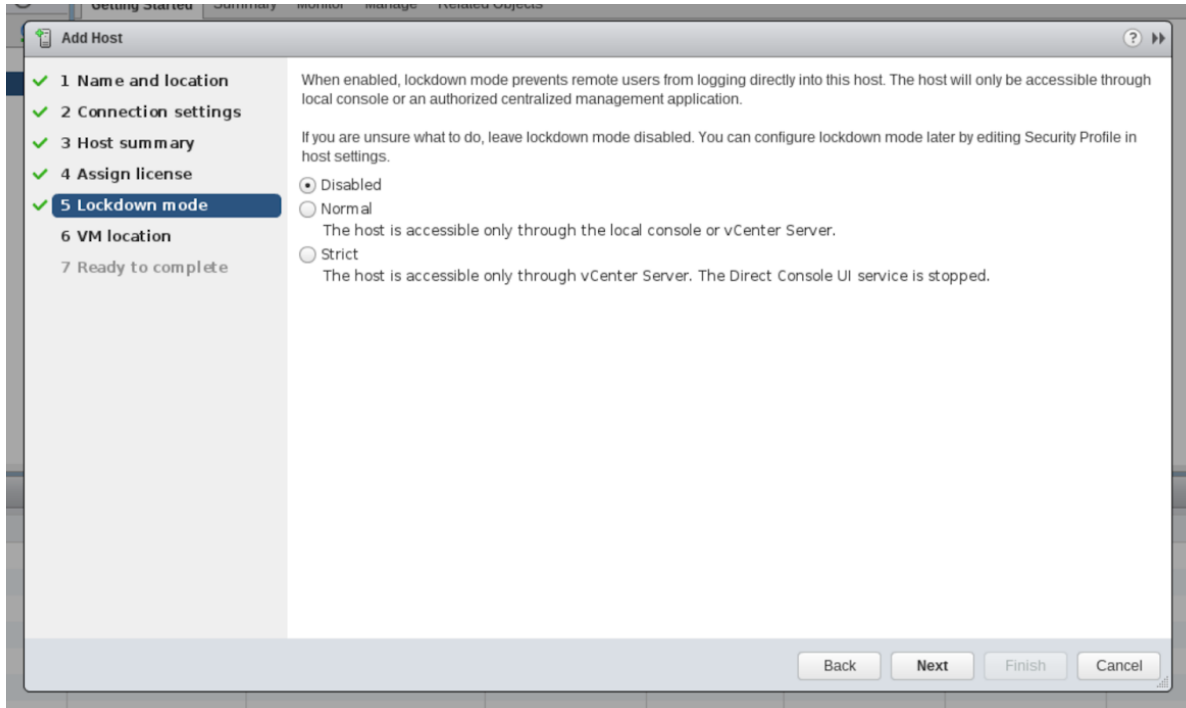


This window lets you configure the lockdown mode.

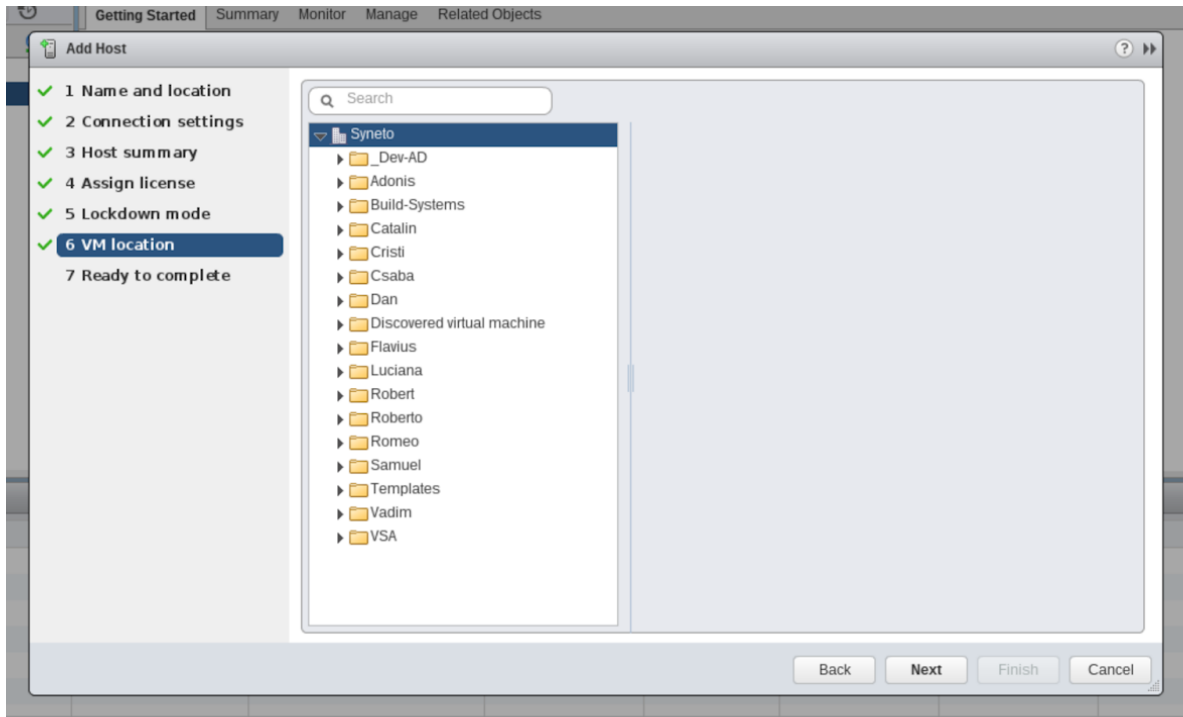
When enabled, lockdown mode prevents remote users from logging directly into this host. The host will only be accessible through local console or an authorized centralized management application.

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.



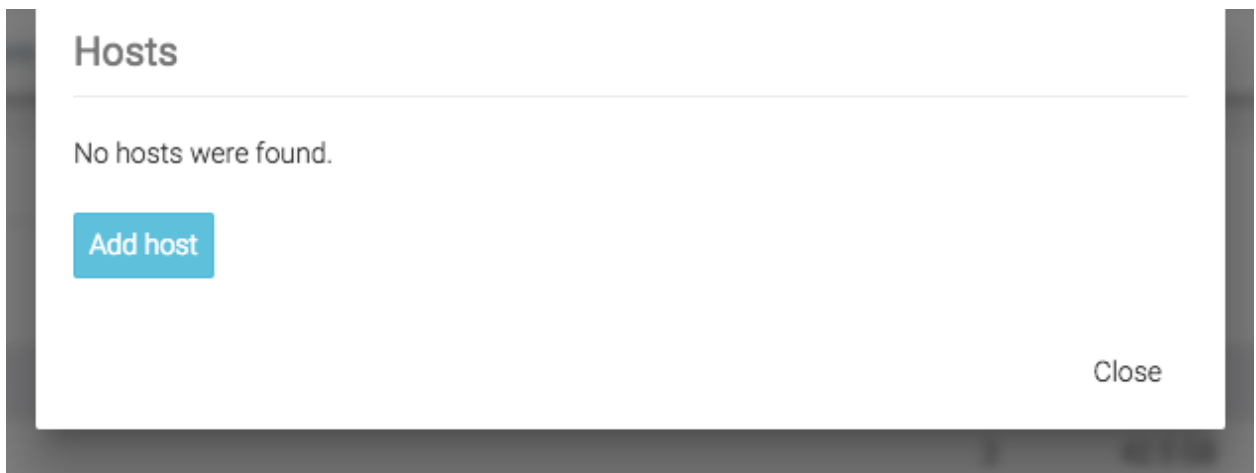
Review your selections and if you do not want to modify anything, click Finish.



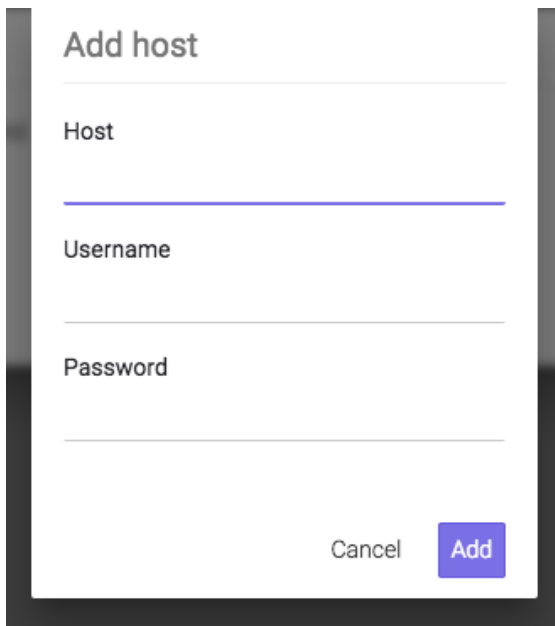
Now, we need to add vCenter to SynetoOS. On the SynetoOS Management Interface, go to the Datastores page.



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.

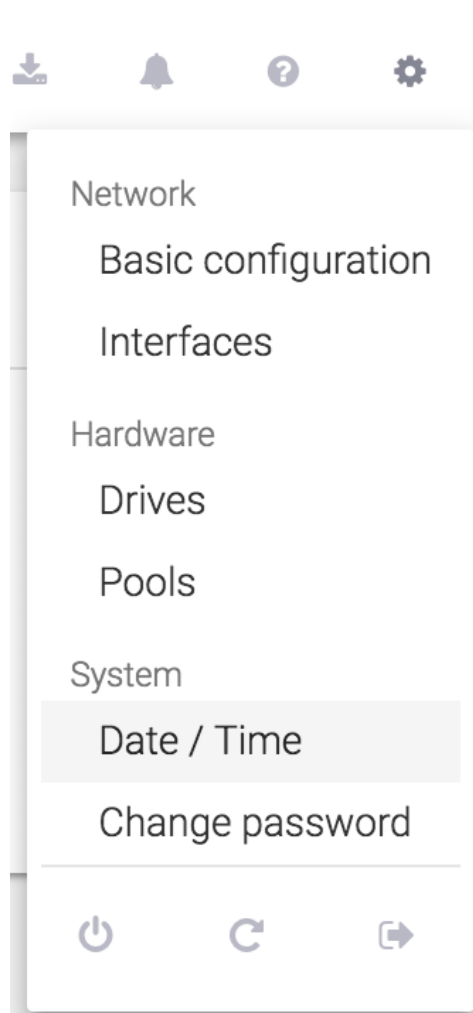
The image shows a form titled 'Add host' with a white background and a dark border. The title is at the top left. Below the title, there are three input fields, each with a label to its left: 'Host', 'Username', and 'Password'. Each label is followed by a horizontal line representing the input field. At the bottom right of the form, there are two buttons: a 'Cancel' button and an 'Add' button.

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

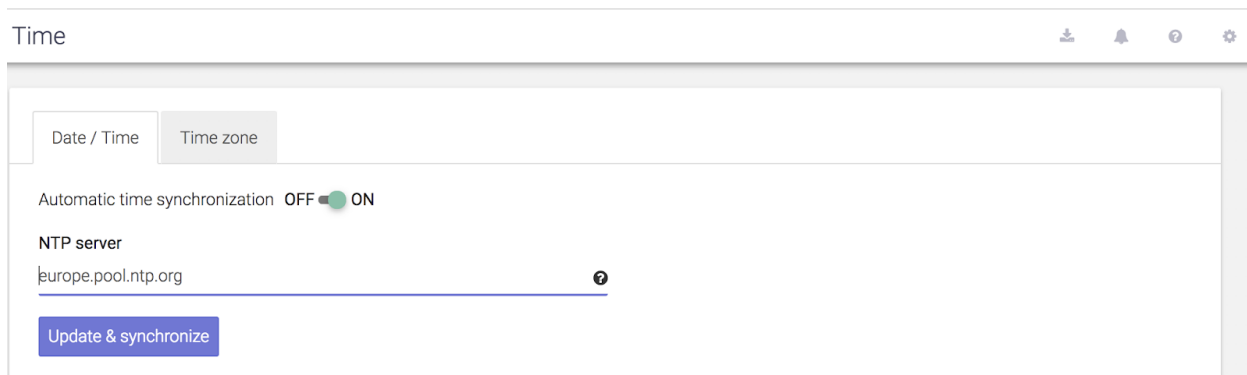
For more information about managing non-Syneto hosts, see chapter [Serving non-Syneto hosts](#)

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 corruptions.

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

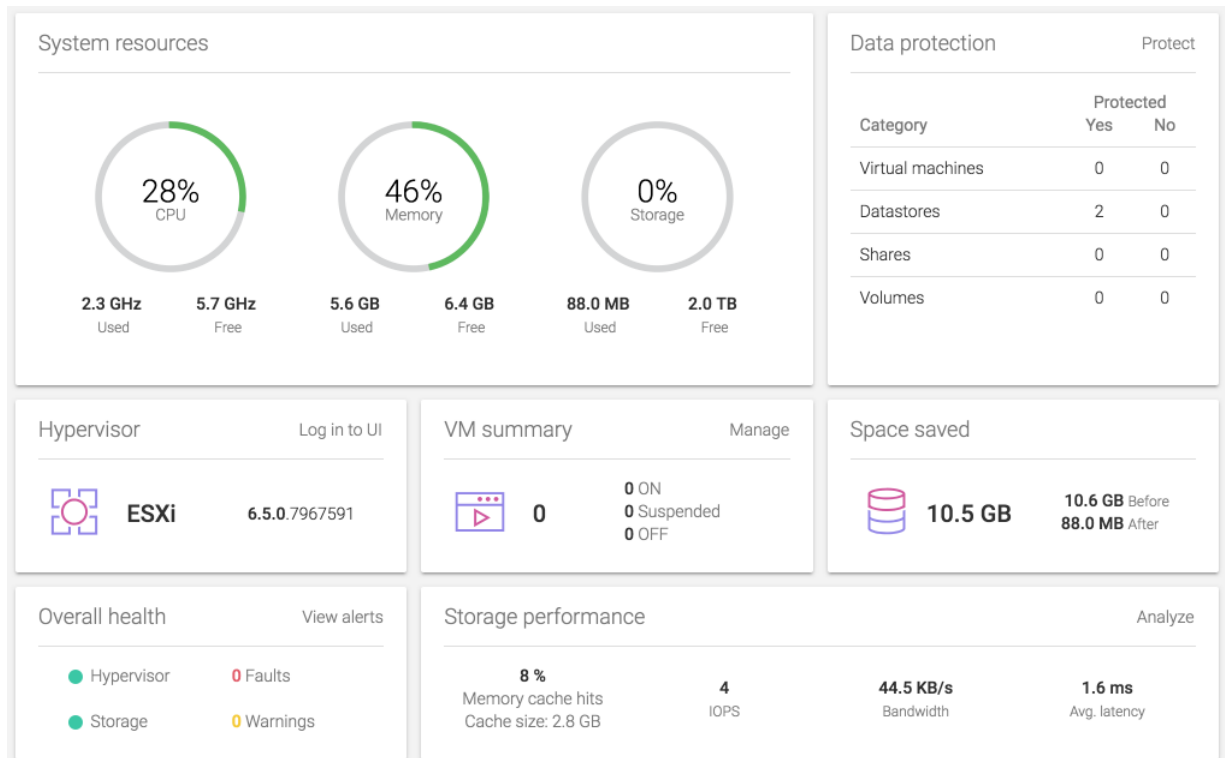
Self-healing protects especially hot data, which is accessed frequently and thus checked frequently for its consistency/integrity. Less frequently accessed (cold) data is at a larger risk because its blocks may accumulate multiple errors, which could render checksums useless. To protect cold data, the 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.

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.

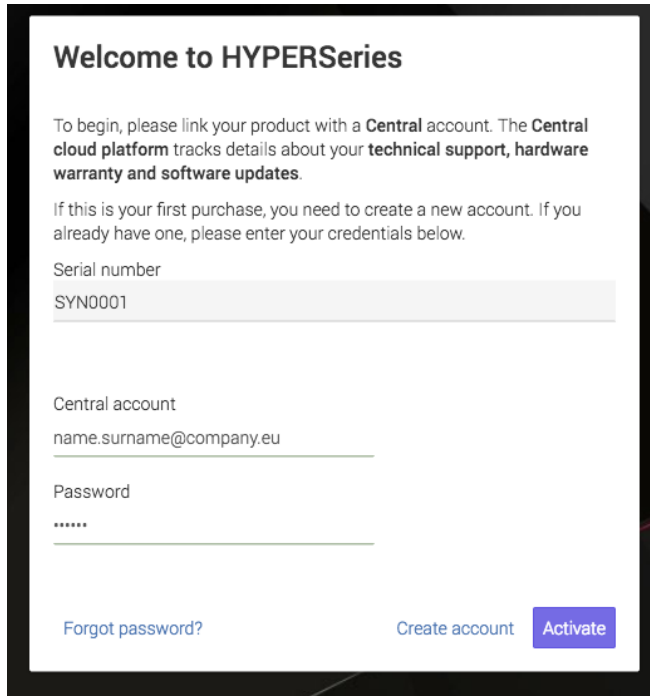
Confirm

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.

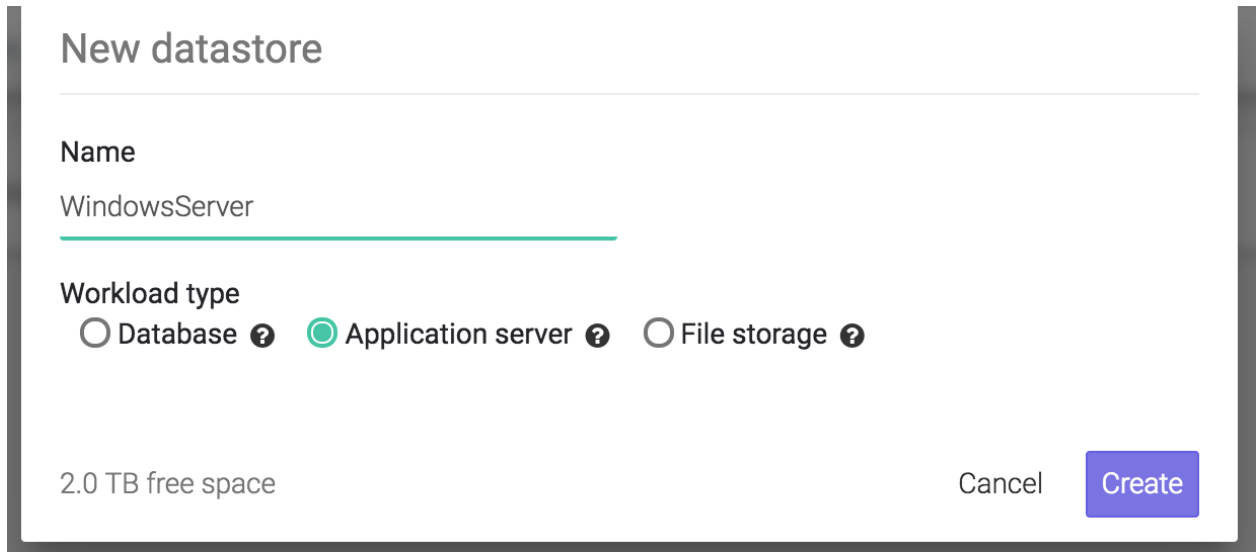
Provisioning and managing

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 "WindowsServer". 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 "2.0 TB 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. At the top, there is a search bar and a 'Show' button with a dropdown menu containing 'Originals', 'Clones', and 'Replicas'. The page number '1-1 of 1' is in the top right corner. Below this is a table with columns for 'Name', 'VMs', 'Used space', and 'Protection'. The table contains one entry: 'WindowsServer' with 0 VMs, 23.0 KB used space, and a 'NOT PROTECTED' status.



Below the table is a detailed view for the 'WindowsServer' datastore. It includes a 'Delete' button in the top right. The details are organized into sections:

- Virtual machines:** none
- Workload:** Application server (with a 'Change' link)
- Protection:** NOT PROTECTED (with a 'Protect' link)
- Snapshots:** none
- Space effective used:** A bar chart showing 11.50 KB used and 23.00 KB total. A legend indicates 'Data: 23.00 KB' and 'Snapshots: 0.00 B'.
- Compression:** 1.00x, Saved 0.0 B
- Mounted on:** syneto-esxi-7919f26c.dev.syneto.net (with a 'Change' link)
- Description:** [2018-07-16 13:06:35] Datastore WindowsServer created. (with an 'Edit' link)

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 

Name	Mounts	Size	Mounted on	Pool	Protection
EmailServer-Linux	3		1 host	hybrid	NOT PROTECTED
OracleServer	0	23.0 KB	1 host	hybrid	NOT PROTECTED
WindowsServer	1	100.2 MB	1 host	hybrid	H
WindowsServer_2018_07_17_10_37	1	59.5 KB	1 host	hybrid	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.

New datastore

Name

Workload type

Database ?
 Application server ?
 File storage ?

Store on

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

Cancel Create

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.

Search 1-4 of 4

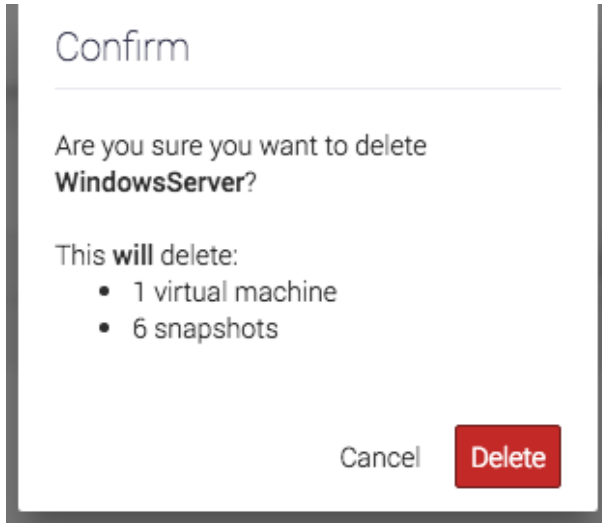
Show Originals, Clones, » Replicas

Name	VMs	Used space	Mounted on	Pool	Protection
EmailServer-Linux	0	23.0 KB	1 host	hybrid	NOT PROTECTED
OracleServer	0	23.0 KB	1 host	hybrid	NOT PROTECTED
WindowsServer	1	100.2 MB	1 host	hybrid	H
WindowsServer_2018_07_17_10_37	1	59.5 KB	1 host	hybrid	NOT PROTECTED

OracleServer Delete

Pool	hybrid	Space effective used	<div style="width: 100%; height: 10px; background: linear-gradient(to right, blue, lightblue);"></div>	11.50 KB 23.00 KB
Virtual machines	none		<div style="width: 100%; height: 10px; background: linear-gradient(to right, blue, lightblue);"></div>	
Workload	Application server Change			
Protection	NOT PROTECTED Protect		Compression 1.00x	Saved 0.0 B
Snapshots	none	Mounted on	• syneto-esxi-7919f26c.dev.syneto.net Change	
		Description	[2018-07-17 10:51:34] Datastore OracleServer 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	Pool	Protection
EmailServer-Linux	0	23.0 KB	hybrid	NOT PROTECTED
OracleServer	0	23.0 KB	hybrid	NOT PROTECTED
WindowsServer	1	1.1 GB	hybrid	H

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	Mounted on	Pool	Protection
EmailServer-Linux	0	23.0 KB	1 host	hybrid	NOT PROTECTED
OracleServer	0	23.0 KB	1 host	hybrid	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	Mounted on	Pool	Protection
WindowsServer	1	100.2 MB	1 host	hybrid	H
WindowsServer_2018_07_17_10_37	1	59.5 KB	1 host	hybrid	NOT PROTECTED

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.

Search Show Originals, Clones, » Replicas 1-3 of 3

Name	VMs	Used space	Pool	Protection
EmailServer-Linux	0	23.0 KB	hybrid	NOT PROTECTED
OracleServer	0	23.0 KB	hybrid	NOT PROTECTED
WindowsServers	2	121.5 KB	hybrid	M

WindowsServers Delete

Pool	hybrid	Space effective used	<div style="width: 100%; height: 10px; background: linear-gradient(to right, #007bff, #6c757d);"></div>	64.50 KB 121.50 KB
Virtual machines	2 View		<div style="width: 100%; height: 10px; background: linear-gradient(to right, #007bff, #6c757d);"></div>	■ Data: 42.00 KB — ■ Snapshots: 79.50 KB
Workload	Application server Change		Compression 1.01x	Saved 0.0 B
Protection	M Protect	Mounted on	• syneto-esxi-7919f26c.dev.syneto.net Change	
Snapshots	3 Browse	Description	[2018-07-17 16:15:34] Datastore WindowsServers 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 Browse in the snapshots section of the detailed view of a datastore.

Snapshots

< November 2017 >

Mo	Tu	We	Th	Fr	Sa	Su
30	31	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	1	2	3
4	5	6	7	8	9	10

November 8, 2017 6 snapshots

14:40	139.5 KB
14:39	2.4 MB
14:38	1.9 MB
14:37	2.1 MB
14:36	1.7 MB
14:35	6.2 MB

Summary

Datastore	WindowsServer
Oldest snapshot on	2017/11/08
Total snapshots	6

Close

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

Snapshots

< November 2017 >

Mo	Tu	We	Th	Fr	Sa	Su
30	31	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	1	2	3
4	5	6	7	8	9	10

November 8, 2017 6 snapshots

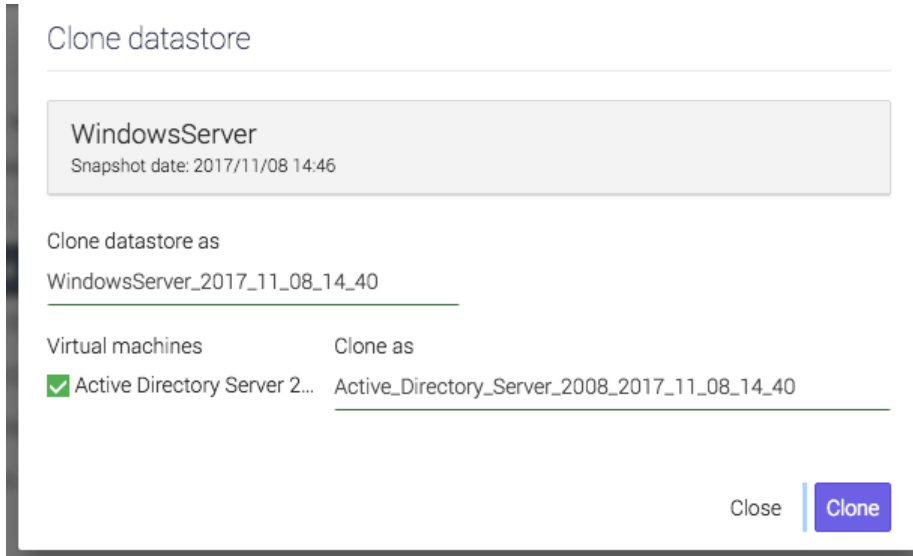
14:40	139.5 KB
14:39	2.4 MB
14:38	1.9 MB
14:37	2.1 MB
14:36	1.7 MB
14:35	6.2 MB

Summary

Datastore	WindowsServer
Oldest snapshot on	2017/11/08
Total snapshots	6

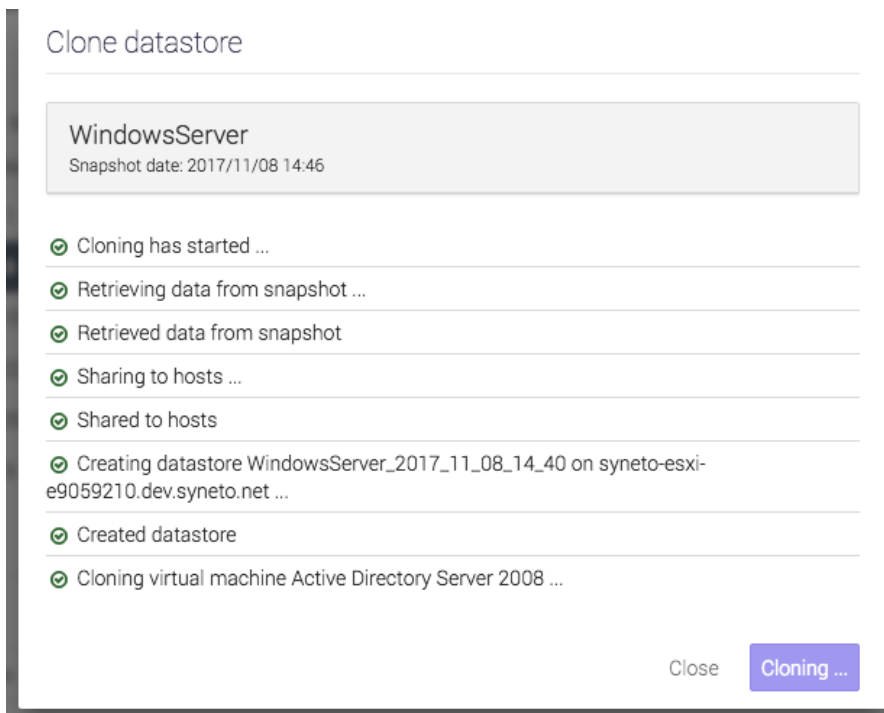
Close

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. At the top, there are controls for 'Create / Register VM', 'Console', 'Power on', 'Power off', 'Suspend', 'Refresh', and 'Actions'. A search bar is also present. Below this is a table of virtual machines:

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

Below the table, there is a 'Quick filters...' dropdown and a '6 items' indicator. The details view for the selected VM 'Active_Directory_Server_2008_2017_11_08_14_40' is shown below:

Active_Directory_Server_2008_2017_11_08_14_40


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

On the right side, there are resource usage indicators: CPU (0 MHz), MEMORY (0 B), and 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.

The screenshot shows the vSphere Datastore browser. At the top, there is a search bar and a dropdown menu set to 'Show Originals, Clones, » Replicas'. The page number '1-4 of 4' is shown in the top right. Below this is a table of datastores:

Name	VMs	Used space	Pool	Protection
EmailServer-Linux	0	23.0 KB	hybrid	NOT PROTECTED
OracleServer	0	23.0 KB	hybrid	NOT PROTECTED
WindowsServers	2	95.0 KB	hybrid	M
 WindowsServers_2018_07_17_16_46	0	1.0 KB	hybrid	NOT PROTECTED

Below the table, the details for the selected datastore 'WindowsServers_2018_07_17_16_46' are shown. At the top right, there are 'Revert' and 'Delete' buttons. A blue information box is highlighted, containing the text: 'Cloned from WindowsServers (2018/07/17 16:46)'. Below this, there are several rows of information:

- Pool: hybrid
- Virtual machines: 0/2 (with a 'Register additional' link)
- Workload: Application server (with a 'Change' link)
- Space effective used: 512.00 B (with a progress bar showing 1.00 KB used)
- Compression: 1.00x (with 'Saved 0.0 B')
- Mounted on: syneto-esxi-7919f26c.dev.syneto.net (with a 'Change' link)
- Description: [2018-07-17 16:47:52] Datastore WindowsServers_2018_07_17_16_46 created. (with an 'Edit' link)

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	<div style="width: 100%;"><div style="width: 55%;"></div></div>	18.0 GB 9.9 GB
Protection	D			<div style="width: 100%;"><div style="width: 55%;"></div></div>	Data: 6.1 GB — Snapshots: 3.8 GB
Snapshots	3	Browse		Compression 1.81x	Saved 8.1 GB
Clones			Mounted on	• syneto-esxi-70d3bdb4.dev.syneto.net Change	
			Description	[2017-07-22 11:56:03] Datastore WindowsServer created. Edit	

Name	VMs	Snapshot date
WindowsServer_2018_02_09_14_40	1/1	2018/02/09 14:40
WindowsServer_2018_02_08_16_40	1/1	2018/02/08 16:40

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.

WindowsServer Delete

Virtual machines	1	View	Space effective used	<div style="width: 100%;"><div style="width: 55%;"></div></div>	47.5 KB 90.5 KB
Protection	H			<div style="width: 100%;"><div style="width: 55%;"></div></div>	Data: 32.0 KB — Snapshots: 58.5 KB
Snapshots	3	Browse		Compression 1.01x	Saved 0.0 B
Clones			Mounted on	• syneto-esxi-70d3bdb4.dev.syneto.net Change	
			Description	[2018-01-23 15:35:03] Datastore WindowsServer created. Edit	

Name	VMs	Snapshot date
WindowsServer_2017_11_08_14_40	1/1	2018/01/23 16:00

Register additional VMs
 Delete clone

Or select the cloned datastore and click Delete.

WindowsServer_2017_11_08_14_40 1 54 KB NOT PROTECTED

WindowsServer_2017_11_08_14_40 [Revert](#) [Delete](#)

Cloned from *WindowsServer* (2018/01/23 16:00)

Virtual machines 1 [View](#) Space effective used 274.5 KB
53.5 KB
■ Data: 53.5 KB — ■ Snapshots: 0.0 B

Compression 5.84x Saved 221.0 KB

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

Description [2018-01-23 16:32:31] Datastore WindowsServer_2017_11_08_14_40 created. [Edit](#)

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.

Confirm

Are you sure you want to delete the cloned datastore with its virtual machines?

[Cancel](#) [Delete clone](#)

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.

Search 6-6 of 6 < >

Name	VMs	Used space	Protection
ds1_2018_01_23_16_09	0	1 KB	NOT PROTECTED

ds1_2018_01_23_16_09 [Revert](#) [Delete](#)

Cloned from ds1 (2018/01/23 16:09)

Virtual machines 0

Space effective used 512.0 B

■ Data: 1.0 KB ■ Snapshots: 0.0 B 1.0 KB

Compression 1.00x Saved 0.0 B

Mounted on [Change](#)

- syneto-esxi-02c7d82c.dev.syneto.net

Description [Edit](#)

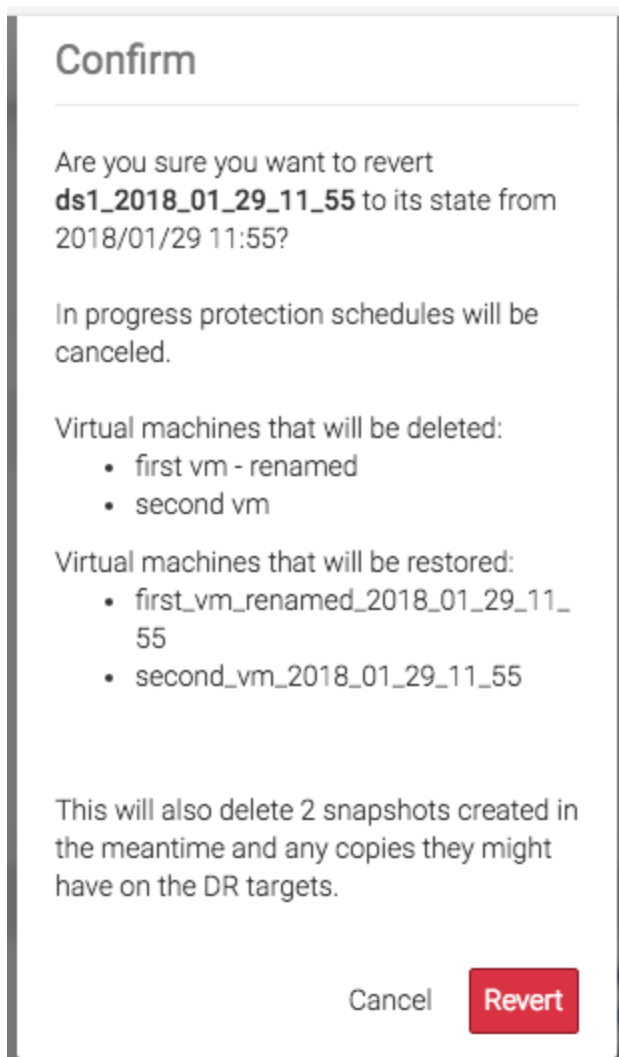
[2018-01-23 16:19:33] Datastore ds1_2018_01_23_16_09 created.

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

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 needed don't 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 shows the SynetoOS interface. At the top, there is a search bar and a filter dropdown set to 'Show Originals, Clones, » Replicas'. Below this is a table of clones:

Name	VMs	Used space	Pool	Protection
EmailServer-Linux	0	23.0 KB	hybrid	NOT PROTECTED
OracleServer	0	23.0 KB	hybrid	NOT PROTECTED
WindowsServers	2	95.0 KB	hybrid	M
WindowsServers_2018_07_17_16_18	0	1.0 KB	hybrid	NOT PROTECTED

The clone 'WindowsServers_2018_07_17_16_18' is selected. Below the table, the details for this clone are shown:

WindowsServers_2018_07_17_16_18 Revert Delete

Cloned from *WindowsServers* (2018/07/17 16:18)

Pool	hybrid	Space effective used	<div style="width: 100%; height: 10px; background-color: #0070c0;"></div> 512.00 B
Virtual machines	0/2 Register additional		<div style="width: 100%; height: 10px; background-color: #0070c0;"></div> 1.00 KB
Workload	Application server Change		<div style="width: 100%; height: 10px; background-color: #0070c0;"></div> Data: 1.00 KB — Snapshots: 0.00 B
		Compression 1.00x	Saved 0.0 B
Mounted on	• syneto-esxi-7919f26c.dev.syneto.net Change		
Description	[2018-07-17 16:18:45] Datastore WindowsServers_2018_07_17_16_18 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 the SynetoOS interface. At the top, there is a search bar and a filter menu showing 'Show Originals, Clones, » Replicas' with '1-4 of 4' items. Below this is a table of datastores:

Name	VMs	Used space	Pool	Protection
EmailServer-Linux	0	23.0 KB	hybrid	NOT PROTECTED
OracleServer	0	23.0 KB	hybrid	NOT PROTECTED
WindowsServers	2	121.5 KB	hybrid	M
WindowsServers_2018_07_17_16_18	1	60.0 KB	hybrid	NOT PROTECTED

The selected datastore, **WindowsServers_2018_07_17_16_18**, is shown in detail below. It is a clone of *WindowsServers (2018/07/17 16:18)*. The details include:

- Pool:** hybrid (with a dropdown menu showing *Microsoft_SQL_Server_2018_07_17_16_18*)
- Virtual machines:** 1/2 (with [View](#) and [Register additional](#) links)
- Workload:** Application server (with [Change](#) link)
- Space effective used:** 277.50 KB total, 60.00 KB used. Legend: Data: 60.00 KB, Snapshots: 0.00 B.
- Compression:** 5.55x, Saved 217.5 KB
- Mounted on:** syneto-esxi-7919f26c.dev.syneto.net (with [Change](#) link)
- Description:** [2018-07-17 16:18:45] Datastore WindowsServers_2018_07_17_16_18 created. (with [Edit](#) link)

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).

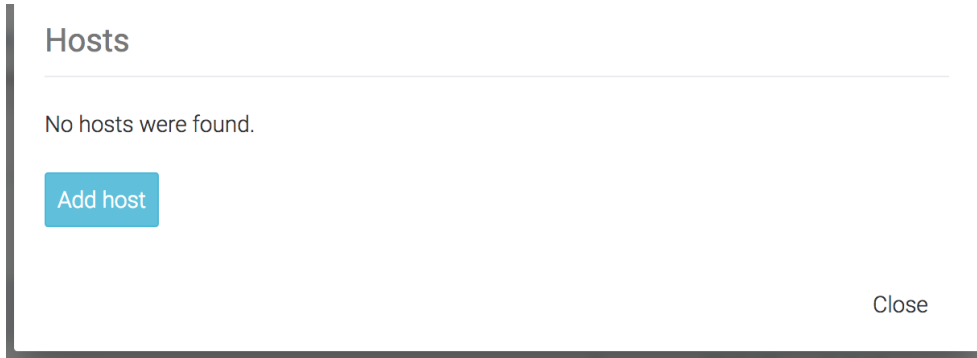
If you click [Change](#) you will be able to unmount it from this local hypervisor.

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

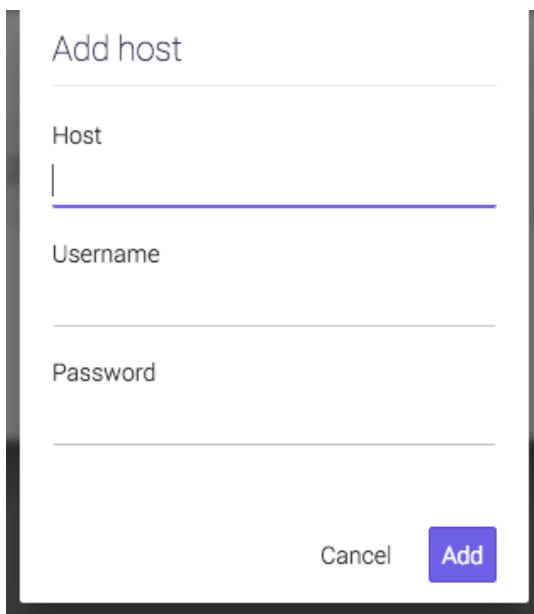
Go to Main menu → Datastores. On the action bar you will see a button named External hosts.



By default the dialog will show no host.

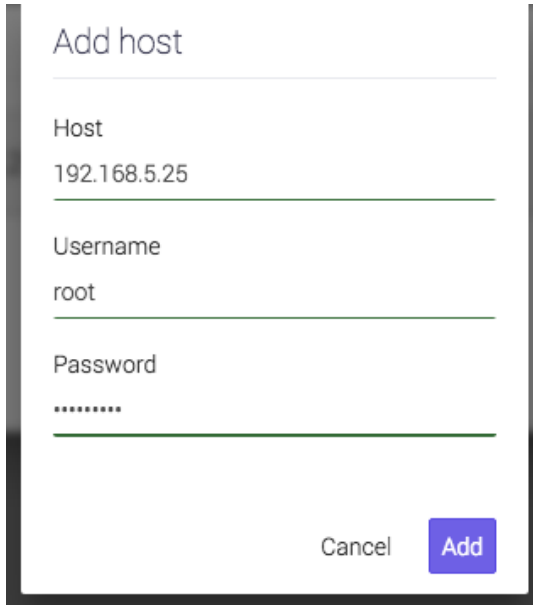
NOTE: Local ESXi hypervisor is not displayed on the External hosts list.

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

A screenshot of a dialog box titled "Add host". The dialog has a white background and a dark border. At the top, the title "Add host" is displayed. Below the title, there are three input fields: "Host", "Username", and "Password". The "Host" field has a blue underline. At the bottom of the dialog, there are two buttons: "Cancel" and "Add".

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.

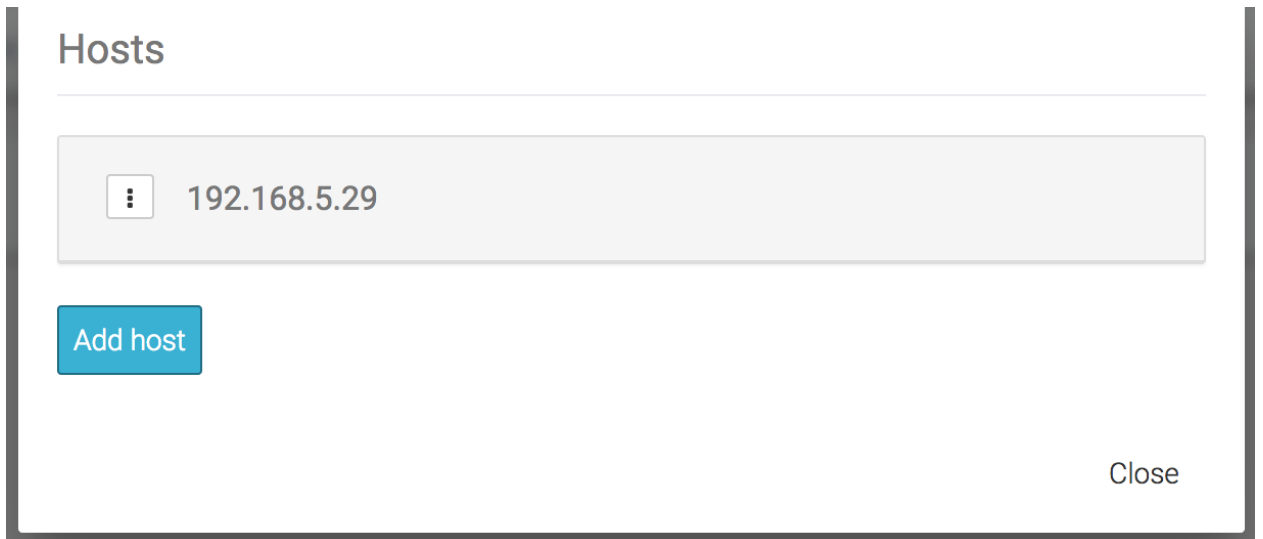
WARNING: If your ESXi is managed by a vCenter, you must specify the vCenter at the Host field. Adding the ESXi host may lead to partial or unexpected interaction between the Syneto HYPER and VMware.




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

- Add host** (Title)
- Host**: 192.168.5.25
- Username**: root
- Password**:
- 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
- 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 1-4 of 4

Show Originals, Clones, » Replicas

Name	VMs	Used space	Pool	Protection
EmailServer-Linux	0	23.0 KB	hybrid	NOT PROTECTED
OracleServer	0	23.0 KB	hybrid	NOT PROTECTED
WindowsServers	2	95.0 KB	hybrid	M
<input type="checkbox"/> WindowsServers_2018_07_17_16_46	0	1.0 KB	hybrid	NOT PROTECTED

WindowsServers Delete

Pool	hybrid	Space effective used	<div style="width: 100%;"><div style="width: 50%;"></div></div>	51.00 KB 95.00 KB
Virtual machines	2 View		<div style="width: 100%;"><div style="width: 42%; background-color: #007bff;"></div><div style="width: 13%; background-color: #6c757d;"></div></div>	Data: 42.00 KB — Snapshots: 53.00 KB
Workload	Application server Change			Compression 1.02x Saved 0.0 B
Protection	M Protect			
Snapshots	2 Browse			

Mounted on [Change](#)

Description [2018-07-17 16:15:34] Datastore WindowsServers created. [Edit](#)

Clones

Name	VMs	Snapshot date
<input type="checkbox"/> WindowsServers_2018_07_17_16_46	0/2	2018/07/17 16:46

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

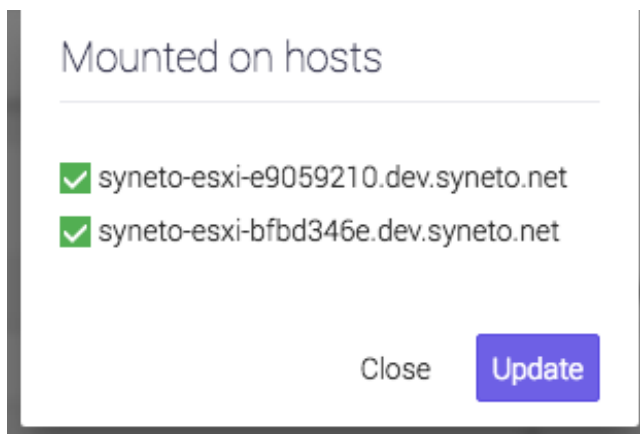
[Close](#) [Update](#)

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 preconfigured with a VMware ESXi local hypervisor. You can provision, manage, run, protect, analyze virtual machines on the device itself.

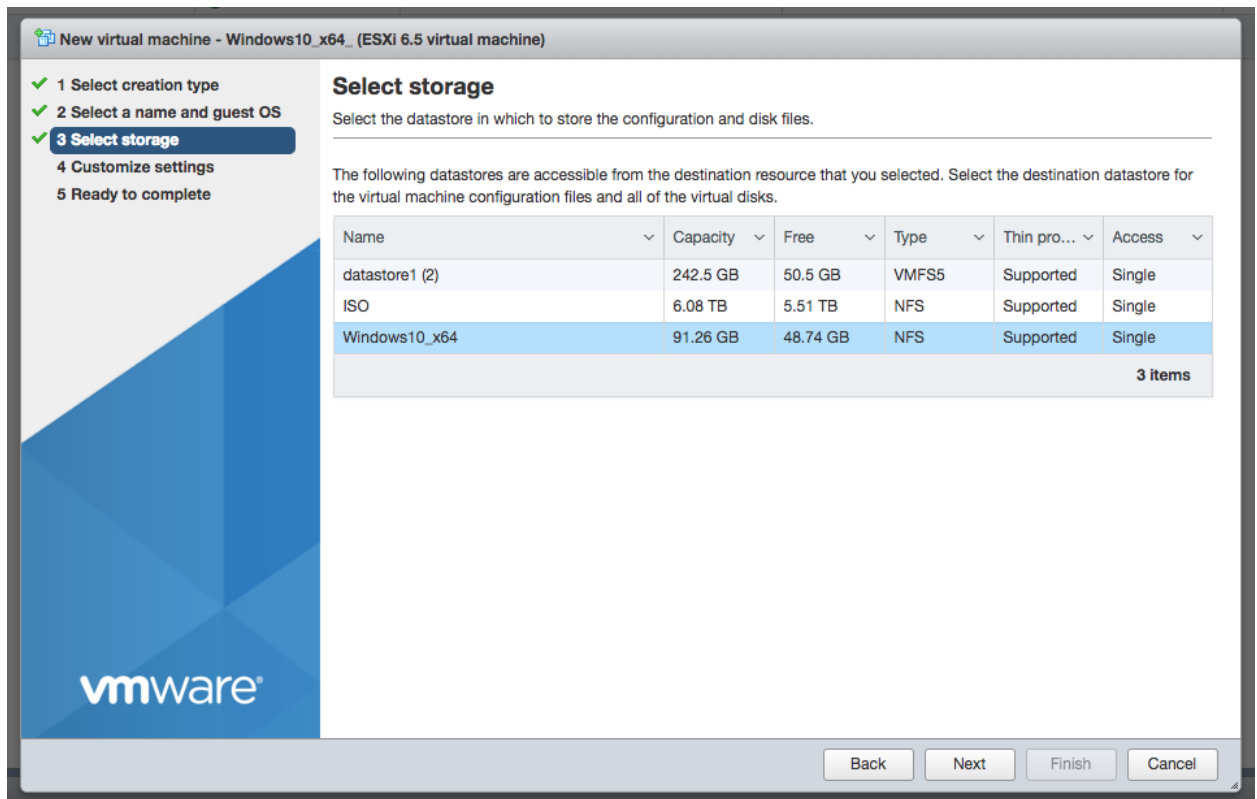
Provisioning and managing

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.

The screenshot displays the SynetoOS interface for managing virtual machines. At the top, the title is 'Virtual machines' with a 'Rescan' button. A search bar is located above a table of VMs. The table has columns for Name, CPU, Memory, Used space, and Protection. One VM, 'Windows10_x64', is selected and highlighted. Below the table, the details pane for 'Windows10_x64' is shown, indicating it is 'Powered off'. The details pane includes a 'Power on' button and various configuration options:

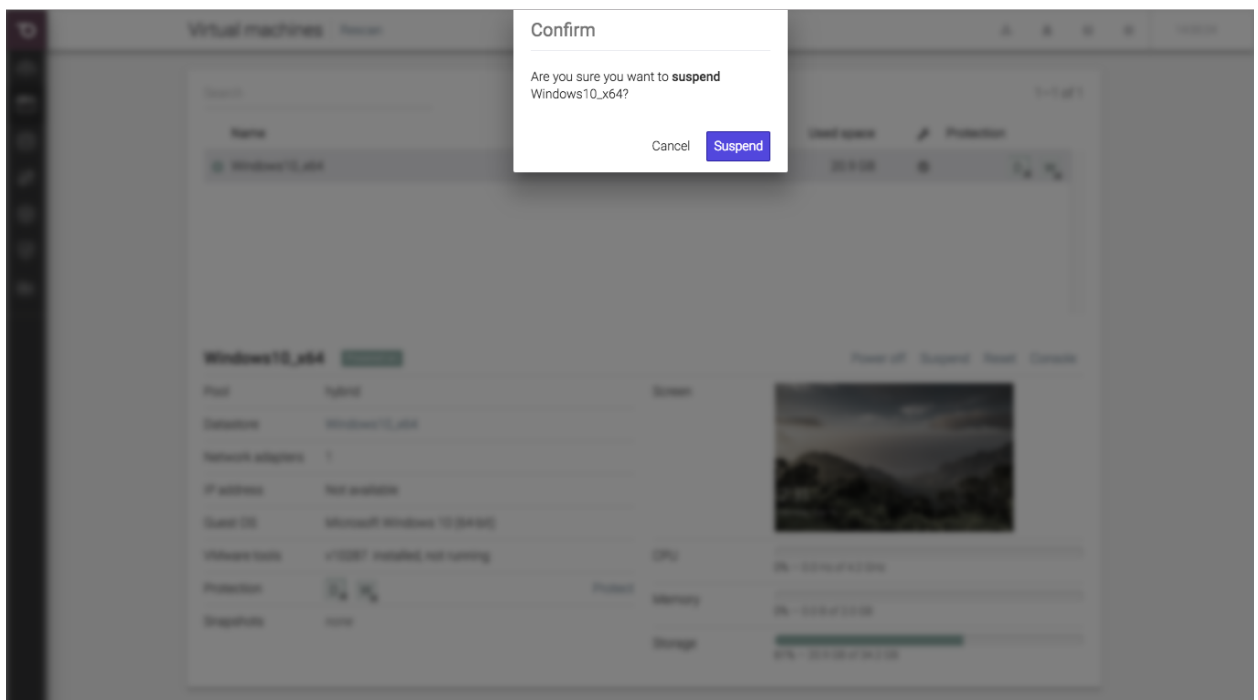
Property	Value	Resource	Usage
Pool	hybrid	CPU	0% — 0.0 Hz of 4.2 GHz
Datastore	Windows10_x64	Memory	0% — 0.0 B of 2.0 GB
Network adapters	1	Storage	61% — 20.8 GB of 34.2 GB
IP address	Not available		
Guest OS	Microsoft Windows 10 (64-bit)		
VMware tools	v10287 installed, not running		
Protection			Protect
Snapshots	none		

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.

Windows10_x64 Powered on Power off Suspend Reset Console

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.

Windows10_x64 Suspended Resume Power off

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 'Virtual machines' management interface. A table lists the VMs, with 'Windows10_x64' selected. Below the table, the details for this VM are displayed, including its configuration and resource usage.

Name	CPU	Memory	Used space	Protection
Windows10_x64	2.8 GHz	1.5 GB	20.9 GB	[D] [W]

Property	Value	Resource Usage
Pool	hybrid	
Datastore	Windows10_x64	
Network adapters	1	
IP address	192.168.5.142	
Guest OS	Microsoft Windows 10 (64-bit)	
VMware tools	v10287 running, managed by VMware	
Protection	[D] [W] Protect	
Snapshots	none	
CPU	67% — 2.8 GHz of 4.2 GHz	<div style="width: 67%;"></div>
Memory	75% — 1.5 GB of 2.0 GB	<div style="width: 75%;"></div>
Storage	65% — 20.9 GB of 32.0 GB	<div style="width: 65%;"></div>

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 of the Windows10_x64 VM. The interface displays the Windows 10 desktop environment, including the Recycle Bin, a 'This PC' window showing the file explorer, and the taskbar with the Start button and system tray. The system tray shows the time as 2:43 PM on 7/16/2018. The top of the console window shows 'syneto | HYPERSeries' and 'Connected (encrypted) to Windows10_x64'.

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:

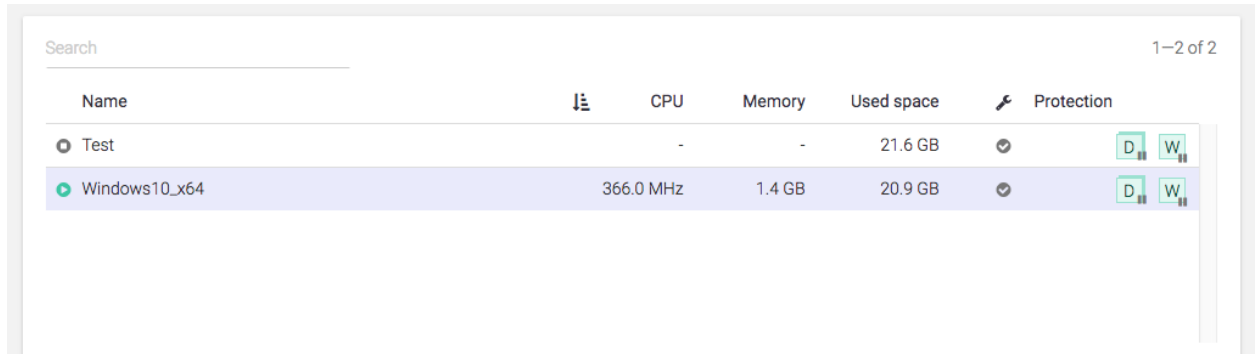
1. Crash consistent - When restoring, the virtual machine will be in a state similar after a power outage.
Recommended for virtual machines resilient to forced reboots. (ie. Linux or Solaris servers or workstations)
2. Application consistent - Before taking the snapshot, a message will be sent to the 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.



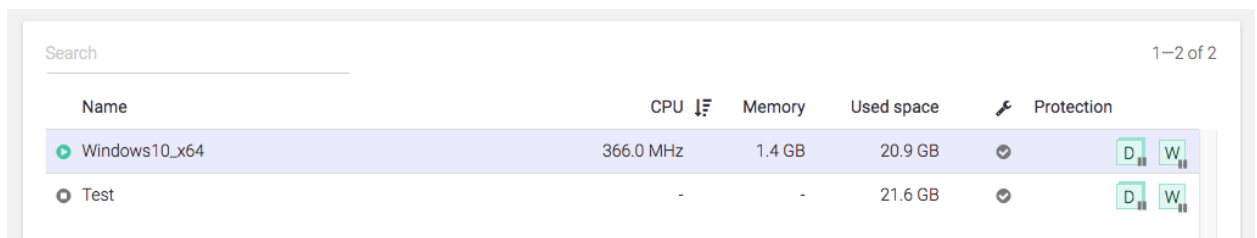
Name	CPU	Memory	Used space	Protection
Test	-	-	21.6 GB	✓
Windows10_x64	366.0 MHz	1.4 GB	20.9 GB	✓

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	☑	D W
Test	-	-	21.6 GB	☑	D W

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	☑	D W
Windows10_x64	366.0 MHz	1.4 GB	20.9 GB	☑	D W

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.

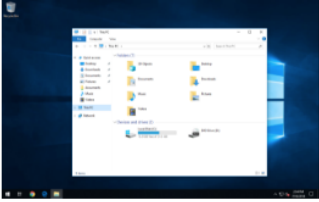
Additional information

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.

Windows10_x64 Powered on

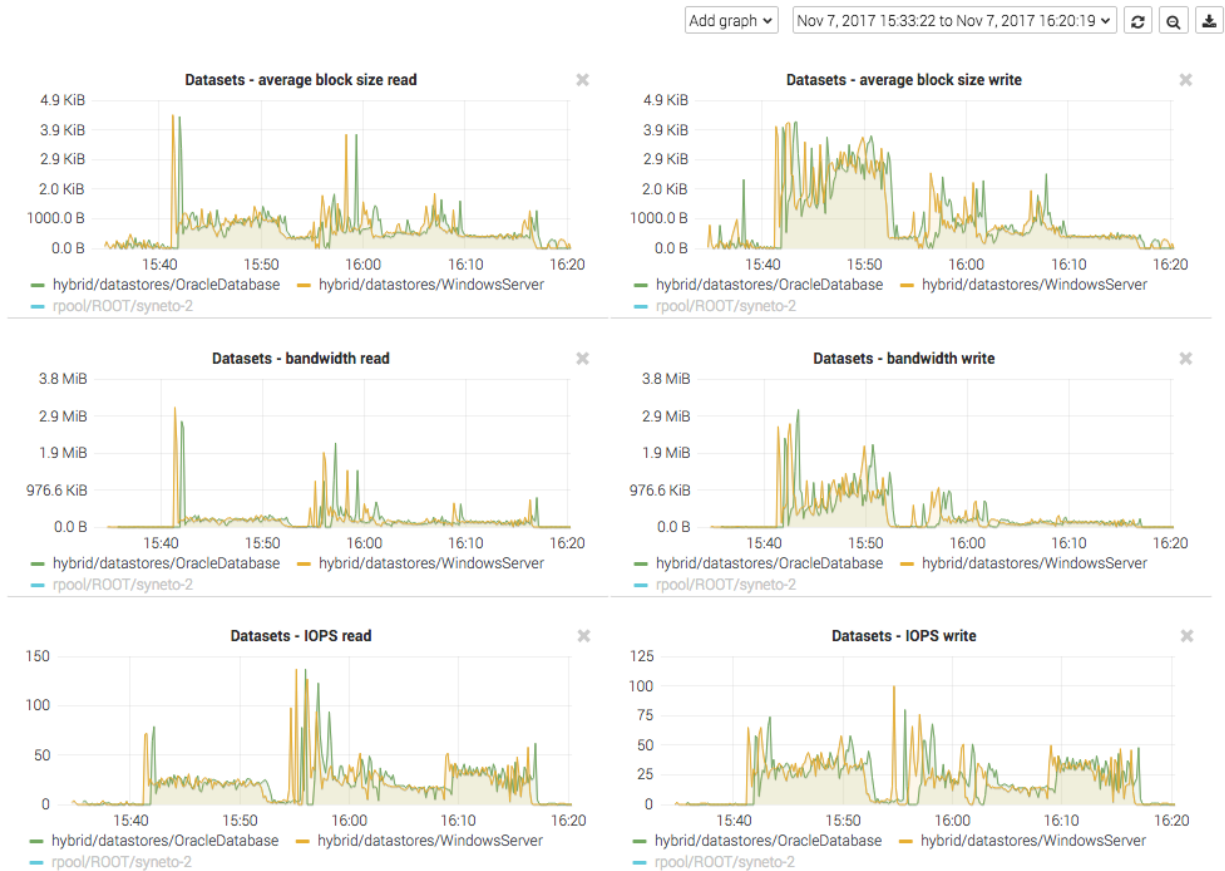
[Power off](#) [Suspend](#) [Reset](#) [Console](#)

Pool	hybrid	Screen	
Datastore	Windows10_x64		
Network adapters	1		
IP address	192.168.5.142		
Guest OS	Microsoft Windows 10 (64-bit)		
VMware tools	v10287 running, managed by VMware	CPU	<div style="width: 9%; background-color: #28a745; height: 10px; margin-bottom: 2px;"></div> 9% – 366.0 MHz of 4.2 GHz
Protection	D W	Memory	<div style="width: 68%; background-color: #28a745; height: 10px; margin-bottom: 2px;"></div> 68% – 1.4 GB of 2.0 GB
Snapshots	none	Storage	<div style="width: 65%; background-color: #28a745; height: 10px; margin-bottom: 2px;"></div> 65% – 20.9 GB of 32.0 GB

Analyzing performance

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

Dataset IOPS



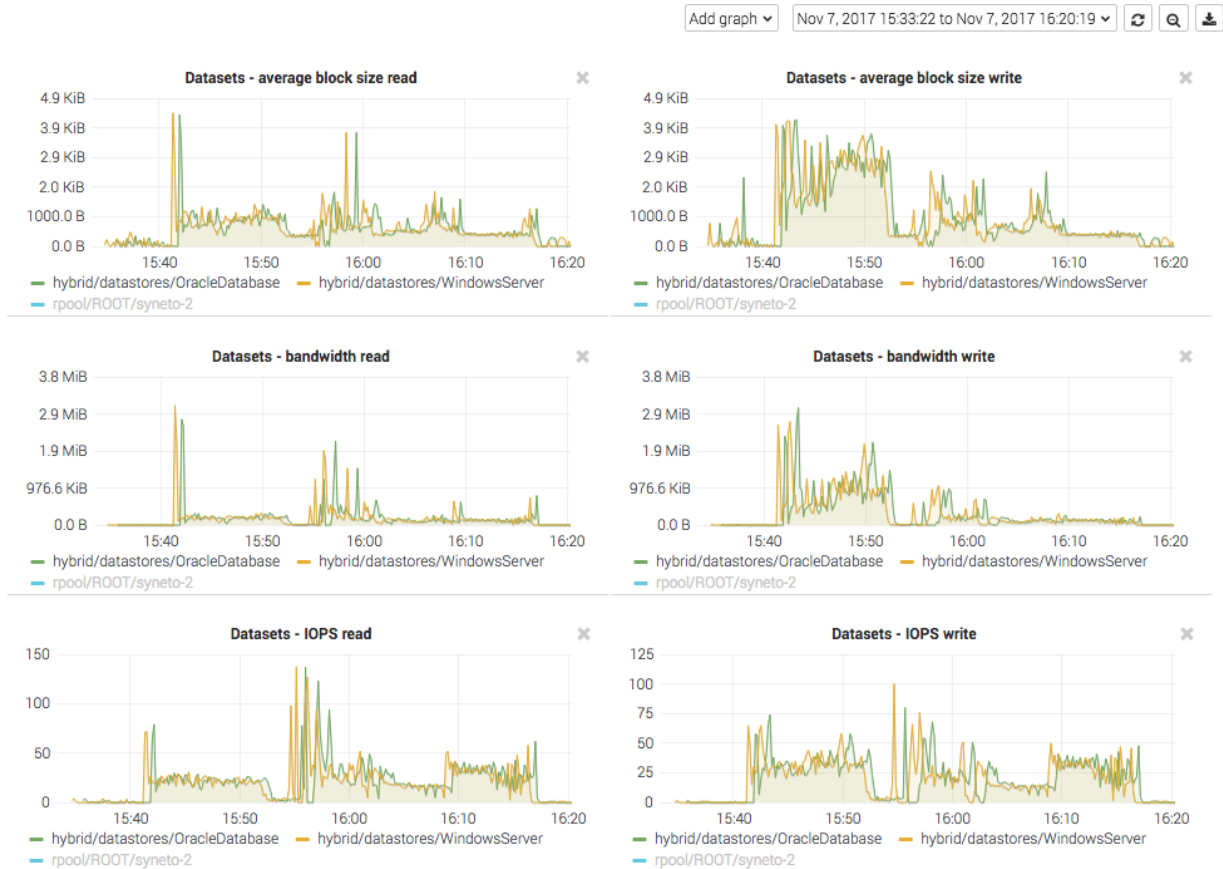
Go to Main menu → Analytics, click Add graph, select Dataset → IOPS.

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.

Search 1-4 of 4

Name	CPU	Memory	Used space	Protection
Active Directory Server 2008	34.0 MHz	1.0 GB	8.4 GB	M
Email Server	23.0 MHz	538.0 MB	4 KB	NOT PROTECTED
Email Server MX	7.0 MHz	24.0 MB	4 KB	NOT PROTECTED
Oracle Server	-	-	3 KB	NOT PROTECTED

Active Directory Server 2008

Datastore	WindowsServer	CPU	<div style="width: 1%;"><div style="width: 1%;"></div></div> 1% – 34.0 MHz of 2.5 GHz
Network adapters	1	Memory	<div style="width: 100%;"><div style="width: 100%;"></div></div> 100% – 1.0 GB of 1.0 GB
IP address	Not available	Storage	<div style="width: 20%;"><div style="width: 20%;"></div></div> 20% – 8.4 GB of 42.7 GB
Guest OS	Microsoft Windows Server 2008 R2 (64-bit)		
Protection	M	Snapshots	6
			Browse

For all these, and many more cases, you can simply click Browse in the snapshots section of the detailed view of a virtual machine.

Snapshots

November 2017

Mo	Tu	We	Th	Fr	Sa	Su
30	31	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	1	2	3
4	5	6	7	8	9	10

November 9, 2017 6 snapshots

11:12	crash consistent	87.0 KB
11:11	crash consistent	87.0 KB
11:10	crash consistent	87.0 KB
11:09	crash consistent	87.0 KB
11:08	crash consistent	87.0 KB
11:07	crash consistent	87.0 KB

Summary

VM name	Active Directory Server 2008
Oldest snapshot on	2017/11/09
Total snapshots	6

Close

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

Snapshots

November 2017

Mo	Tu	We	Th	Fr	Sa	Su
30	31	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	1	2	3
4	5	6	7	8	9	10

November 9, 2017 6 snapshots

11:12	crash consistent	87.0 KB
11:11	crash consistent	87.0 KB
11:10	crash consistent	87.0 KB
11:09	crash consistent	87.0 KB
11:08	crash consistent	87.0 KB
11:07	crash consistent	87.0 KB

Summary

VM name	Active Directory Server 2008
Oldest snapshot on	2017/11/09
Total snapshots	6

Close

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

Clone virtual machine

Active Directory Server 2008
Snapshot date: 2017/11/09 11:12

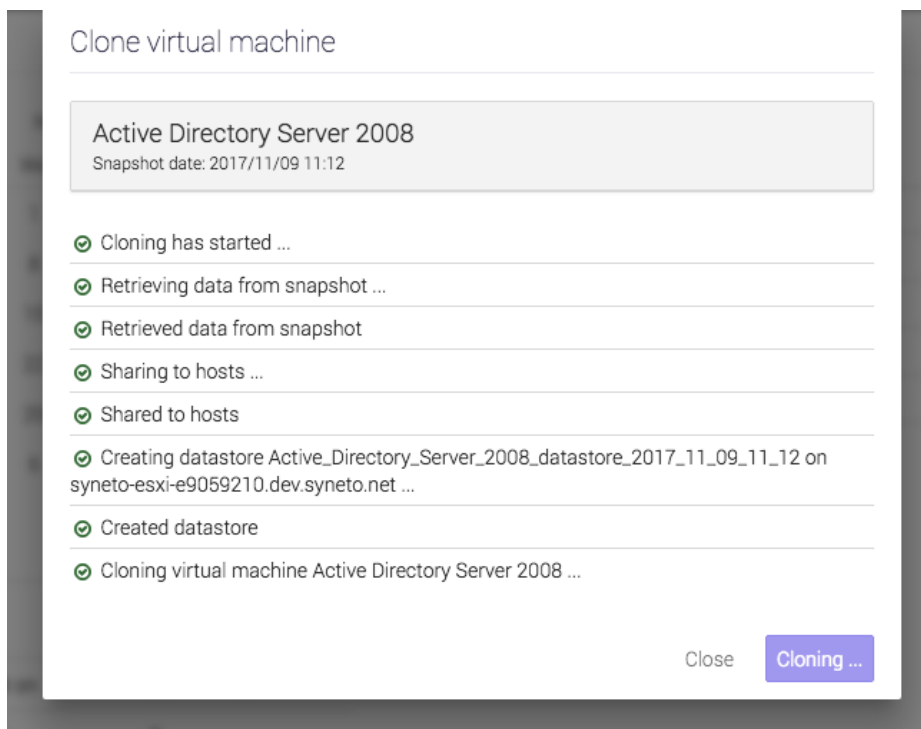
Clone virtual machine as
Active_Directory_Server_2008_2017_11_09_11_12

Access on host
syneto-esxi-e9059210.dev.syneto.net

Clone datastore as
Active_Directory_Server_2008_datastore_2017_11_09_11_

Close 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.

Search 1—5 of 5

Name	CPU	Memory	Used space	Protection
Active Directory Server 2008	40.0 MHz	1.0 GB	8.4 GB	M
Active_Directory_Server_2008_2017_11_09_11_12	-	-	5.8 GB	NOT PROTECTED
Email Server	24.0 MHz	538.0 MB	4 KB	NOT PROTECTED
Email Server MX	12.0 MHz	24.0 MB	4 KB	NOT PROTECTED
Oracle Server	-	-	3 KB	NOT PROTECTED

Active_Directory_Server_2008_2017_11_09_11_12

Datastore	Active_Directory_Server_2008_datastore_2017_11_09_11_12	CPU	0% — 0.0 Hz of 2.5 GHz
Network adapters	1	Memory	0% — 0.0 B of 1.0 GB
IP address	Not available	Storage	14% — 5.8 GB of 41.2 GB
Guest OS	Microsoft Windows Server 2008 R2 (64-bit)		
Protection	NOT PROTECTED		
Snapshots	0		

The screenshot displays the 'Virtual Machines' management interface for a host named 'syneto-esxi-e9059210.dev.syneto.net'. At the top, there are control buttons for 'Create / Register VM', 'Console', 'Power on', 'Power off', 'Suspend', 'Refresh', and 'Actions', along with a search bar. Below this is a table listing several virtual machines:

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

Below the table, a 'Quick filters...' dropdown is visible, and a '6 items' indicator is shown. The main content area features a detailed view for the selected VM 'Active_Directory_Server_2008_2017_11_09_11_12'. This view includes a thumbnail of the Windows logo, a list of properties (Guest OS: Microsoft Windows Server 2008 R2 (64-bit), Compatibility: ESXi 6.5 and later (VM version 13), VMware Tools: No, CPUs: 1, Memory: 1 GB), and resource usage indicators for CPU (0 MHz), MEMORY (0 B), and STORAGE (5.82 GB).

At the bottom, the 'Recent tasks' section provides a log of operations:

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.

Search
1-5 of 5

Name	CPU	Memory	Used space	Protection
Active Directory Server 2008	40.0 MHz	1.0 GB	8.4 GB	M
Active_Directory_Server_2008_2017_11_09_11_12	-	-	5.8 GB	NOT PROTECTED
Email Server	24.0 MHz	538.0 MB	4 KB	NOT PROTECTED
Email Server MX	12.0 MHz	24.0 MB	4 KB	NOT PROTECTED
Oracle Server	-	-	3 KB	NOT PROTECTED

Active Directory Server 2008

Datastore	WindowsServer	CPU	<div style="display: flex; align-items: center;"> <div style="width: 100px; height: 10px; background: linear-gradient(to right, #ccc, #ccc);"></div> <div style="margin-left: 5px;">2% - 40.0 MHz of 2.5 GHz</div> </div>
Network adapters	1	Memory	<div style="display: flex; align-items: center;"> <div style="width: 100px; height: 10px; background: linear-gradient(to right, #ccc, #ccc);"></div> <div style="margin-left: 5px;">100% - 1.0 GB of 1.0 GB</div> </div>
IP address	Not available	Storage	<div style="display: flex; align-items: center;"> <div style="width: 100px; height: 10px; background: linear-gradient(to right, #ccc, #ccc);"></div> <div style="margin-left: 5px;">20% - 8.4 GB of 42.7 GB</div> </div>
Guest OS	Microsoft Windows Server 2008 R2 (64-bit)		
Protection	M	Storage	
Snapshots	6	Browse	

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.

Snapshots

November 2017

Mo	Tu	We	Th	Fr	Sa	Su
30	31	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	1	2	3
4	5	6	7	8	9	10

November 9, 2017 7 snapshots

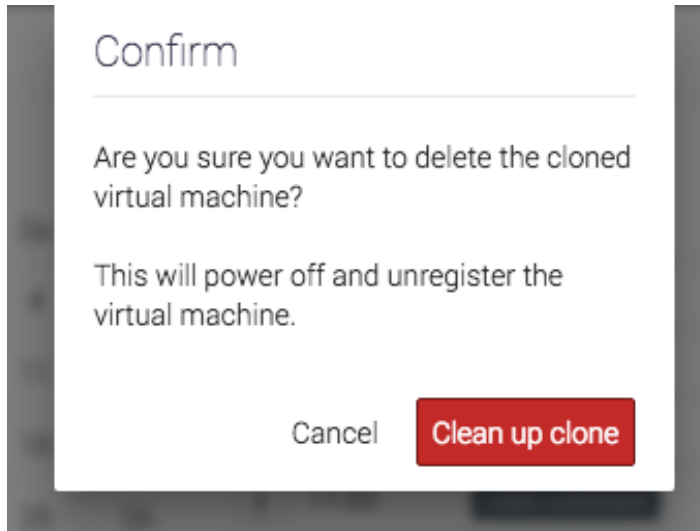
⋮	11:35	crash consistent	87.0 KB
⋮	11:34	crash consistent	87.0 KB
⋮	11:33	crash consistent	87.0 KB
⋮	11:32	crash consistent	87.0 KB
⋮	11:31	crash consistent	87.0 KB
⋮	11:30	crash consistent	97.0 KB
⋮	11:12	crash consistent	721.0 KB
<div style="border: 1px solid #ccc; padding: 2px; display: inline-block;">Clean up clone</div>			

Summary

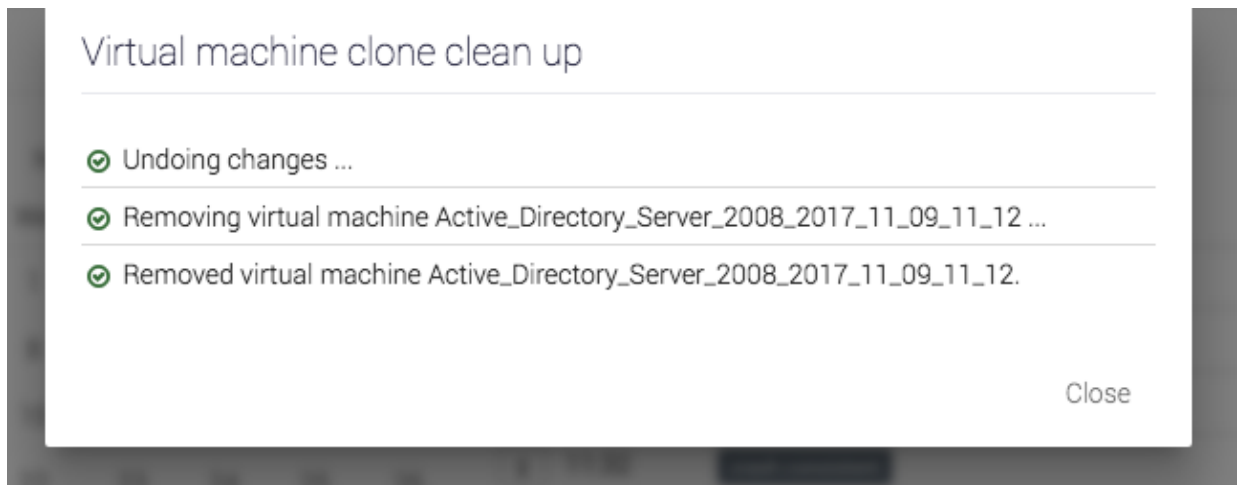
VM name	Active Directory Server 2008
Oldest snapshot on	2017/11/09
Total snapshots	7

[Close](#)

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.

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 context: 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.

The screenshot shows a 'Join to' dialog box with two tabs: 'Workgroup' (which has a 'joined' status) and 'Active Directory'. The 'Active Directory' tab is selected, and the following fields are visible:

- Domain name:** DN.MYCOMPANY (with a help icon)
- Domain controller:** dc.dn.mycompany (with a help icon)
- Username:** Administrator
- Password:** (masked with dots)
- LAN manager authentication level:** Radio buttons for levels 0, 1, 2, 3, 4 (selected), and 5 (with a help icon). Below this, it says: "Select lmauth level 4 if using Windows 2008 SP2."
- Limit organizational units access:** An unchecked checkbox.

At the bottom right, there are 'Cancel' and 'Join' buttons.

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.

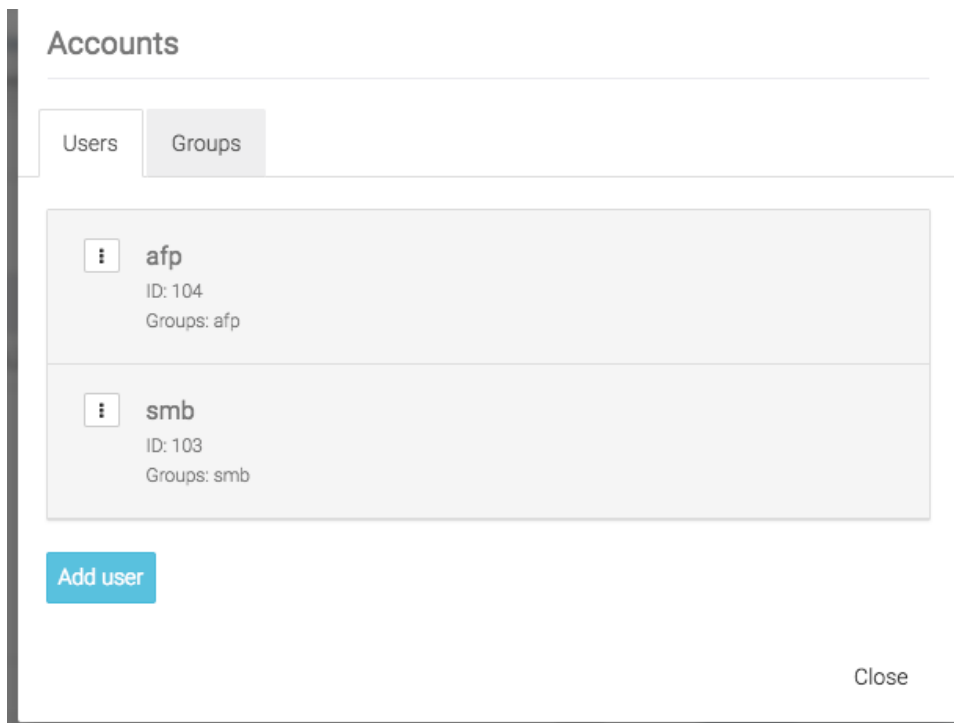
Managing local users and groups

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.

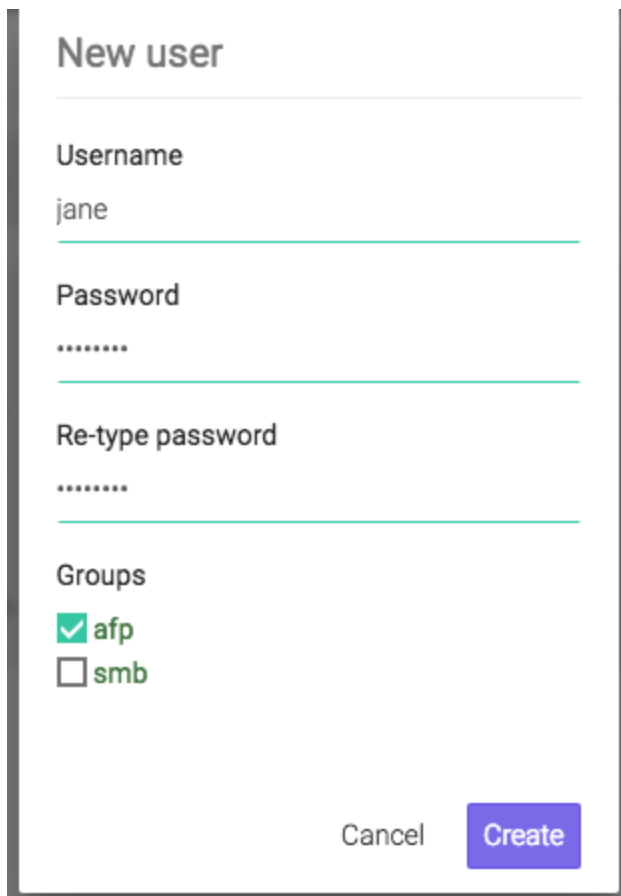
When you are joined to a workgroup, you can manage local users and groups. Click on Accounts in the action bar. The dialog will present two tabs: Users and Groups.



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

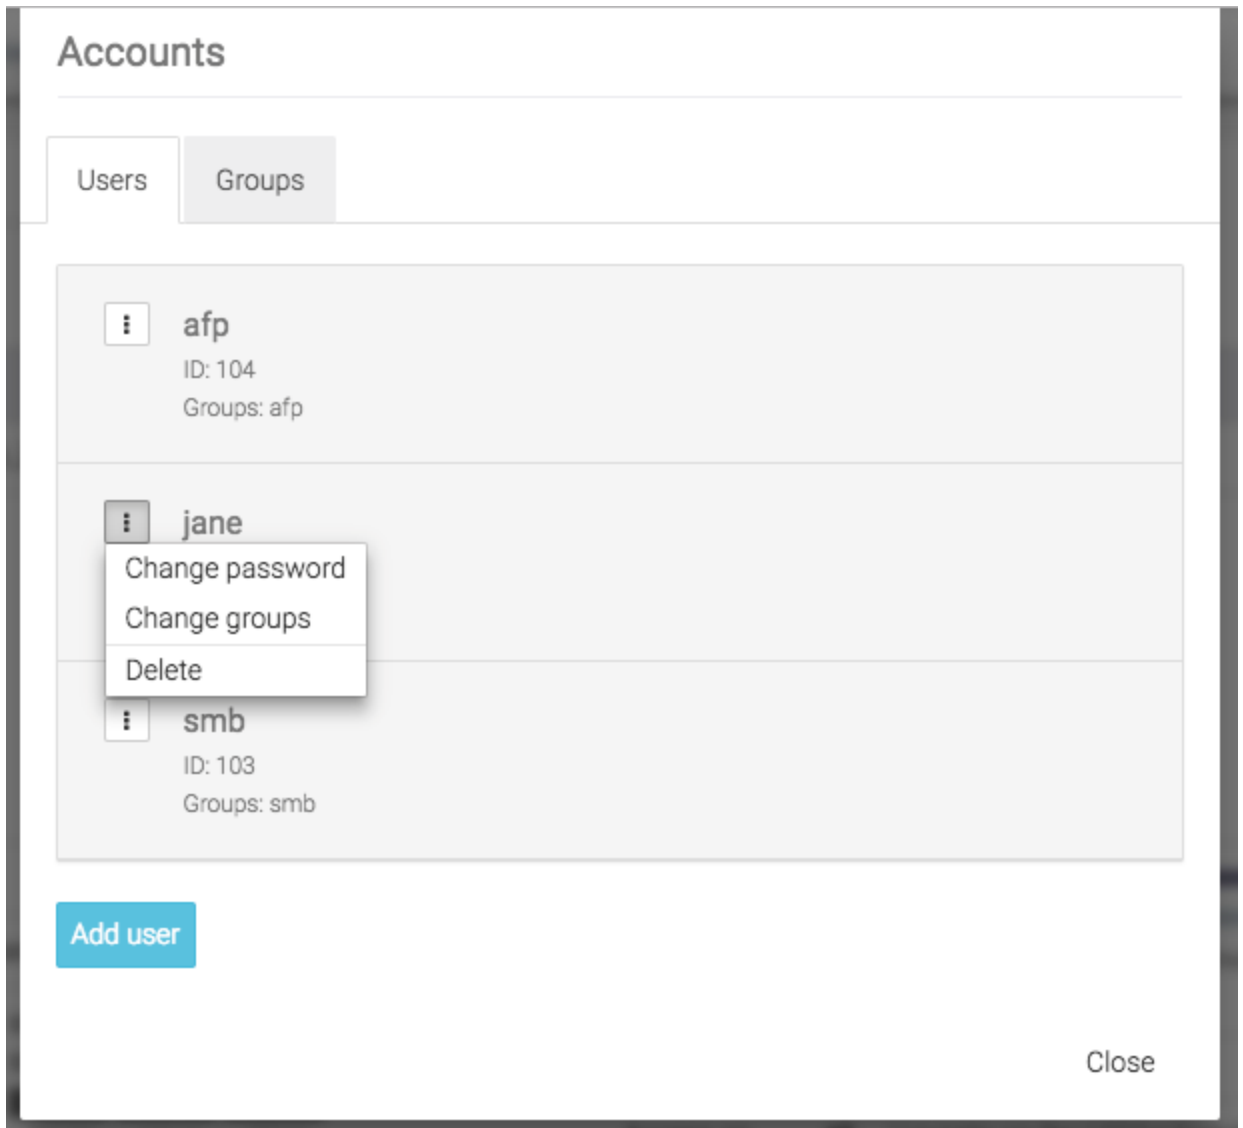
afp

smb

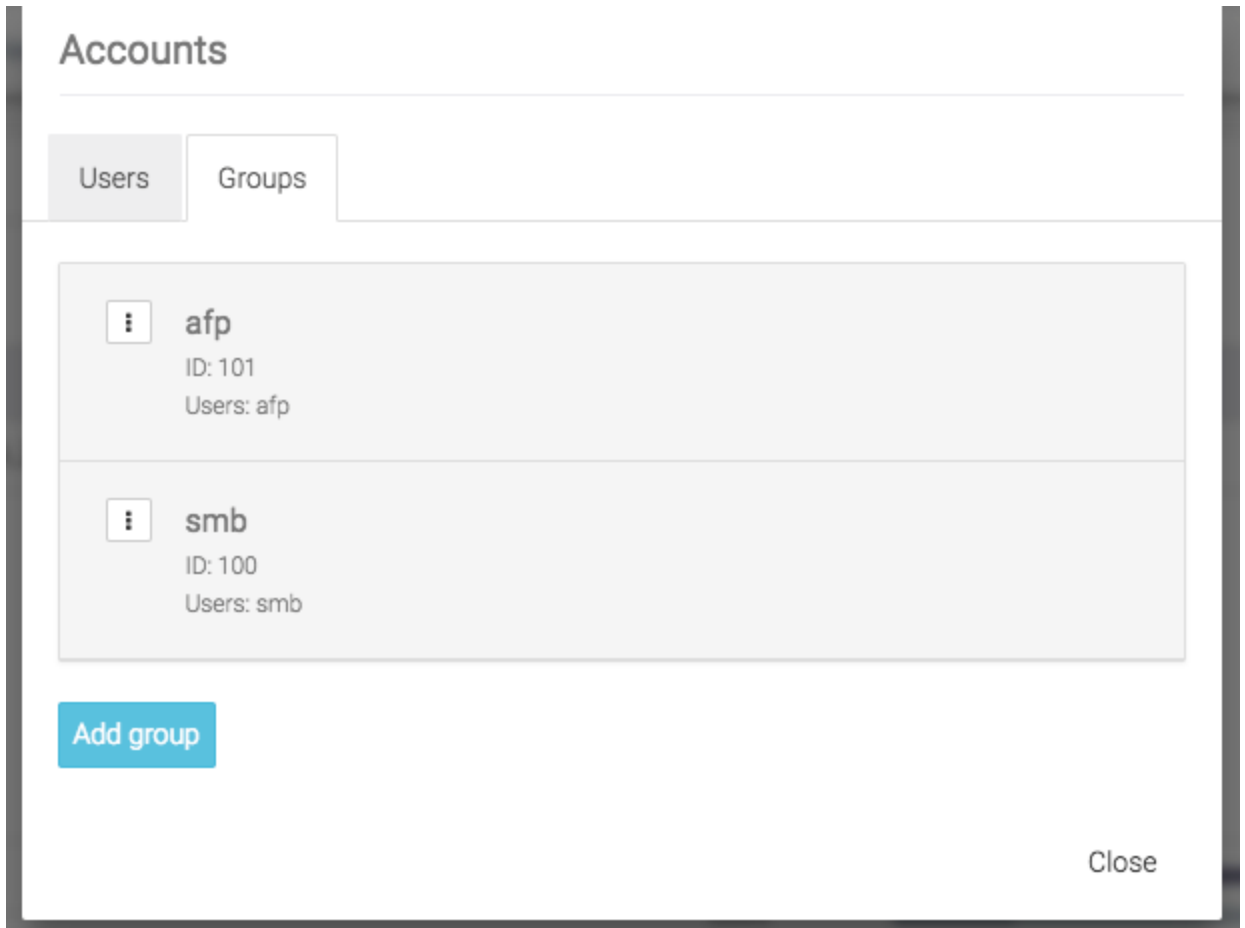
Cancel Create

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.

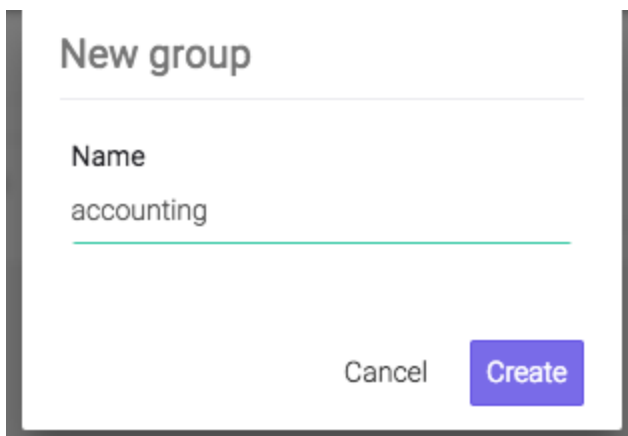
In the user list you can click the context menu and manage the user's groups, password, and delete the user.



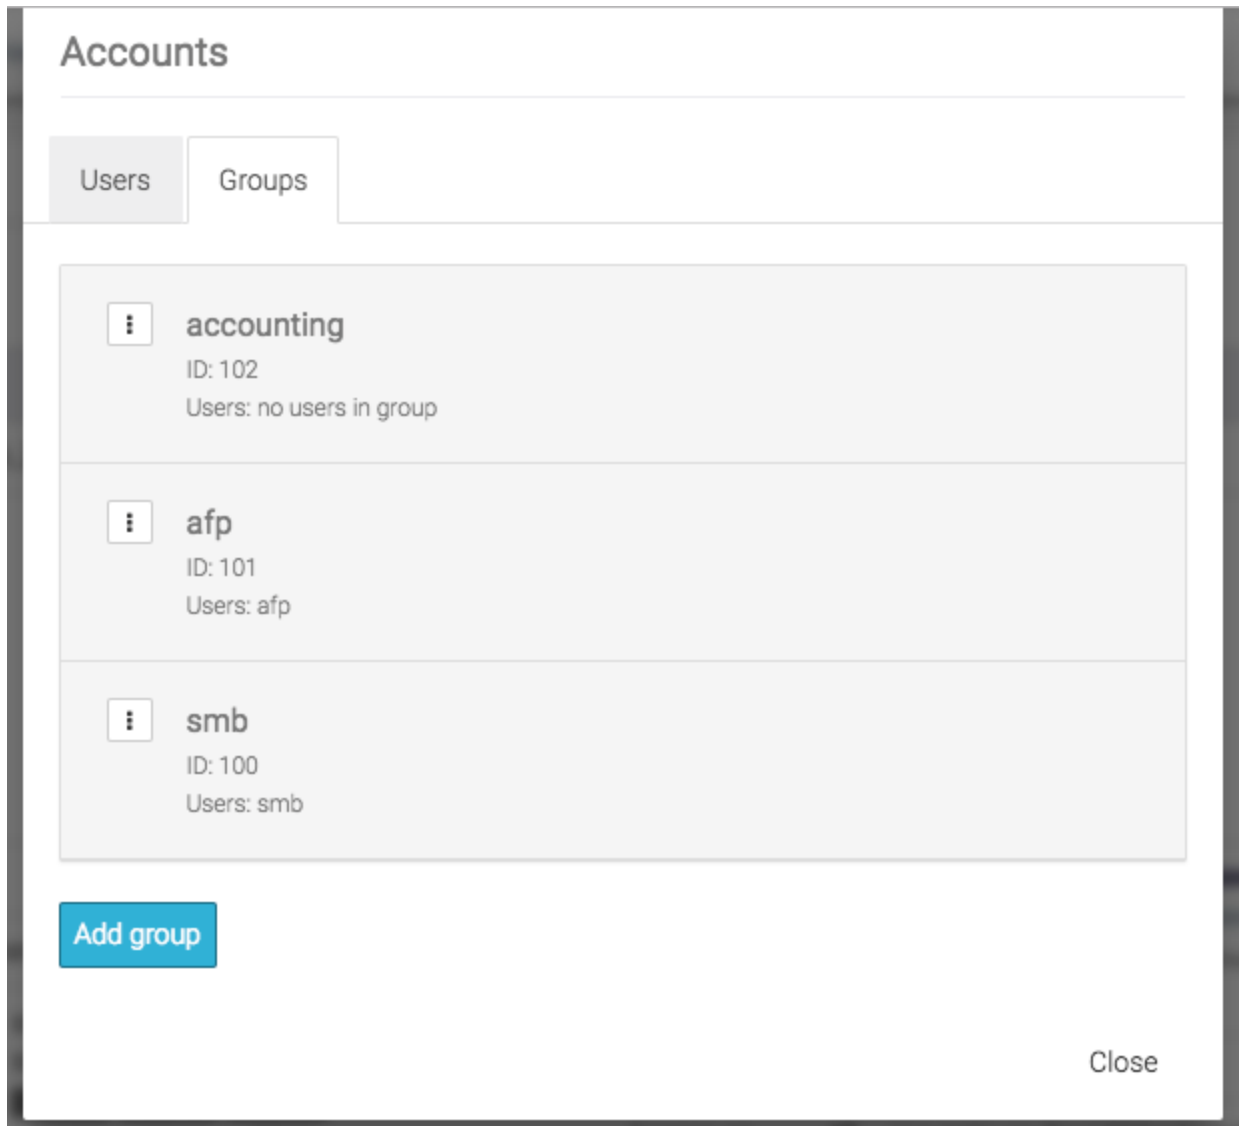
Adding and removing groups is similar to users.



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



And then, click Create to create the group.

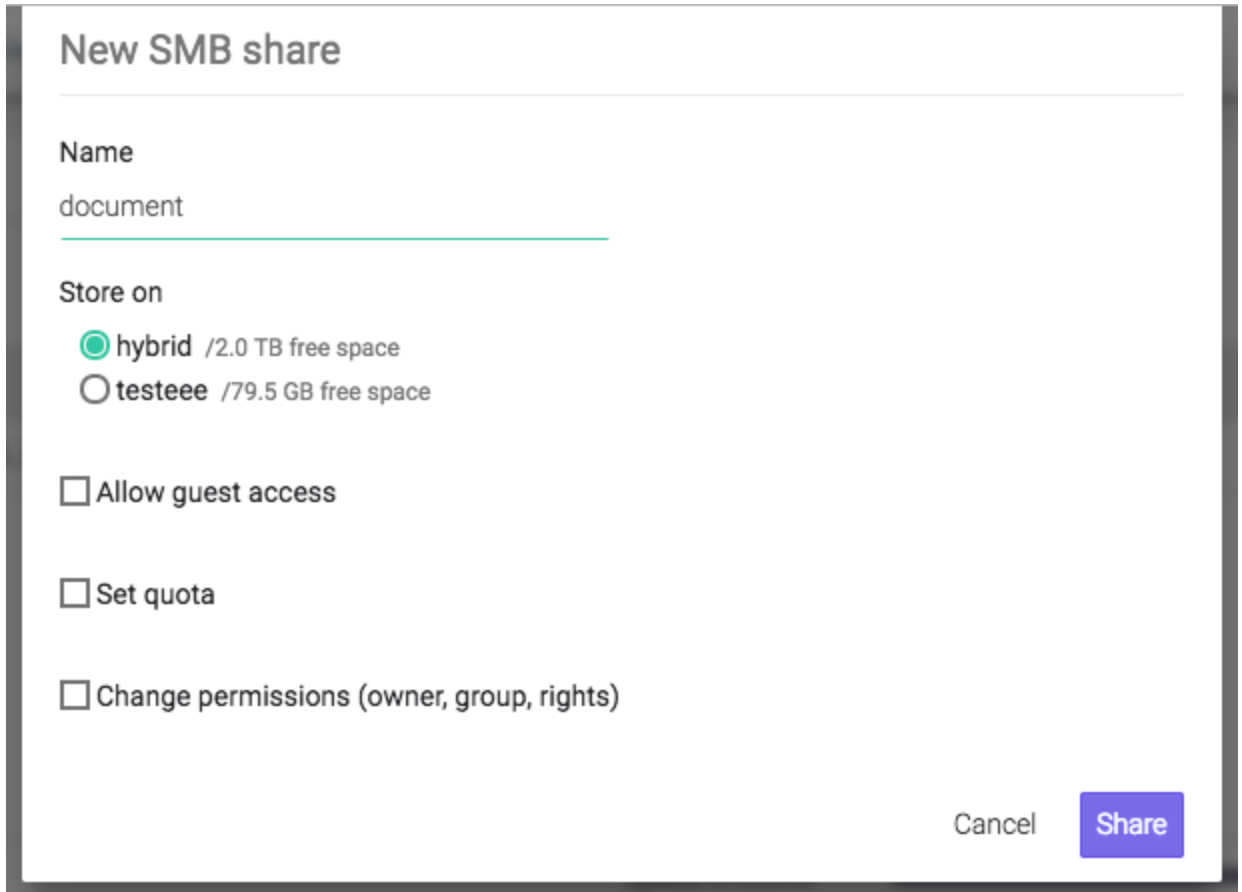


The new group will appear in the list.

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 1-1 of 1

Show Originals, Clones, » Replicas

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

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

Pool	hybrid	Space effective used	<div style="width: 100%;"><div style="width: 100%;"></div></div>	12.00 KB
Type	SMB share		<div style="width: 100%;"><div style="width: 100%;"></div></div>	24.00 KB
Permissions	Owner smb Group smb Rights rwx rw- r--	Change	Compression 1.00x	Saved 0.0 B
Quota	1.0 GB	Change	Access on \\syneto-os-0ec1027e\document 📄	
Protection	NOT PROTECTED	Protect	smb://syneto-os-0ec1027e/document 📄	
Snapshots	none		Guest access OFF <input checked="" type="checkbox"/> ON	
		Description	[2018-07-13 18:04:35] 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 shows the SynetoOS interface. At the top, there is a search bar and a dropdown menu set to 'Show Originals, Clones, » Replicas'. Below this is a table with columns: Name, Type, Used space, Quota, Pool, and Protection. A single row is visible for a share named 'document' with Type 'SMB', Used space '24.0 KB', Quota '1.0 GB', Pool 'hybrid', and Protection 'NOT PROTECTED'. Below the table, the detailed view for the 'document' share is shown. It includes fields for Pool (hybrid), Type (SMB share), Permissions (Owner smb, Group smb, Rights rwx for owner, rw- for group, r-- for others), Quota (1.0 GB), Protection (NOT PROTECTED), and Snapshots (none). On the right side, there are sections for Space effective used (12.00 KB for effective, 24.00 KB for total), Compression (1.00x, Saved 0.0 B), and Guest access (OFF/ON toggle). A description field contains the text: '[2018-07-13 18:04:35] 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, » Replicas 1-2 of 2

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

video Pause Delete

Pool	hybrid	Space effective used	<div style="width: 100%;"><div style="width: 100%;"></div></div>	12.00 KB
Type	AFP share		<div style="width: 100%;"><div style="width: 100%;"></div></div>	23.00 KB
Permissions	Owner afp Group afp Rights rwx rw- r--	Change	Compression 1.04x	Saved 0.0 B
Quota	200.0 GB	Change	Access on afp://syneto-os-0ec1027e/video	
Protection	NOT PROTECTED	Protect	Time machine OFF <input checked="" type="checkbox"/> ON	
Snapshots	none		Spotlight search OFF <input checked="" type="checkbox"/> ON	
		Description	[2018-07-13 18:14:12] 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.

The screenshot displays the SynetoOS interface. At the top, there is a search bar and navigation options: 'Show Originals, Clones, » Replicas'. The page number '1-2 of 2' is in the top right corner. Below this is a table of shares:

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

The 'video' share is selected, and its details are shown below:

video Resume Delete

Pool	hybrid	Space effective used		12.00 KB 23.00 KB					
Type	AFP share								
Permissions	Owner afp Group afp Rights <table border="1"><tr><td>rwx</td><td>rw-</td><td>---</td></tr><tr><td>owner</td><td>group</td><td>others</td></tr></table>	rwx	rw-	---	owner	group	others	Change	Compression 1.04x Saved 0.0 B
rwx	rw-	---							
owner	group	others							
Quota	200.0 GB	Change	Time machine	OFF <input type="checkbox"/> ON <input checked="" type="checkbox"/>					
Protection	NOT PROTECTED	Protect	Spotlight search	OFF <input type="checkbox"/> ON <input checked="" type="checkbox"/>					
Snapshots	none		Description	[2018-07-13 18:14:12] 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 navigation options: 'Show Originals, Clones, > Replicas'. The main area shows a table of shares with columns for Name, Type, Used space, Quota, and Protection. The 'jane-backup' share is highlighted in blue. Below the table, a detailed view for the 'jane-backup' share is shown, including its type (AFP share), permissions (Owner: afp, Group: afp, Rights: rwx for owner, rw- for group, r-- for others), space effective used (88.50 KB data, 138.00 KB snapshots), quota (unlimited), access on (afp://syneto-os-0ec1027e/jane-backup), time machine (OFF), spotlight search (OFF), and a description: '[2018-07-16 14:59:54] AFP share jane-backup created.'

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
jane-backup	AFP	171.0 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

Type	AFP share	Space effective used	88.50 KB
Permissions	Owner afp Group afp Rights rwx owner rw- group r-- others		171.00 KB
Quota	unlimited	Compression 1.03x	Saved 0.0 B
Protection	M	Access on	afp://syneto-os-0ec1027e/jane-backup
Snapshots	6	Time machine	OFF ON
		Spotlight search	OFF ON
		Description	[2018-07-16 14:59:54] AFP share jane-backup created.

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

Type

Specifies 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

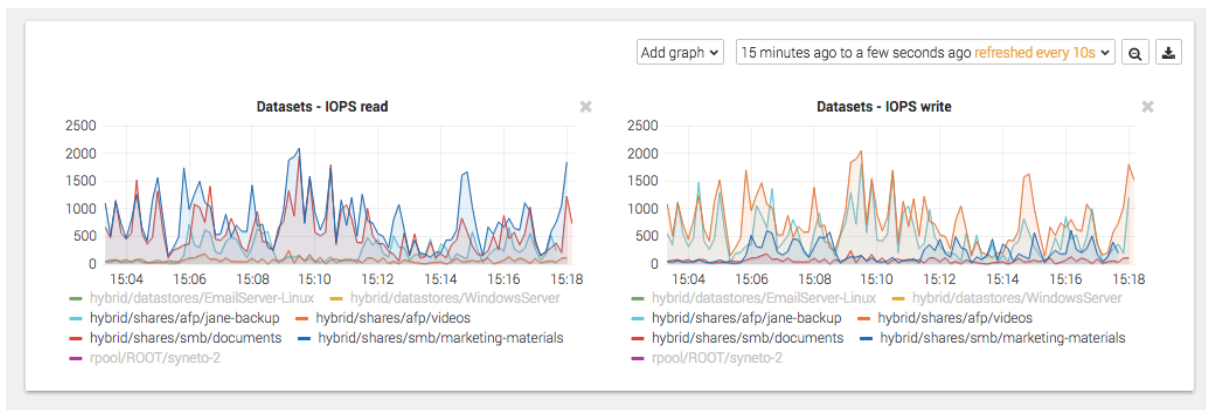
Shows the quota set on the share. It can be used along with Used space to determine if space limits are reached.

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

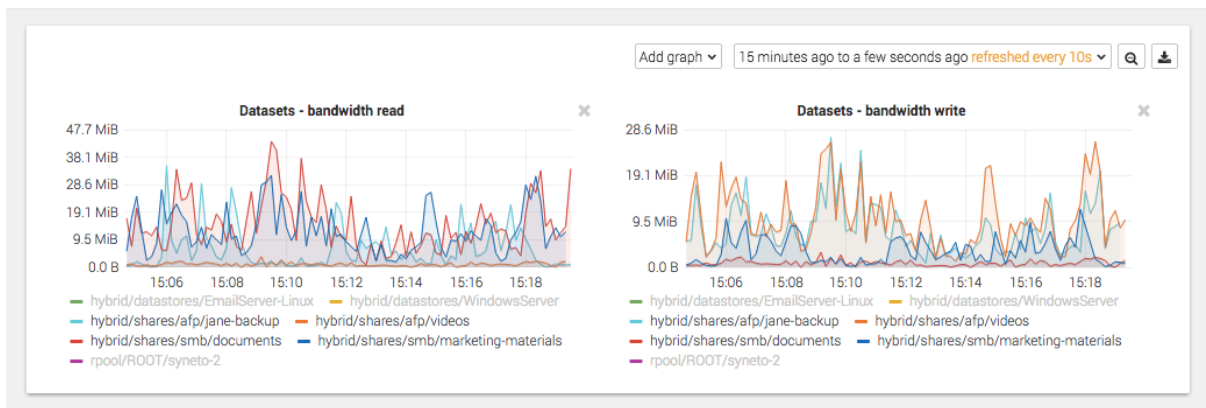
Go to Main menu → Analytics, click Add graph, select Dataset → IOPS.



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.

Search
Show Originals, Clones, » Replicas
1-4 of 4

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
jane-backup	AFP	171.0 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

Type	AFP share	Space effective used	<div style="display: flex; align-items: center;"> <div style="width: 100px; height: 10px; background: linear-gradient(to right, #007bff, #6c757d);"></div> 88.50 KB </div> <div style="display: flex; align-items: center;"> <div style="width: 100px; height: 10px; background: linear-gradient(to right, #007bff, #6c757d);"></div> 171.00 KB </div> <div style="margin-left: 10px; font-size: 8px;"> ■ Data: 33.00 KB — ■ Snapshots: 138.00 KB </div>
Permissions	Owner afp Group afp Rights rwx rw- r-- <small>owner group others</small>	Change	Compression 1.03x Saved 0.0 B
Quota	unlimited	Change	Access on <code>afp://syneto-os-0ec1027e/jane-backup</code> 📄
Protection	M	Protect	Time machine OFF <input checked="" type="checkbox"/> ON
Snapshots	6	Browse	Spotlight search OFF <input checked="" type="checkbox"/> ON
		Description	[2018-07-16 14:59:54] AFP share jane-backup created. Edit

For all these, and many more cases, you can simply click Browse in the snapshots section of the detailed view of a share.

Snapshots

July 2018

Mo	Tu	We	Th	Fr	Sa	Su
25	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

July 16, 2018 6 snapshots

- ▶ ⓘ 12:53 24.0 KB
- ▶ ⓘ 12:52 24.0 KB
- ▶ ⓘ 12:51 24.0 KB
- ▶ ⓘ 12:50 24.0 KB
- ▶ ⓘ 12:49 24.0 KB
- ▶ ⓘ 12:48 24.0 KB

Summary

Share	jane-backup
Oldest snapshot on	2018/07/16
Total snapshots	6

Close

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

Snapshots

< July 2018 >

Mo	Tu	We	Th	Fr	Sa	Su
25	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

July 16, 2018 6 snapshots

▶ ⓘ 12:53	24.0 KB
▶ ⓘ 12:52	24.0 KB
▶ ⓘ 12:51	24.0 KB
▶ ⓘ 12:50	24.0 KB
▶ ⓘ 12:49	24.0 KB
▶ ⓘ 12:48	24.0 KB

Clone
Delete

Summary

Share	jane-backup
Oldest snapshot on	2018/07/16
Total snapshots	6

Close

Select clone.

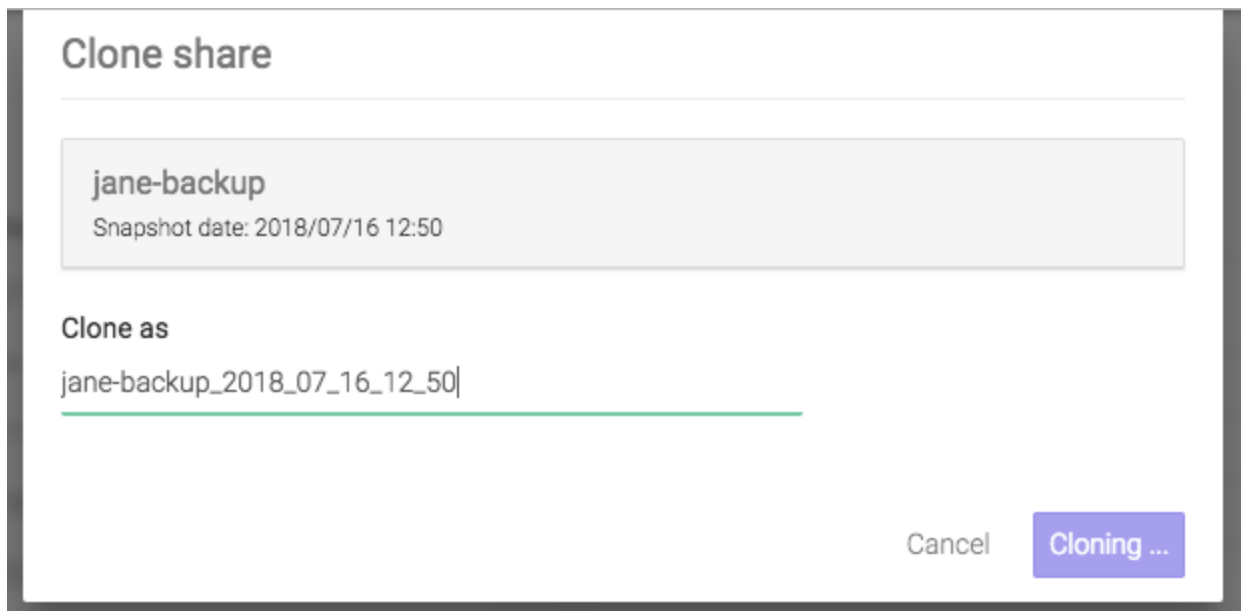
Clone share

jane-backup
Snapshot date: 2018/07/16 12:51

Clone as
jane-backup_2018_07_16_12_51|

Cancel **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 HYPHER 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 HYPHER snapshots.

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

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 HYPHER 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.

The screenshot displays the SynetoOS Shares management interface. At the top, there is a search bar and a filter dropdown set to 'Show Originals, Clones, > Replicas'. A table lists several shares:

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
jane-backup	AFP	178.0 KB	unlimited	M
jane-backup_2018_07_16_12_50	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

The 'jane-backup' share is selected, and its detailed settings are shown below:

- Name:** jane-backup (with 'Pause' and 'Delete' actions)
- Type:** AFP share
- Space effective used:** 92.50 KB (Data: 34.00 KB, Snapshots: 144.00 KB)
- Permissions:** Owner: afp, Group: afp, Rights: owner (rwx), group (rw-), others (r--)
- Quota:** unlimited
- Protection:** M (with 'Protect' action)
- Snapshots:** 6 (with 'Browse' action)
- Access on:** afp://syneto-os-0ec1027e/jane-backup
- Time machine:** OFF → ON
- Spotlight search:** OFF → ON
- Description:** [2018-07-16 12:30:50] AFP share jane-backup created. (with 'Edit' action)
- Clones:**

Clone	Cloned on
jane-backup_2018_07_16_12_50	2018/07/16 12:50

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

Snapshots

July 2018

Mo	Tu	We	Th	Fr	Sa	Su
25	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

Summary

Share	jane-backup
Oldest snapshot on	2018/07/16
Total snapshots	6

July 16, 2018 6 snapshots

- ▶ ⓘ 12:53 24.0 KB
- ▶ ⓘ 12:52 24.0 KB
- ▶ ⓘ 12:51 24.0 KB
- ▼ ⓘ 12:50

cloned

Starting automated schedule ...
 12:50:05 AFP Share name: jane-backup ...
 12:50:05 Frequency: every 1 minutes ...
 12:50:05 Host: syneto-os-0ec1027e ...
 12:50:07 Taking AFP Share snapshot ...
 12:50:07 Local snapshot taken.

24.0 KB
- ▶ ⓘ 12:49 24.0 KB
- ▶ ⓘ 12:48 24.0 KB

Close

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

Confirm

Are you sure you want to delete the cloned share?

Cancel
Delete clone

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 the SynetoOS interface with a list of clones. The clone 'jane-backup_2018_07_16_12_50' is selected. Below the list, 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), quota (unlimited), and space effective used (7.00 KB). The clone is noted as being cloned from 'jane-backup (2018/07/16 12:50)'. The interface also shows options to Revert, Pause, and Delete the clone.

Name	Type	Used space	Quota	Protection
document	SMB	24.0 KB	1.0 GB	NOT PROTECTED
jane-backup	AFP	178.0 KB	unlimited	M
jane-backup_2018_07_16_12_50	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_07_16_12_50 [Revert](#) [Pause](#) [Delete](#)

Cloned from *jane-backup* (2018/07/16 12:50)

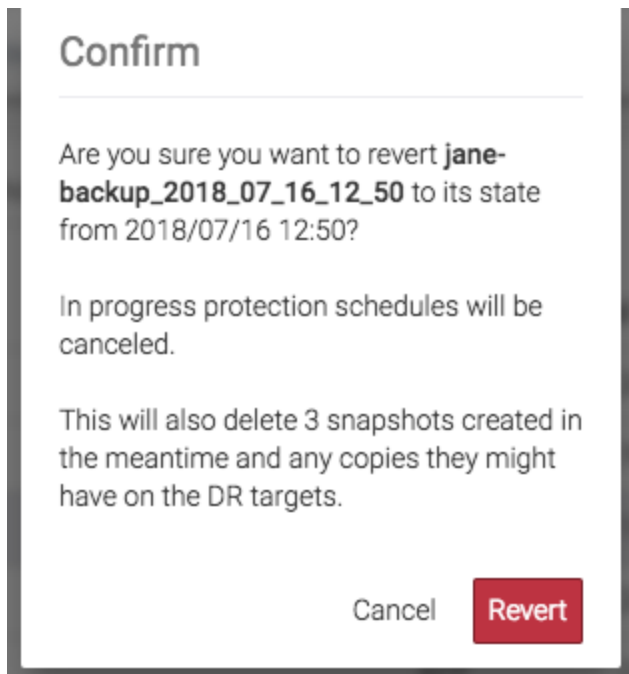
Type	AFP share	Space effective used	7.00 KB
Permissions	Owner afp Group afp Rights rwx rw- r-- owner group others	used	14.00 KB
Quota	unlimited	Compression 1.00x	Saved 0.0 B
		Access on	afp://syneto-os-0ec1027e/jane-backup_2018_
		Time machine	OFF <input checked="" type="checkbox"/> ON
		Spotlight search	OFF <input checked="" type="checkbox"/> ON
		Description	[2018-07-16 14:30:55] AFP share jane-backup_2018_07_16_12_50 created.

A clone revert process will do the following steps:

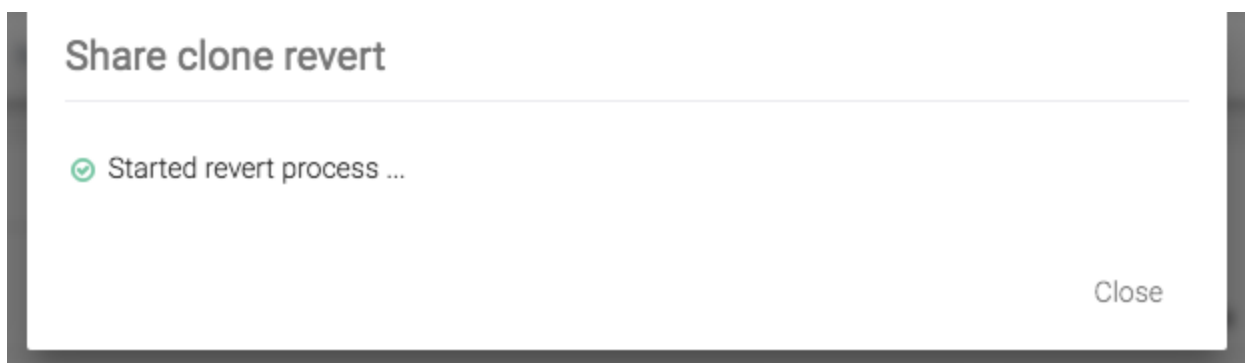
- Delete all snapshots created after the snapshot from which the clone was created
- Replace the original share with the clone

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	Edit
	[2018-07-16 15:12:38] AFP share jane-backup_2018_07_16_12_50 created.	

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.

Initiators (remote)

No initiators were found.

Add initiator

Close

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

Add initiator

Group	New group name
Create new group	▼ Initiators-Group-1

Identifier

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

Cancel

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

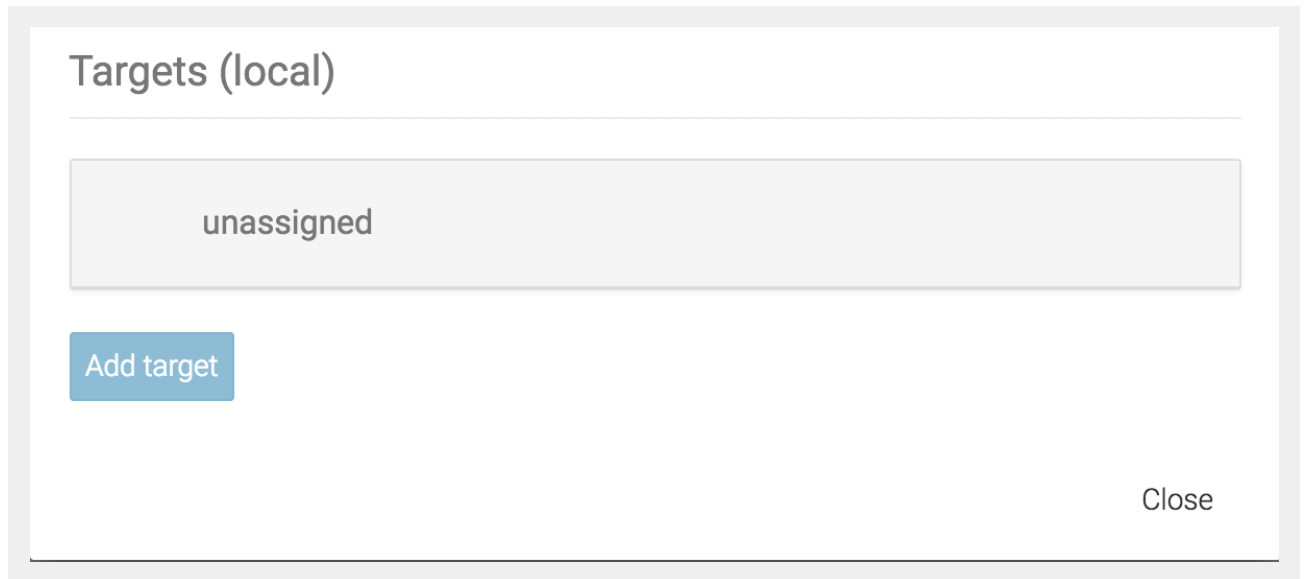
Once you are done, click Add. The new initiator will be shown in the list. From this list you can manage the initiators, move them from one group to another, or delete them.

iSCSI local targets

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

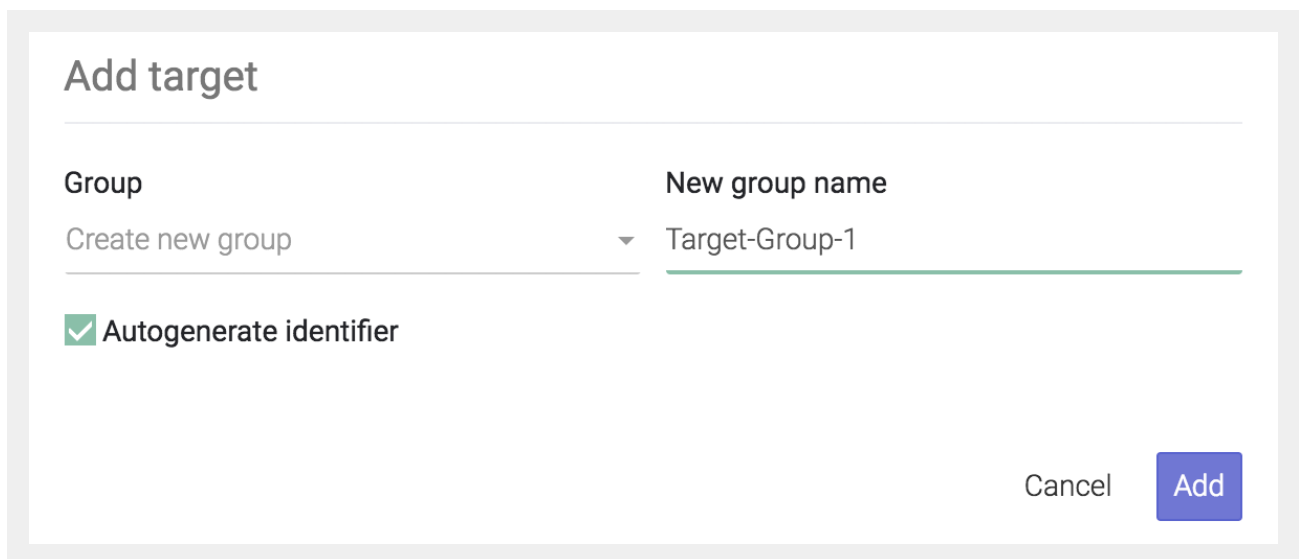
In case you wish more granular access control, you can specify local targets and groups by going to Main menu → Shares → Targets.

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.

Targets (local)

- Target-Group-1
 - iqn.2013-02.net.syneto:02:2cb16da9-33de-45c7-a483-2db779367d57 OFF ON

unassigned

Add target

Close

Once you are done, click Add. The new target will be show in the list. From this list you can manage the targets, move them from one group the 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 Show Originals, Clones, » Replicas 1-1 of 1

Name	Used space	Provisioning	Protection
Volume1	12.0 KB	thin	NOT PROTECTED

Volume1 Delete

Size	2.0 GB	Expand	Space effective used		6.00 KB 12.00 KB
Block size	32.0 KB				
Provisioning	THICK <input type="radio"/> THIN <input checked="" type="radio"/>			Compression 1.00x	Saved 0.0 B
Write back cache	OFF <input type="radio"/> ON <input checked="" type="radio"/>			Description	[2018-07-13 12:59:21] Volume Volume1 created. Edit
Protection	NOT PROTECTED	Protect			
Snapshots	none				

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 displays the 'Volumes' page in SynetoOS. At the top, there is a search bar and a filter dropdown set to 'Show Originals, Clones, » Replicas'. The page shows 1-2 of 2 volumes. A table lists the volumes with columns for Name, Used space, Provisioning, and Protection.

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

Below the table, the configuration for 'Volume1' is shown. It includes fields for Size (2.0 GB), Block size (32.0 KB), Provisioning (THICK/THIN), Write back cache (OFF/ON), Protection (M/W), and Snapshots (2). A 'Space effective used' bar chart shows 6.00 KB of data and 12.00 KB of snapshots. The description is '[2018-07-13 12:59:21] Volume Volume1 created.' and the iSCSI access table shows Initiators-Group-1 connected to Target-Group-1 on LUN 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, » Replicas 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

Size	2.0 GB	Expand	Space effective used		6.00 KB 12.00 KB
Block size	32.0 KB				
Provisioning	THICK <input type="radio"/> THIN <input checked="" type="radio"/>			Compression 1.00x	Saved 0.0 B
Write back cache	OFF <input type="radio"/> ON <input checked="" type="radio"/>			Description	[2018-07-13 12:59:21] Volume Volume1 created. Edit
Protection	M W	Protect			
Snapshots	2	Browse			

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 **Browse** in the snapshots section of the detailed view of a volume.

Snapshots

July 2018

Mo	Tu	We	Th	Fr	Sa	Su
25	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

July 13, 2018 3 snapshots

▼ ⓘ	13:06	0.0 B
	13:06:04	Starting automated schedule ...
	13:06:04	ISCSI Volume name: Volume1 ...
	13:06:05	Frequency: every 1 minutes ...
	13:06:06	Host: syneto-os-44b46bc3 ...
	13:06:07	Taking ISCSI Volume snapshot ...
	13:06:08	Local snapshot taken.
▶ ⓘ	13:05	0.0 B
▶ ⓘ	13:04	0.0 B

Summary

Volume	Volume1
Oldest snapshot on	2018/07/13
Total snapshots	3

Close

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

Snapshots

July 2018

Mo	Tu	We	Th	Fr	Sa	Su
25	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

July 13, 2018 3 snapshots

<div style="display: flex; justify-content: space-between;"> ⌵ 13:06 0.0 B </div> <ul style="list-style-type: none"> <div style="display: flex; justify-content: space-between; align-items: flex-start;"> Clone 13:06:04 Starting automated schedule ... </div> <div style="display: flex; justify-content: space-between; align-items: flex-start;"> Delete 13:06:04 ISCSI Volume name: Volume1 ... </div> <div style="display: flex; justify-content: space-between; align-items: flex-start;"> 13:06:05 Frequency: every 1 minutes ... </div> <div style="display: flex; justify-content: space-between; align-items: flex-start;"> 13:06:06 Host: syneto-os-44b46bc3 ... </div> <div style="display: flex; justify-content: space-between; align-items: flex-start;"> 13:06:07 Taking ISCSI Volume snapshot ... </div> <div style="display: flex; justify-content: space-between; align-items: flex-start;"> 13:06:08 Local snapshot taken. </div>
<div style="display: flex; justify-content: space-between;"> ▶ 13:05 0.0 B </div>
<div style="display: flex; justify-content: space-between;"> ▶ 13:04 0.0 B </div>

Summary

Volume	Volume1
Oldest snapshot on	2018/07/13
Total snapshots	3

Close

Select Clone.

Clone volume

Volume1
Snapshot date: 2018/07/13 13:06

Clone volume as

Volume1_2018_07_13_13_06

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, » Replicas'. Below this is 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_07_13_13_06	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 the selected volume 'Volume1_2018_07_13_13_06' are shown. It is a clone of 'Volume1' (2018/07/13 13:06). The size is 2.0 GB, and it is currently 'THIN' provisioned. The write back cache is 'OFF'. The space effective used is 0.00 B. The description is '[2018-07-13 13:08:15] Volume Volume1_2018_07_13_13_06 created.' There is a 'Grant access' button for iSCSI access, which is currently not accessible.

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, » Replicas 1-3 of 3

Name	Used space	Provisioning	Protection
Volume1	12.0 KB	thin	M W
 Volume1_2018_07_13_13_06	0.0 B	thin	NOT PROTECTED
Volume2	12.0 KB	thin	M H D W

Volume1 Delete

Size	2.0 GB	Expand	Space effective used	<div style="display: flex; align-items: center;"> <div style="width: 100px; height: 10px; background: linear-gradient(to right, blue 60%, grey 60%);"></div> <div style="margin-left: 10px;">6.00 KB 12.00 KB</div> </div> <small>■ Data: 12.00 KB — ■ Snapshots: 0.00 B</small>
Block size	32.0 KB			
Provisioning	THICK <input checked="" type="radio"/> THIN			Compression 1.00x Saved 0.0 B
Write back cache	OFF <input checked="" type="radio"/> ON			Description [2018-07-13 12:59:21] Volume Volume1 created. Edit
Protection	M W	Protect		
Snapshots	5	Browse		iSCSI access Grant access

Clones		Initiator	Target	LUN
	Volume1_2018_07_13_13_06	Initiators-Group-1	Target-Group-1	1

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

Snapshots

July 2018

Mo	Tu	We	Th	Fr	Sa	Su
25	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

Summary

Volume	Volume1
Oldest snapshot on	2018/07/13
Total snapshots	6

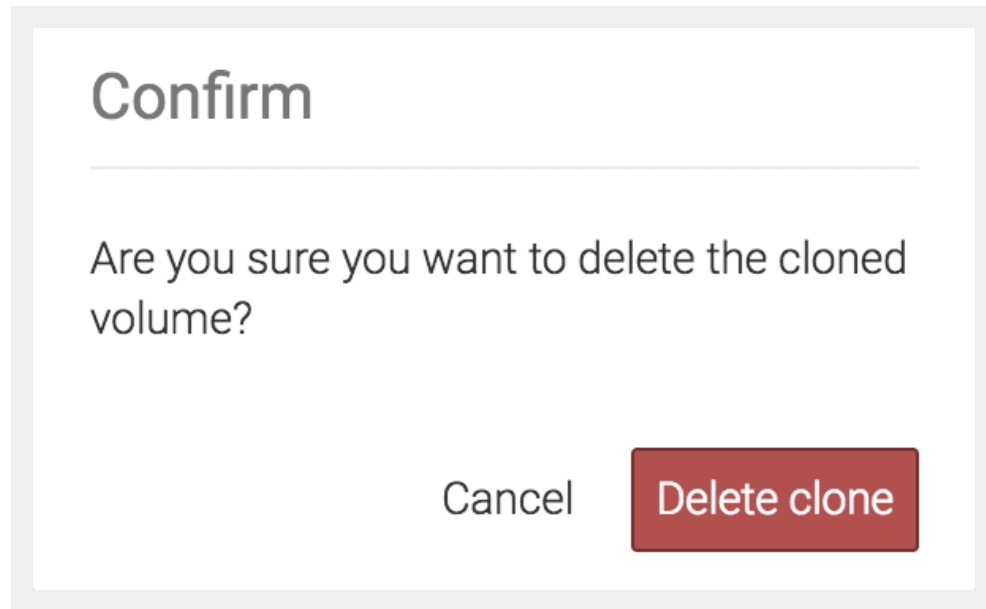
July 13, 2018

6 snapshots

▶	⋮	13:09	0.0 B
▶	⋮	13:08	0.0 B
▶	⋮	13:07	0.0 B
▼	⋮	13:06	0.0 B
		cloned	
		13:06:04	Starting automated schedule ...
		13:06:05	ISCSI Volume name: Volume1 ...
		13:06:06	Frequency: every 1 minutes ...
		13:06:07	Host: syneto-os-44b46bc3 ...
		13:06:08	Taking ISCSI Volume snapshot ...
			Local snapshot taken.
▶	⋮	13:05	0.0 B
▶	⋮	13:04	0.0 B

Close

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.

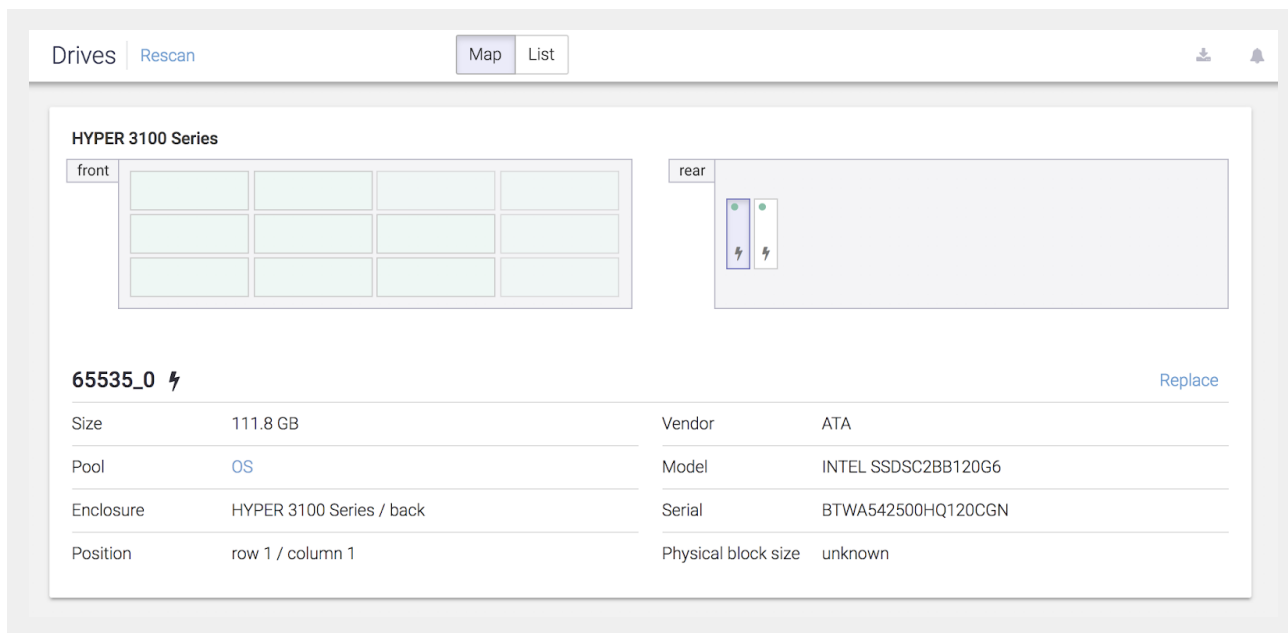
That's it, all that was created and shared for that snapshot clone was reverted. No more mess left behind by the recovery process.

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 displays the 'Drives' page in SynetoOS. At the top, there are tabs for 'Map' and 'List', with 'Map' selected. Below the tabs, there are two diagrams: 'front' and 'rear'. The 'front' diagram shows a 3x4 grid of drive bays, with the top-left bay highlighted in green. The 'rear' diagram shows two drive bays, with the left one highlighted in blue and the right one in white. Below the diagrams, the details for the selected drive '65535_0' are shown. The drive is a 111.8 GB Intel SSDSC2BB120G6 SATA drive located in row 1, column 1 of the Hyper 3100 Series enclosure. A 'Replace' button is visible in the top right corner of the details panel.

HYPER 3100 Series			
front		rear	
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 shows the 'Drives' 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 drives with columns for Name, Size, Pool, Enclosure, and Serial. The drive 'c0t5000CCA255615D54d0' is selected, and its details are shown below the table.

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	K1HS0KAD
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.

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

OS

Data 111.8 GB 93.2 GB

5.5 GB of 79.5 GB used

Healthy

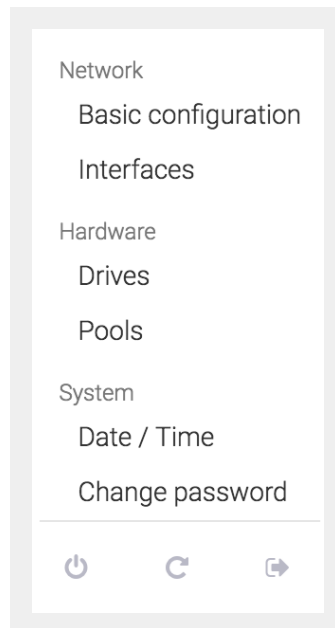
Integrity check [Check now](#)

Last check: never checked

Scheduled check: disabled

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)	

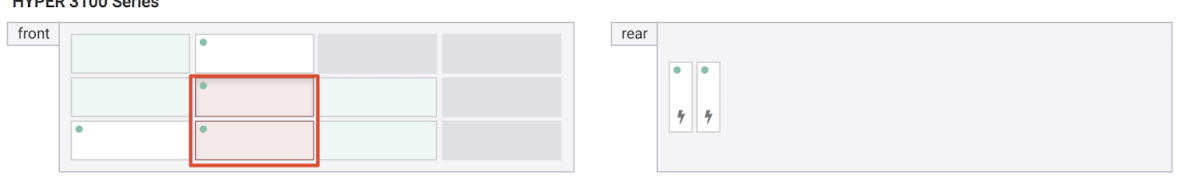
Name	Size	Used space	Scheduled check
✖ hybrid	3.6 TB	364.0 KB	in 4 days
✔ OS	79.5 GB	5.5 GB	disabled

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

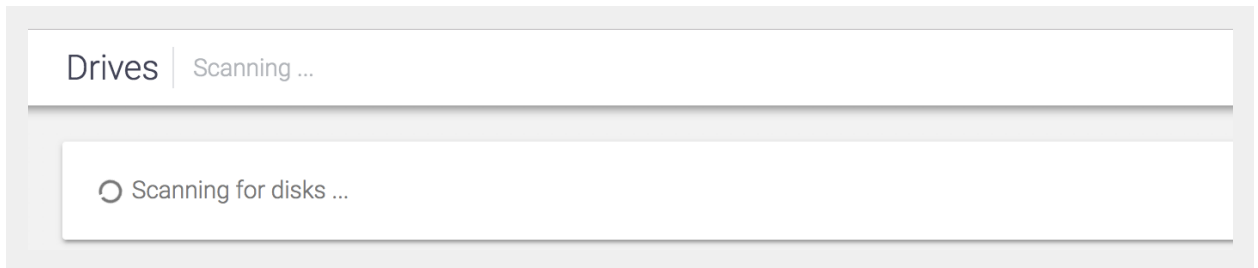
Size	5.5 TB	Vendor	HGST
Pool	hybrid	Model	HUS726060AL5210
Enclosure	Expansion Shelf 12 Bays / front	Serial	K1HRJYSD
Position	row 1 / column 2	Physical block size	4096

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.

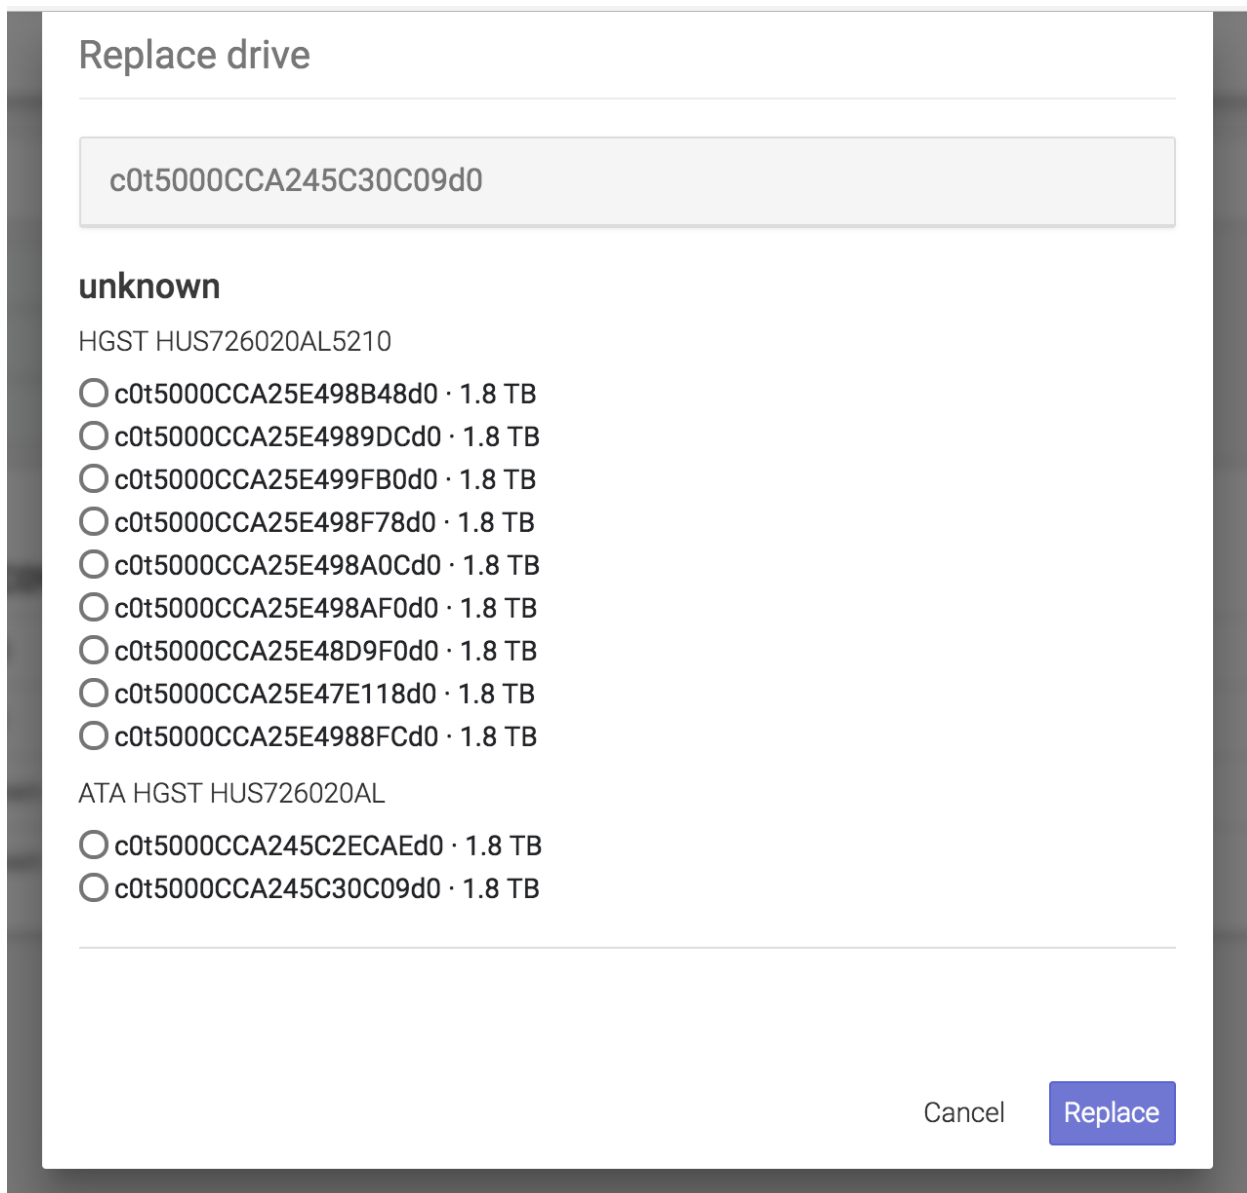
Place the bay in the machine.

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



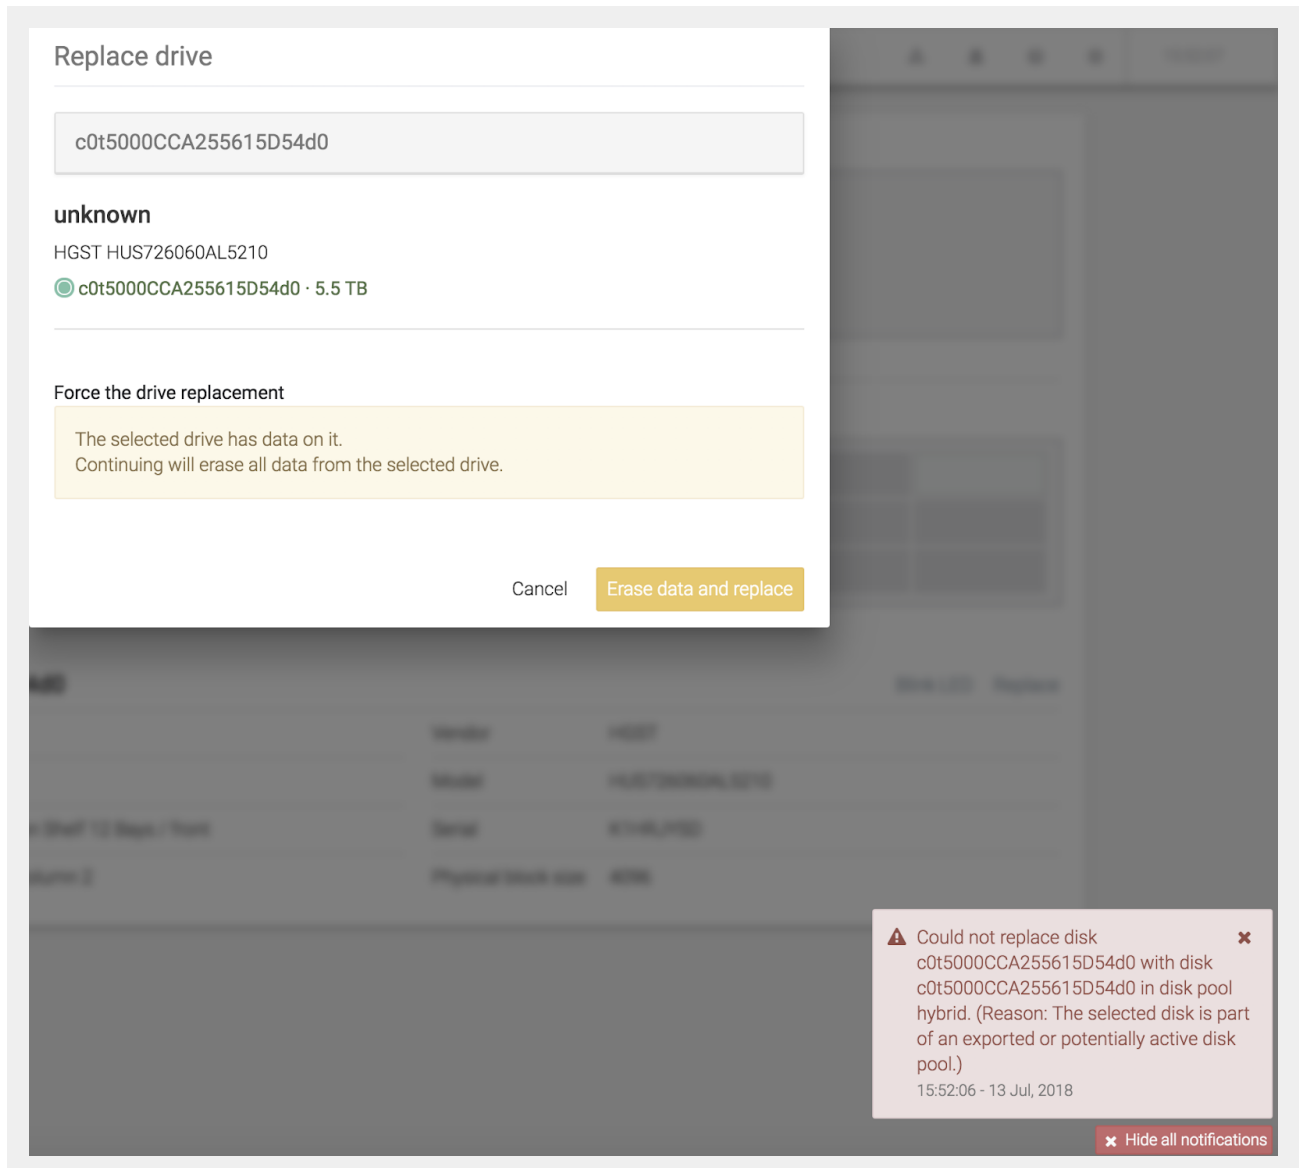
Click on Drive details → Replace .

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 shows the 'Pools' management page in SynetoOS. 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 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 shows 'Data' usage with two bars: 93.2 GB and 111.8 GB. A progress bar indicates '5.2 GB of 79.5 GB used'. A 'Healthy' status indicator with a green checkmark is present. 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

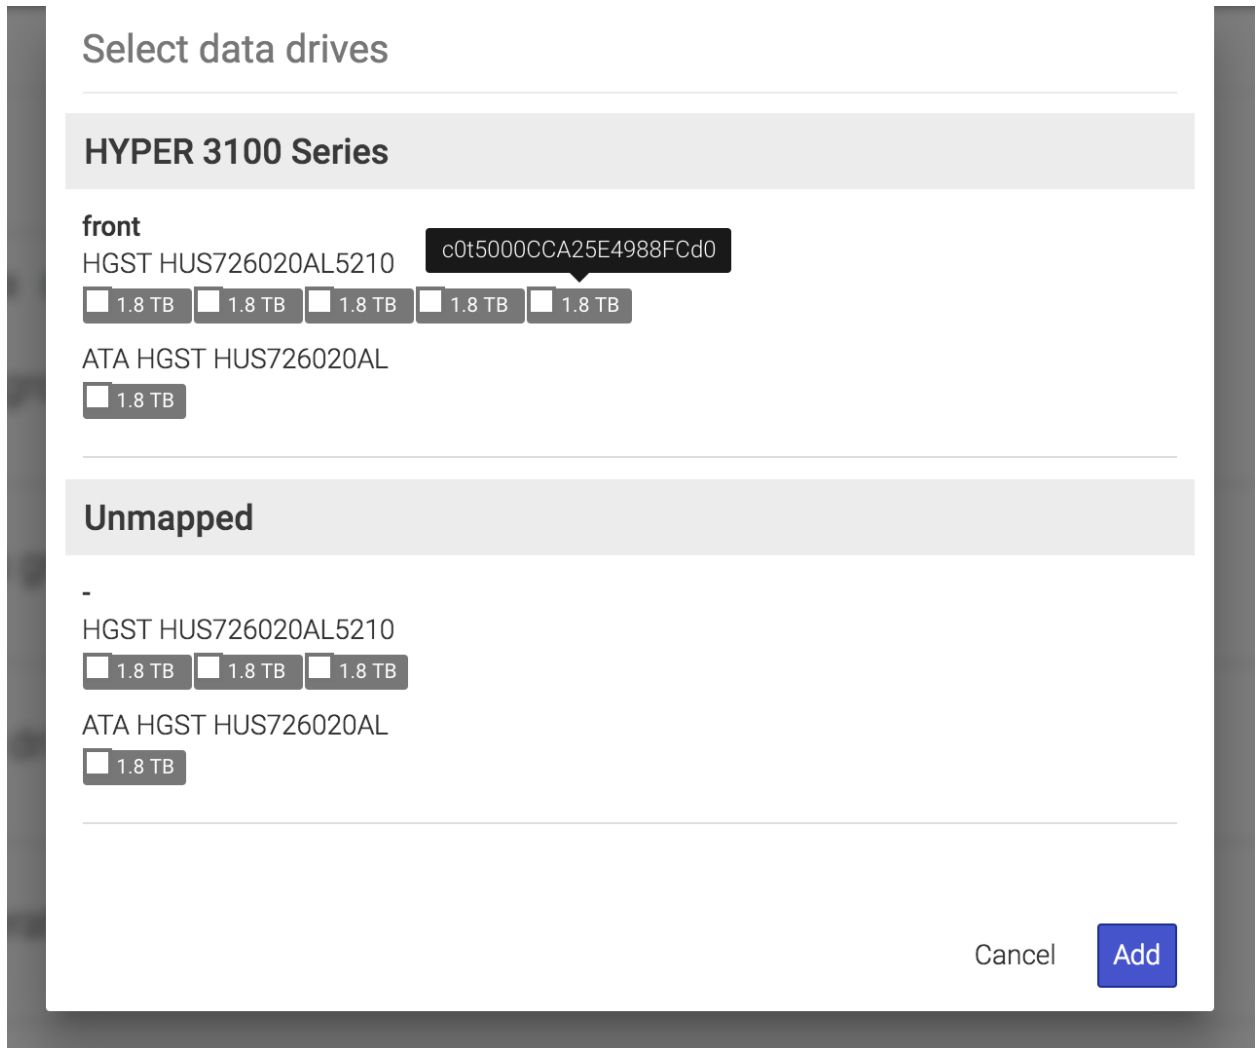
Write acceleration drives group

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

Write acceleration Manage

Read acceleration Manage

Spares Manage

324.0 KB of 1.8 TB used


Healthy

Integrity check Check now

Last check never checked

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

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

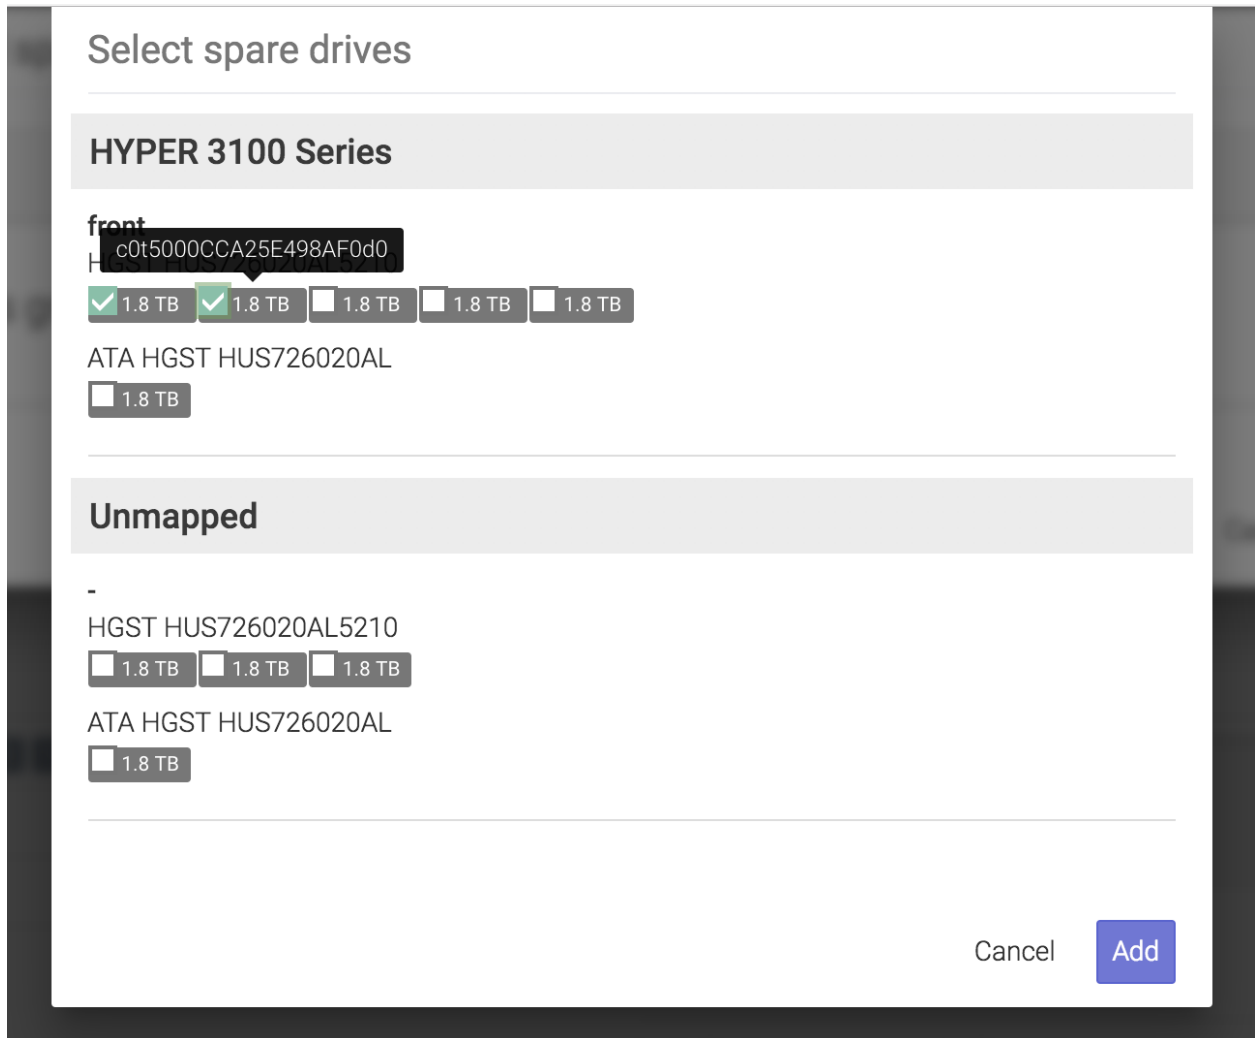
A dialog for adding spare drives will open.

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.

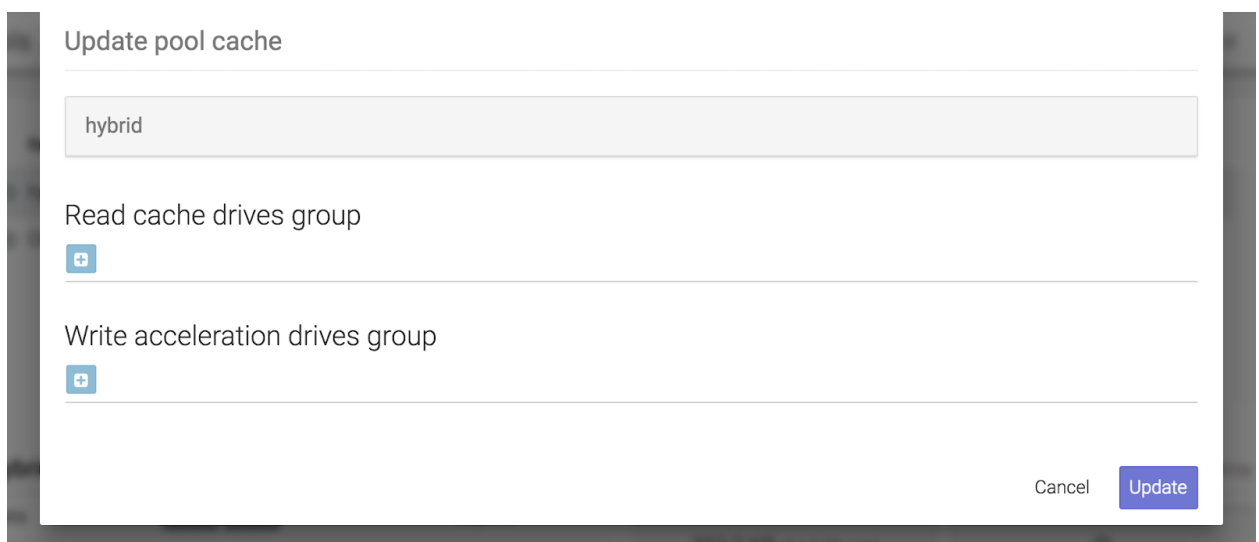



After selecting the drives you want, click Add then Update.

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.

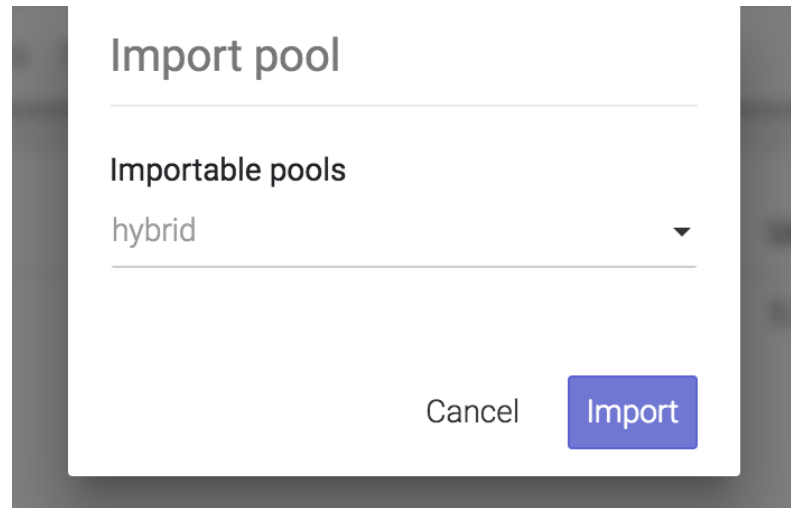
Importing a disk pool

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.

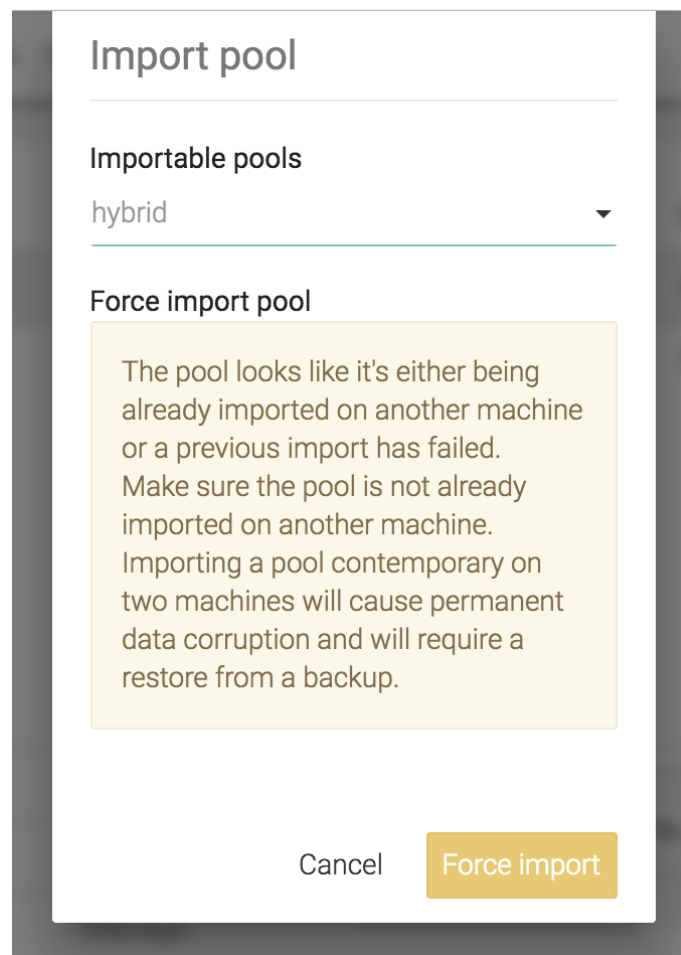
Click Import existing in the action bar. If there are drive pools that can be imported, a list will be shown.

Pools | [New pool](#) [Integrity check](#) [Rescan](#)



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.

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.

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.

Drives Rescan
Map List
⌵ ⌵ ⌵

HYPER 3100 Series

front

rear

Expansion Shelf 12 Bays

front

front

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.

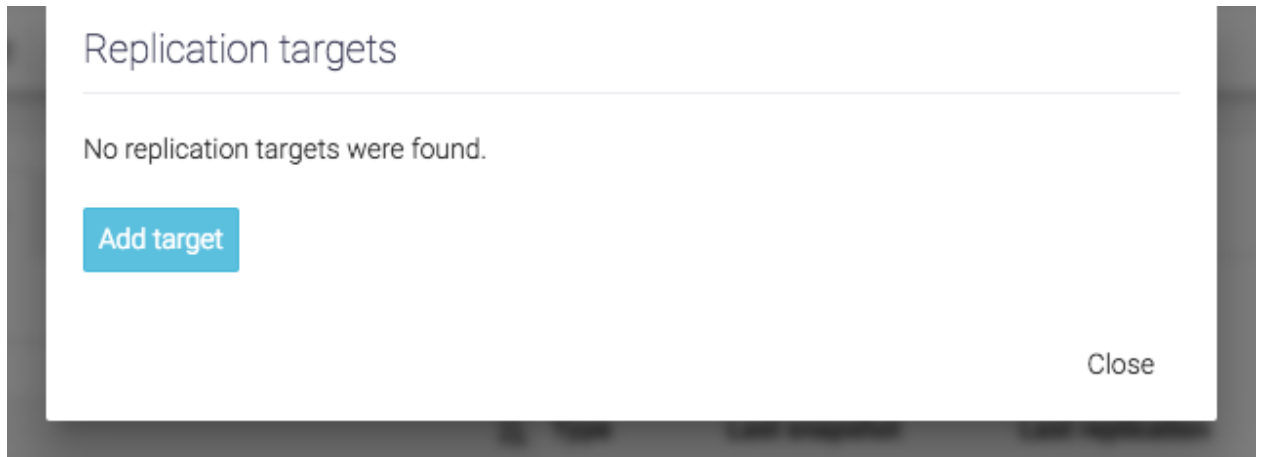
Adding a replication target

The screenshot shows the 'Protect and replicate' page with the 'Replication targets' dialog open. The dialog has tabs for 'Datastores', 'Shares', and 'Volumes'. It includes a search bar, a 'Show All' dropdown, and a '1-2 of 2' indicator. The table below lists 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, there are buttons for 'Progress', 'Snapshots', and 'Schedules'. A message at the bottom states: 'No snapshots have been taken yet.'

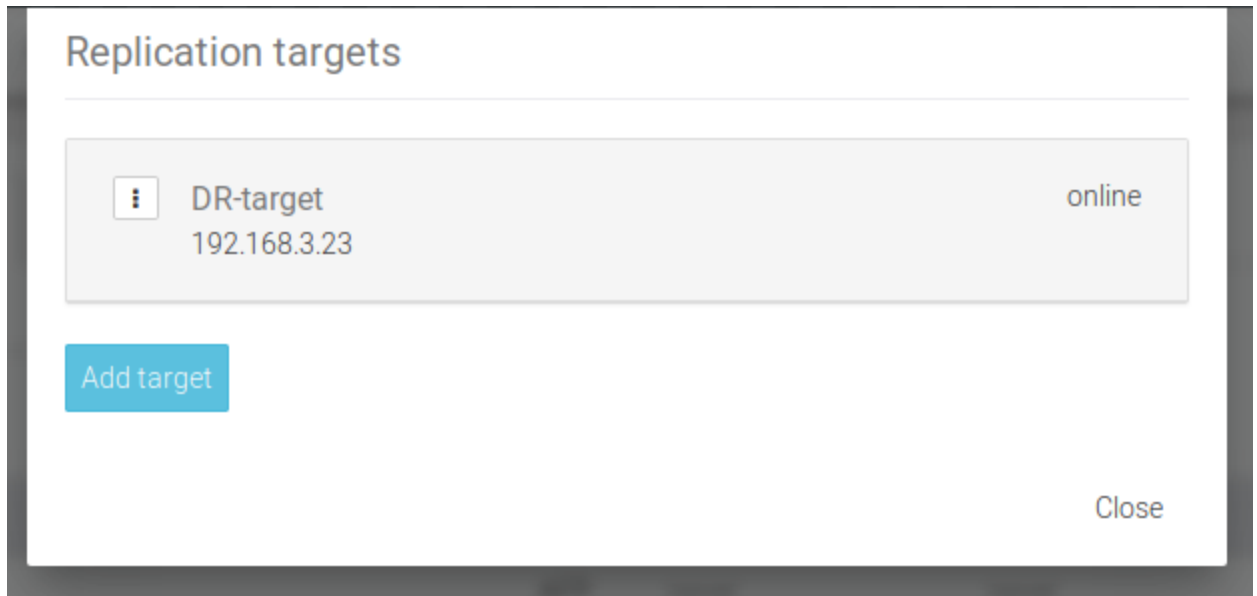
On the Protect and replicate page, click Replication targets in the action bar. A dialog will open, listing the available replication targets.



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.

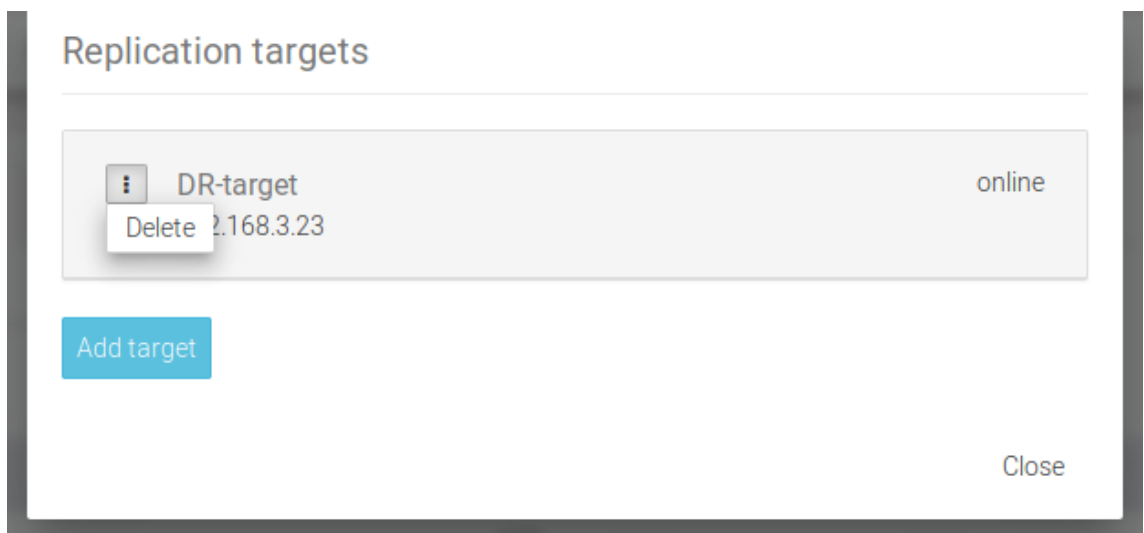


WARNING: It will fail when adding as replication target a HYPERseries machine with an OS below ver. 3.2.8.

Removing a replication target

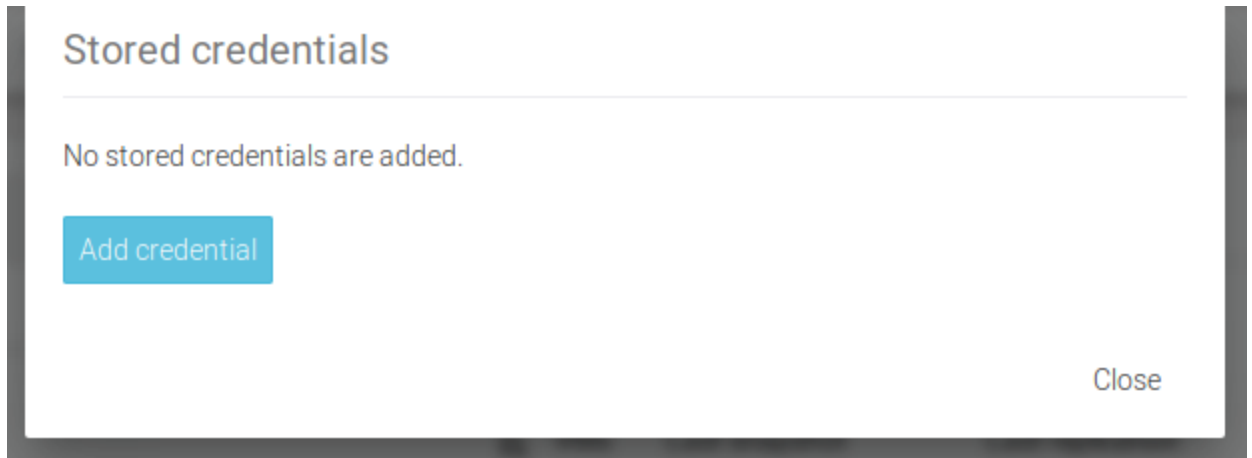
On the Protect and replicate page, click Replication targets in the action bar. A dialog will open listing the available replication targets. For the replication target you want to remove, click on Context Menu → Delete.

WARNING: Deleting a replication target, will also remove it from the snapshot schedules.

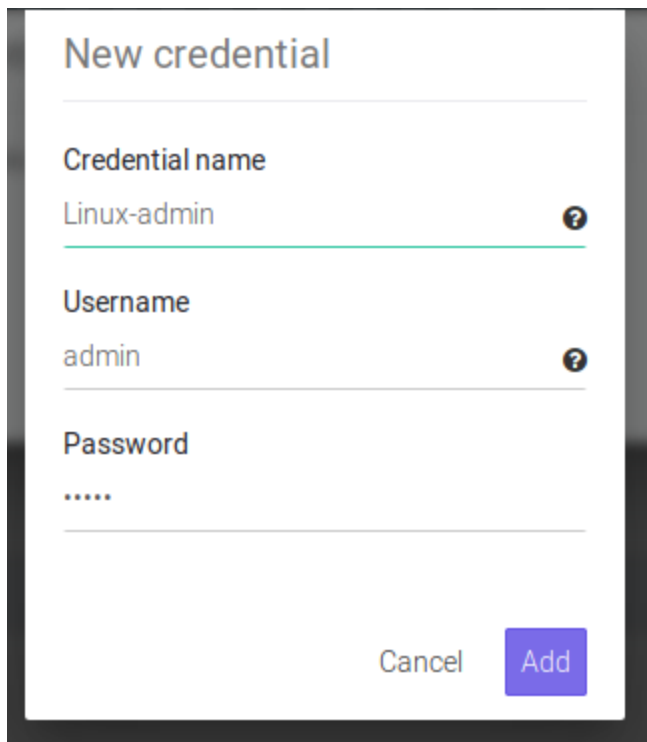


Configuring virtual machine credentials

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.



New credential

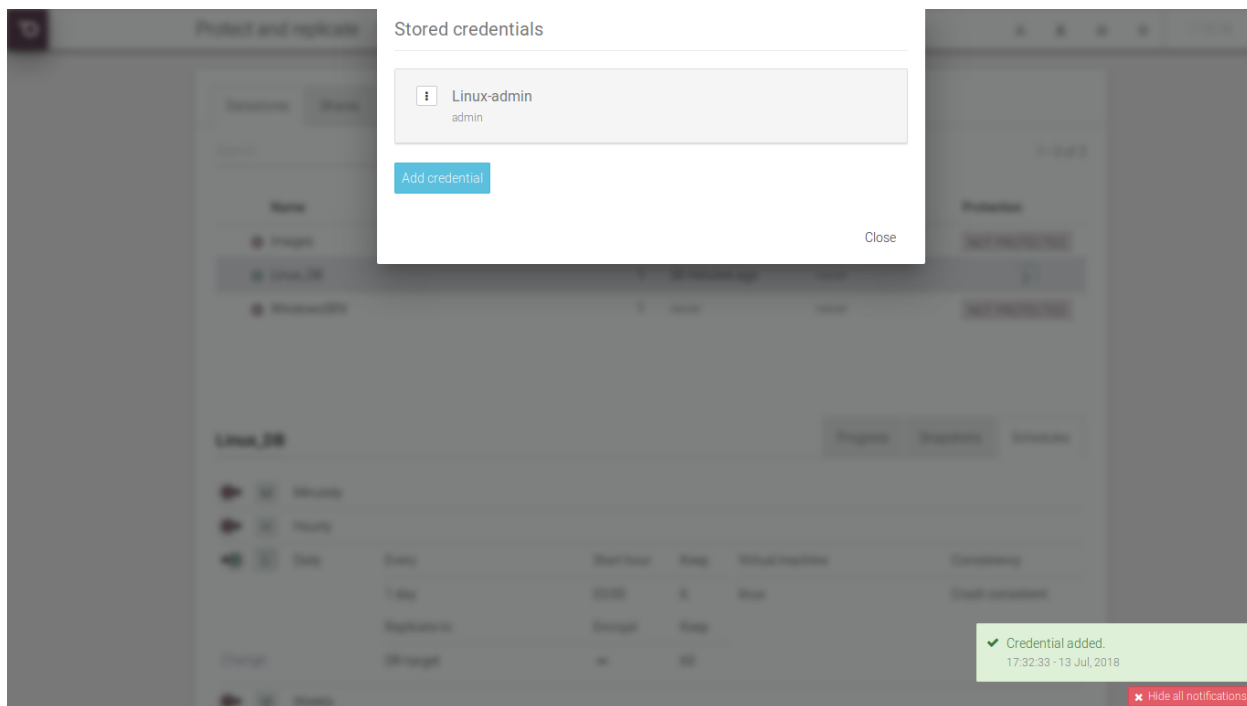
Credential name
Linux-admin ?

Username
admin ?

Password
.....

Cancel Add

Click Add. The credentials will appear in the list.



Stored credentials

Linux-admin
admin

Add credential

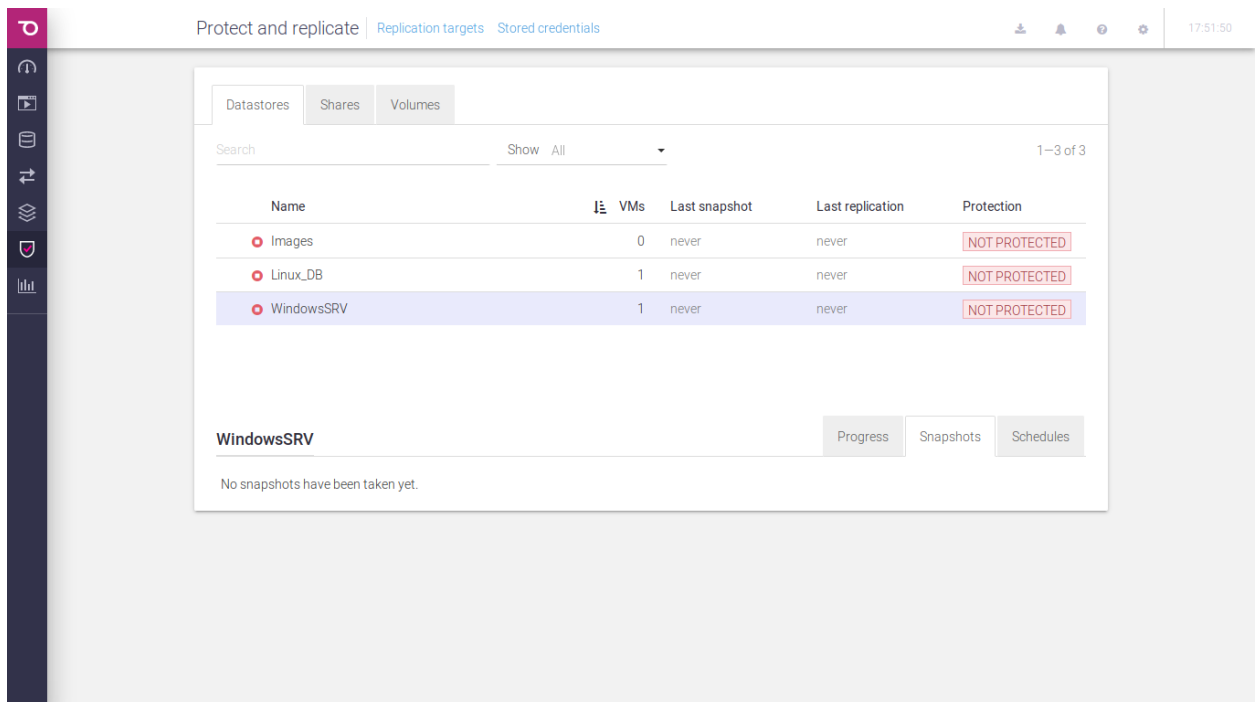
Close

✓ Credential added.
17:32:33 - 13 Jul, 2018
Hide all notifications

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.



The screenshot shows the 'Protect and replicate' page in SynetoOS. The page has a dark sidebar on the left with various icons. The main content area is titled 'Protect and replicate' and has three tabs: 'Datastores', 'Shares', and 'Volumes'. The 'Datastores' tab is selected. Below the tabs is a search bar and a 'Show All' dropdown menu. A table lists three datastores: 'Images', 'Linux_DB', and 'WindowsSRV'. The 'WindowsSRV' row is highlighted. Below the table, there are buttons for 'Progress', 'Snapshots', and 'Schedules'. A message at the bottom states 'No snapshots have been taken yet.'

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

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

Datastores | 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.

Protect and replicate | [Replication targets](#) | [Stored credentials](#) | 15:27:57

Datastores | 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

Linux_DB | 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 datastore and frequency.

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**

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

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**

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 targets: 'Images', 'Linux_DB', and 'WindowsSRV'. The 'Linux_DB' target is highlighted, and its 'Schedules' sub-tab is active. The 'Linux_DB' schedule is configured as 'Daily' with the following details:

Frequency	Every	Start hour	Keep	Virtual machine	Consistency
Daily	1 day	03:00	6	linux	Crash consistent
Replicate to	DR-target	Encrypt	Keep		

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

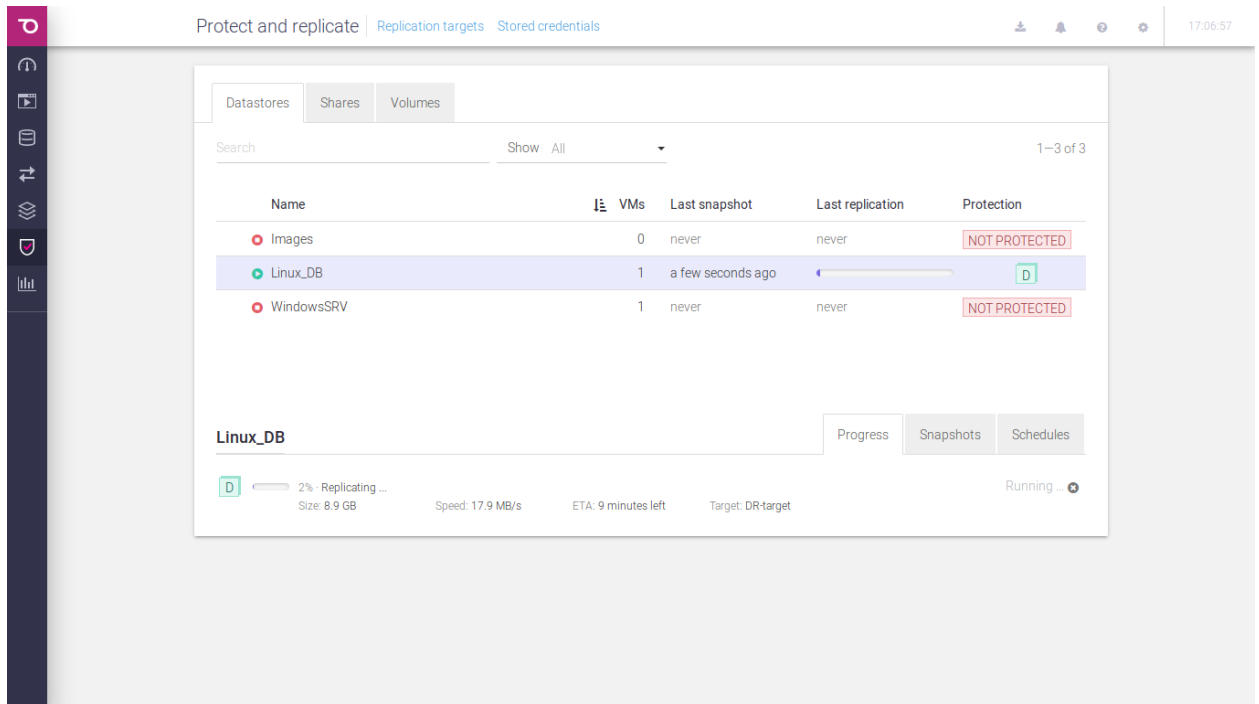
Name	VMs	Last snapshot	Last replication	Protection
Images	0	never	never	NOT PROTECTED
Linux_DB	1	a few seconds ago		D
WindowsSRV	1	never	never	NOT PROTECTED

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.




The screenshot shows 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 three replication targets:

Name	VMs	Last snapshot	Last replication	Protection
Images	0	never	never	NOT PROTECTED
Linux_DB	1	a few seconds ago		
WindowsSRV	1	never	never	NOT PROTECTED

Below the table, the 'Linux_DB' target is selected, and its details are shown. It includes a 'Progress' tab, 'Snapshots' tab, and 'Schedules' tab. The progress bar shows '2% - Replicating ...' with a size of '8.9 GB', speed of '17.9 MB/s', and 'ETA: 9 minutes left'. The target is 'DR-target' and is '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 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, there are tabs for 'Datastores', 'Shares', and 'Volumes'. A search bar and 'Show All' dropdown are present. The main table lists replication targets:

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

Below the table, there are tabs for 'Progress', 'Snapshots', and 'Schedules'. The 'Snapshots' tab is active, showing a calendar for July 2018. The date July 13, 2018, is selected. To the right of the calendar, a list of snapshots is shown:

Date	Time	Size
July 13, 2018	17:06	20.7 MB
July 13, 2018	16:00	20.9 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 selected for 'Linux_DB'. The calendar shows July 13, 2018, selected. The snapshot list is expanded to show the history for the 16:00 snapshot:

Date	Time	Size	Details
July 13, 2018	17:06	20.7 MB	
July 13, 2018	16:00	20.9 MB	<ul style="list-style-type: none"> 16:00:05 Starting automated schedule... 16:00:06 VMware Datastore name: Linux_DB ... 16:00:06 Frequency: every 1 hours ... 16:00:07 Host: syneto-os-710833dd ... 16:00:07 Virtual machines to snapshot: 1 16:00:08 Virtual machine 1/1: linux 16:00:08 Taking crash consistent snapshot ... 16:00:09 Taking VMware Datastore snapshot ... 16:00:10 Local snapshot taken.

Below the snapshot list, there is a 'Summary' section:

Oldest snapshot on	2018/07/13
Total snapshots	2 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	Start hour	Keep
1 days	03:00	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

Crash consistent

Application consistent

Windows VSS

Windows VSS and hooks

Hooks only

Live snapshot (RAM 1.0 GB)

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

The details for 'Linux_DB' are shown below:

Schedule	Frequency	Start hour	Keep	Virtual machine	Consistency
<input type="checkbox"/> M	Minutely				
<input type="checkbox"/> H	Hourly				
<input checked="" type="checkbox"/> D	Daily	Every	1 day	03:00	6
		Replicate to	Encrypt	Keep	
		DR-target	-	60	
<input type="checkbox"/> 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

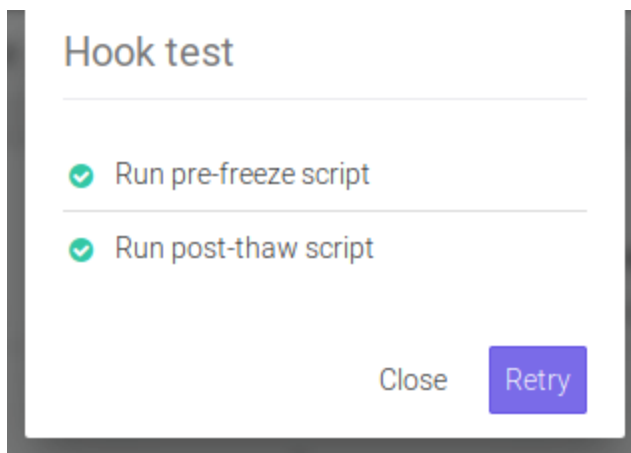
Extra space used for live snapshots: **0.0 B**

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.

Protecting a share

Note: Cloned shares cannot be protected because they are temporary objects.

The Protect and replicate page contains three tabs: Datastores, 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

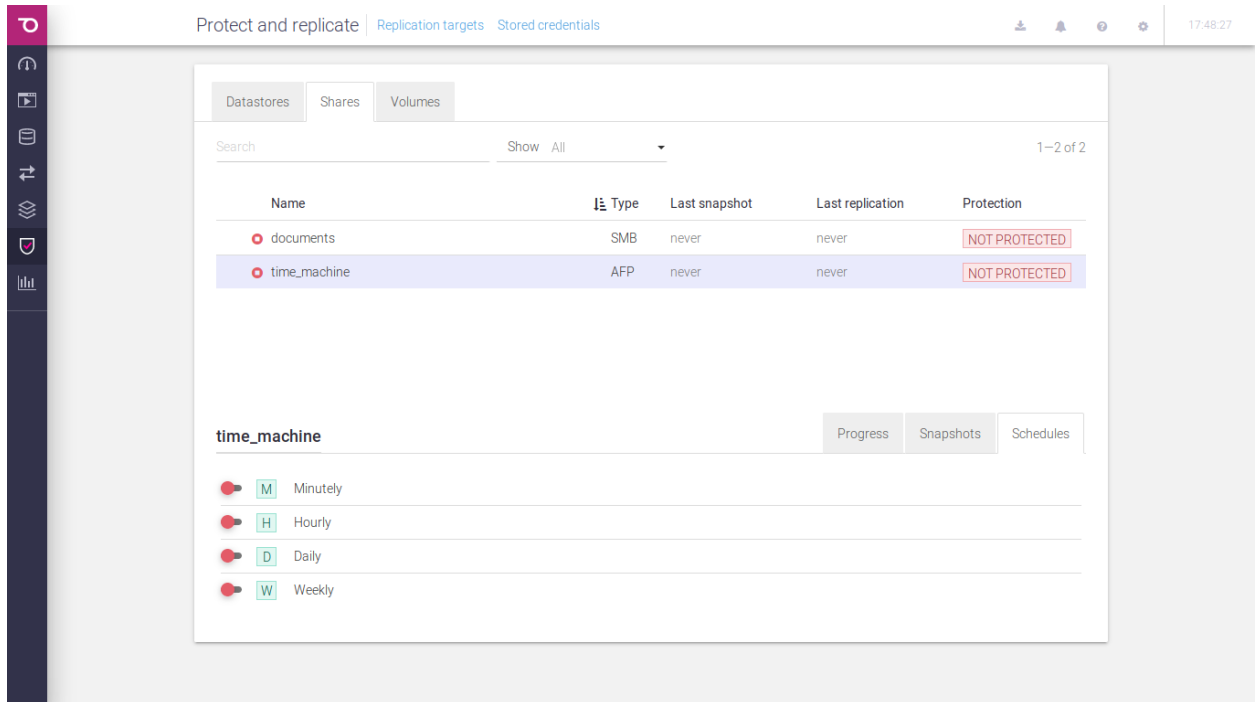
time_machine | Progress | Snapshots | Schedules

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 shows the SynetoOS interface for configuring a share. The main content area displays a table of shares and their replication status. Below the table, the 'time_machine' share is selected, and the 'Schedules' tab is active. The 'Schedules' section shows four frequency options: Minutely, Hourly, Daily, and Weekly, each with a toggle switch.

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.

The screenshot shows a configuration window titled "Hourly schedule". At the top, there is a text field containing "time_machine". Below this is a section titled "Protection". Under "Protection", there are three fields: "Every" with a value of "1" and a unit of "hours", "Start minute" with a value of "00", and "Keep" with a value of "6". Below these is a "Replicate to" section with a dropdown menu currently showing "Don't replicate". At the bottom right of the dialog are two buttons: "Cancel" and "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	Keep
1 ▼ hours	00	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 1 ▼ hours	Start minute 00	Keep 6
Replicate to DR-target ▼	Encrypt <input type="checkbox"/>	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.

Protect and replicate | Replication targets | Stored credentials

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	H

time_machine Progress | Snapshots | Schedules

M Minutely
 H Hourly
 D Daily
 W Weekly

Every	Start minute	Keep
1 hour	00	6
Replicate to	Encrypt	Keep
DR-target	<input type="checkbox"/>	48

Change

✓ Protection schedule saved for **time_machine**
12:05:15 - 16 Jul, 2018

✖ Hide all notifications

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	Keep
1 hours	00	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.

The screenshot shows the 'Protect and replicate' section of the SynetoOS interface. It displays a table of replication targets with the following data:

Name	Type	Last snapshot	Last replication	Protection
documents	SMB	never	never	NOT PROTECTED
time_machine	AFP	a minute ago	a few seconds ago	[H]

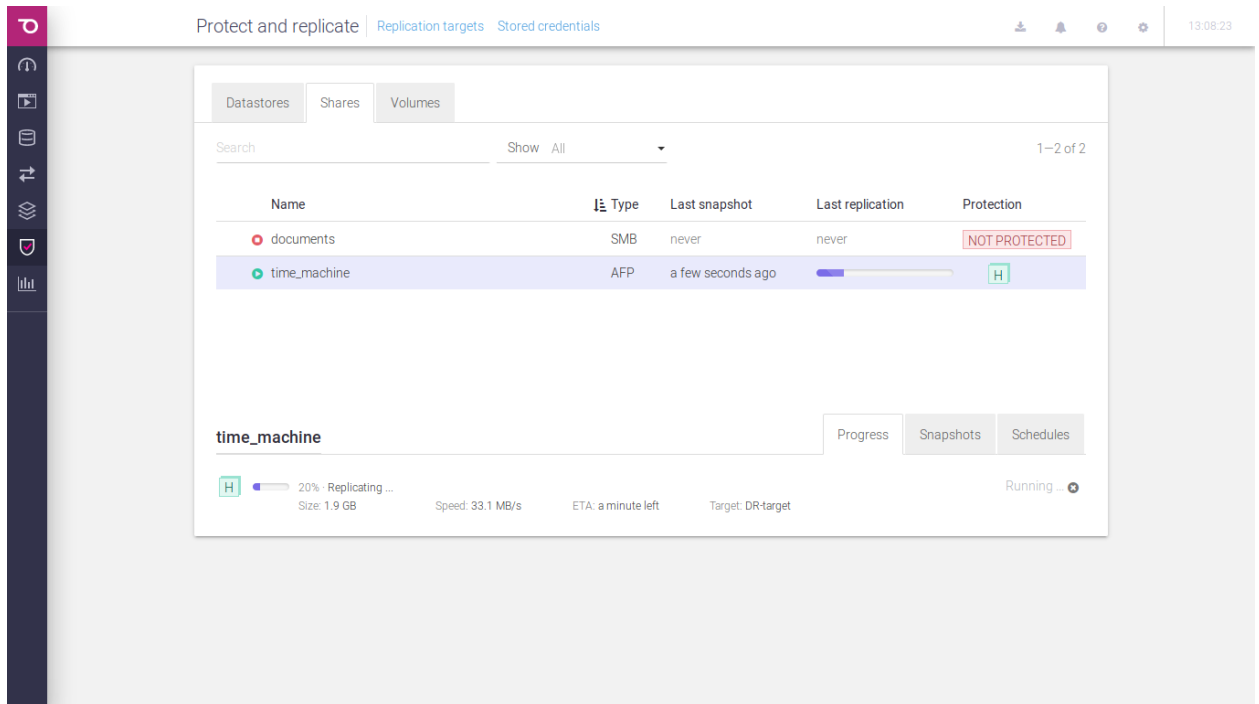
Below the table, the details for the 'time_machine' share are shown. The 'Progress' tab is active, displaying a progress bar that is 100% filled. The status message reads: 'Replication completed a few seconds ago'. Additional details include 'Size: 18.6 KB' and 'Target: DR-target'. A 'Run now' button is visible in the bottom right corner of the 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.




The screenshot displays the SynetoOS interface for managing replication targets. The main window is titled "Protect and replicate" and contains tabs for "Replication targets" and "Stored credentials". The "Replication targets" tab is active, showing a table of targets. The table has columns for "Name", "Type", "Last snapshot", "Last replication", and "Protection". Two targets are listed: "documents" (SMB, never, never, NOT PROTECTED) and "time_machine" (AFP, a few seconds ago, progress bar, H). The "time_machine" target is selected, and its details are shown below the table. The details include a progress bar at 20% (Replicating ...), size of 1.9 GB, speed of 33.1 MB/s, ETA of a minute left, and target of DR-target. The status is "Running ...".

Name	Type	Last snapshot	Last replication	Protection
documents	SMB	never	never	NOT PROTECTED
time_machine	AFP	a few seconds ago		H

time_machine Progress Snapshots Schedules

 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.

Protect and replicate | [Replication targets](#) | [Stored credentials](#) | 16:32:28

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	3 hours ago	4 hours ago	H

time_machine | [Progress](#) | [Snapshots](#) | [Schedules](#)

< July 2018 >

Mo	Tu	We	Th	Fr	Sa	Su	July 16, 2018
25	26	27	28	29	30	1	▶ 13:08 24.0 KB
2	3	4	5	6	7	8	▶ 13:00 24.0 KB
9	10	11	12	13	14	15	▶ 12:19 24.0 KB
16	17	18	19	20	21	22	
23	24	25	26	27	28	29	
30	31	1	2	3	4	5	

3 snapshots

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](#) | 16:33:10

time_machine | AFP | 3 hours ago | 4 hours ago | H

time_machine | [Progress](#) | [Snapshots](#) | [Schedules](#)

< July 2018 >

Mo	Tu	We	Th	Fr	Sa	Su	July 16, 2018
25	26	27	28	29	30	1	▶ 13:08 24.0 KB
2	3	4	5	6	7	8	▼ 13:00 24.0 KB
9	10	11	12	13	14	15	13:00:11 Starting automated schedule...
16	17	18	19	20	21	22	13:00:12 AFP Share name: time_machine ...
23	24	25	26	27	28	29	13:00:13 Frequency: every 1 hours ...
30	31	1	2	3	4	5	13:00:13 Host: syneto-os-710833dd ...
							13:00:15 Taking AFP Share snapshot ...
							13:00:16 Local snapshot taken.
							13:00:25 Starting replication to 192.168.3.23 ...
							13:00:25 Finished replication.
							▶ 12:19 24.0 KB

Summary

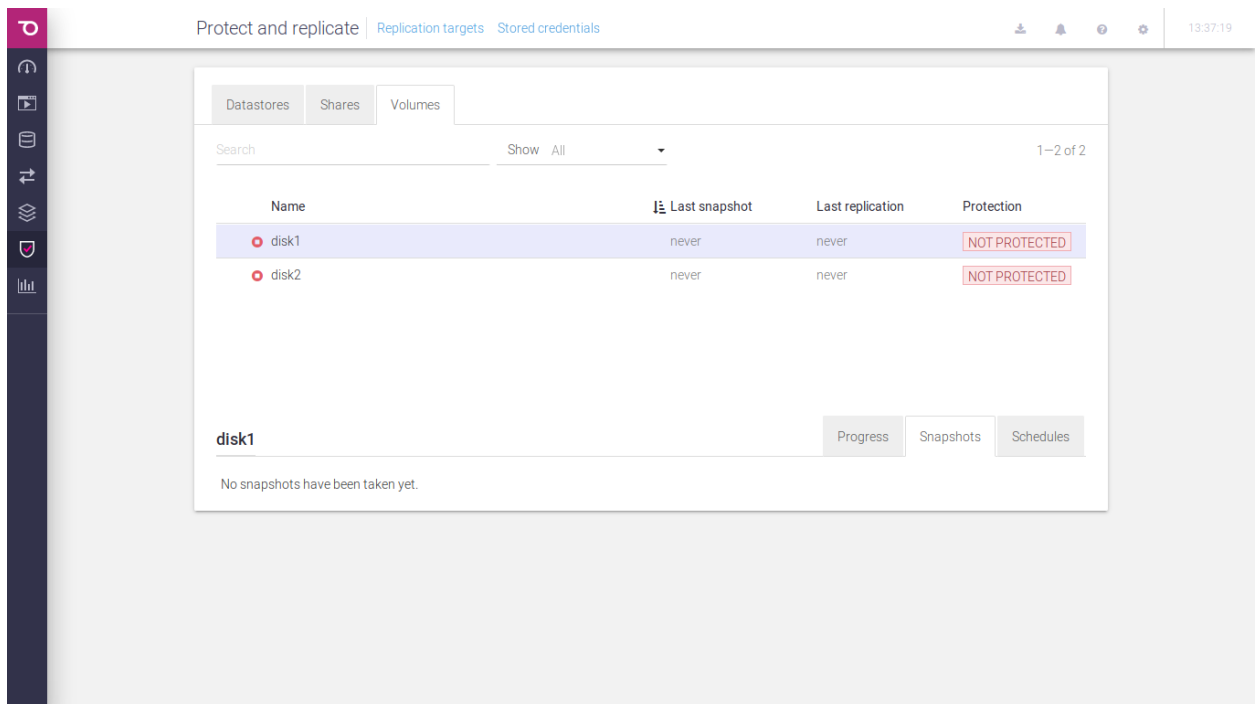
Oldest snapshot on 2018/07/16

Total snapshots 3 [Purge](#)

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' page in SynetoOS. The 'Volumes' tab is selected, 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 buttons for 'Progress', 'Snapshots', and 'Schedules'. A message states: 'No snapshots have been taken yet.'

The table lists all the volumes. There is a Search bar, useful for searching for a specific volume by specifying its name.

Protect and replicate | [Replication targets](#) | [Stored credentials](#) 13:39:50

Datstores | Shares | Volumes

disk1 Show All 1—1 of 1 (filtered from 2 total entries)

Name	Last snapshot	Last replication	Protection
disk1	never	never	NOT PROTECTED

disk1 Progress Snapshots Schedules

No snapshots have been taken yet.

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.

Protect and replicate | [Replication targets](#) | [Stored credentials](#) 13:49:41

Datstores | Shares | **Volumes**

Search Show All 1-2 of 2

Name	Last snapshot	Last replication	Protection
<input type="radio"/> disk1	never	never	NOT PROTECTED
<input type="radio"/> disk2	never	never	NOT PROTECTED

disk1 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 volume and frequency.

The screenshot shows a dialog box titled "Hourly schedule". At the top, a light gray box contains the text "disk1". Below this, the "Protection" section is visible. It includes three input fields: "Every" with a dropdown menu showing "1" and "hours", "Start minute" with the value "00", and "Keep" with the value "6". Below these is a "Replicate to" section with a dropdown menu currently showing "Don't replicate". At the bottom right of the dialog, there are two buttons: "Cancel" and "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	Keep
1 <input type="text"/> hours	00 <input type="text"/>	6 <input type="text"/>

Replicate to

Don't replicate

- Don't replicate
- DR-target

Cancel

After you select the desired replication target, two more fields will appear.

Hourly schedule

disk1

Protection

Every 1 hours Start minute 00 Keep 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.

Protect and replicate | Replication targets | Stored credentials

Search Show All 1-2 of 2

Name	Last snapshot	Last replication	Protection
disk1	never	never	H
disk2	never	never	NOT PROTECTED

disk1 Progress | Snapshots | Schedules

M Minutely
 H Hourly
 D Daily
 W Weekly

Every	Start minute	Keep
1 hour	00	6
Replicate to	Encrypt	Keep
DR-target	<input type="checkbox"/>	48

Change

✓ Protection schedule saved for **disk1**.
15:43:53 - 16 Jul, 2018

✗ Hide all notifications

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 hours

Start minute

Keep

Replicate to

Encrypt

Keep

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 shows 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 two volumes:

Name	Last snapshot	Last replication	Protection
disk1	a few seconds ago	a few seconds ago	[H]
disk2	never	never	NOT PROTECTED

Below the table, the details for 'disk1' are shown. It includes a 'Progress' tab, 'Snapshots', and 'Schedules' tabs. A progress bar indicates 'Replication completed a few seconds ago' with a size of 6.6 KB and a target of DR-target. A 'Run now' button is visible in the bottom right corner of the 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 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 SynetoOS interface for managing replication targets. The main window is titled "Protect and replicate" and includes tabs for "Replication targets" and "Stored credentials". The interface shows a table of replication targets with columns for Name, Last snapshot, Last replication, and Protection. Below the table, there are tabs for "Progress", "Snapshots", and "Schedules".

Name	Last snapshot	Last replication	Protection
disk1	a few seconds ago		
disk2	never	never	NOT PROTECTED

Below the table, the "disk1" details are shown. It includes a "Progress" tab with a progress bar indicating 67% replication. Other details include "Size: 6.6 KB", "Speed: 6.5 MB/s", "ETA: a few seconds left", and "Target: DR-target".

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.

The screenshot shows 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 two replication targets:

Name	Last snapshot	Last replication	Protection
disk1	an hour ago	an hour ago	H D
disk2	never	never	NOT PROTECTED

Below the table, there are tabs for 'Progress', 'Snapshots', and 'Schedules'. The 'Snapshots' tab is active, showing a calendar for July 2018. The date July 16, 2018, is selected, and a list of snapshots is shown for that date:

Date	Time	Size
July 16, 2018	15:52	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.

This screenshot shows the detailed view of the snapshot selected in the previous image. The 'Snapshots' tab is active, and the selected snapshot is expanded to show its history:

Date	Time	Event
July 16, 2018	15:52:00	Starting automated schedule ...
July 16, 2018	15:52:00	ISCSI Volume name: disk1 ...
July 16, 2018	15:52:01	Frequency: every 1 hours ...
July 16, 2018	15:52:01	Host: syneto-os-710833dd ...
July 16, 2018	15:52:03	Taking ISCSI Volume snapshot ...
July 16, 2018	15:52:03	Local snapshot taken.
July 16, 2018	15:52:11	Starting replication to 192.168.3.23 ...
July 16, 2018	15:52:12	Finished replication.

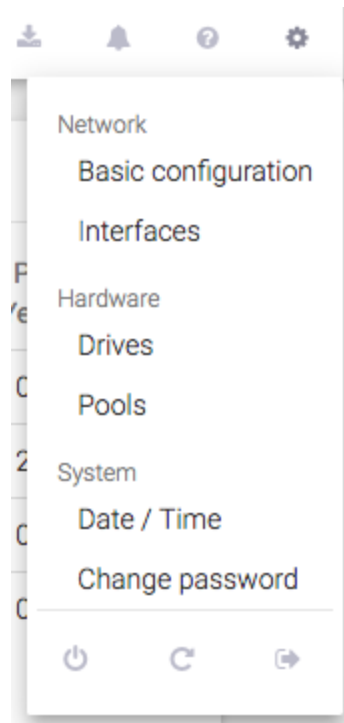
Below the history, there is a 'Summary' section:

Oldest snapshot on	2018/07/16
Total snapshots	1 Purge

SYSTEM CONFIGURATION

Networking

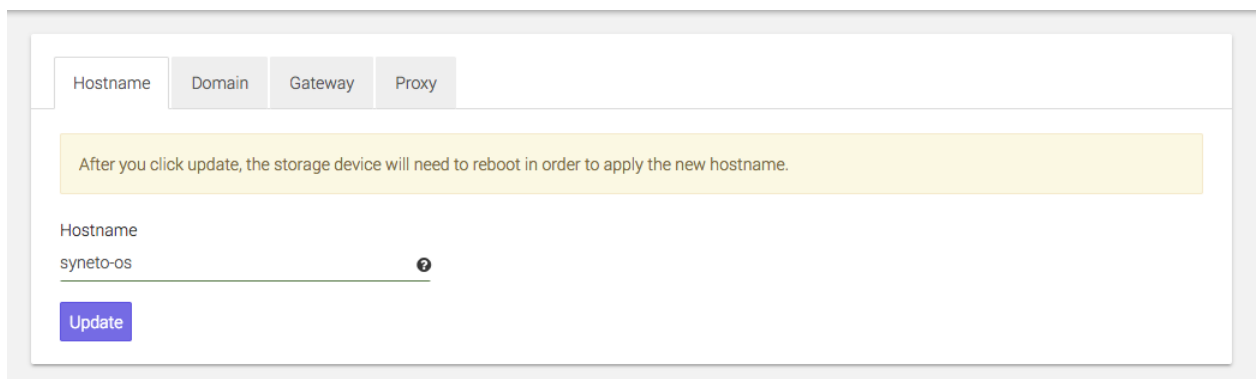
Select Secondary menu → Basic configuration, in the Network section.



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 SynetoOS Basic network configuration page. The page has a tabbed interface with 'Hostname', 'Domain', 'Gateway', and 'Proxy' tabs. The 'Hostname' tab is selected. A yellow warning box at the top states: 'After you click update, the storage device will need to reboot in order to apply the new hostname.' Below this, the 'Hostname' field is labeled 'Hostname' and contains the text 'syneto-os'. There is a question mark icon to the right of the field. 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 the 'Domain' configuration tab in the SynetoOS interface. At the top, there are four tabs: 'Hostname', 'Domain', 'Gateway', and 'Proxy'. The 'Domain' tab is selected. Below the tabs, the 'Domain name' field contains 'company.domain'. Underneath, there are three fields for DNS servers: 'DNS 1' with the value '192.168.1.15', 'DNS 2' with a masked value '#####', and 'DNS 3' with a masked value '#####'. A blue 'Update' button is located at the bottom left of the configuration area.

- Gateway - a router used by the storage to connect to the internet

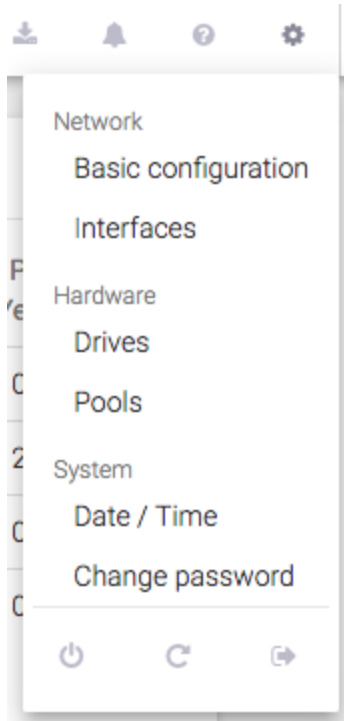
The screenshot shows the 'Gateway' configuration tab in the SynetoOS interface. At the top, there are four tabs: 'Hostname', 'Domain', 'Gateway', and 'Proxy'. The 'Gateway' tab is selected. Below the tabs, the 'Default gateway' field contains '192.168.1.1'. A blue 'Update' button is located at the bottom left of the configuration area.

- Proxy - the host and port of the HTTP proxy server


The screenshot shows the 'Proxy' configuration tab in the SynetoOS interface. At the top, there are four tabs: 'Hostname', 'Domain', 'Gateway', and 'Proxy'. The 'Proxy' tab is selected. Below the tabs, there is a checkbox labeled 'Use a HTTP proxy?' which is currently unchecked. A blue 'Update' button is located at the bottom left of the configuration area.

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.

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

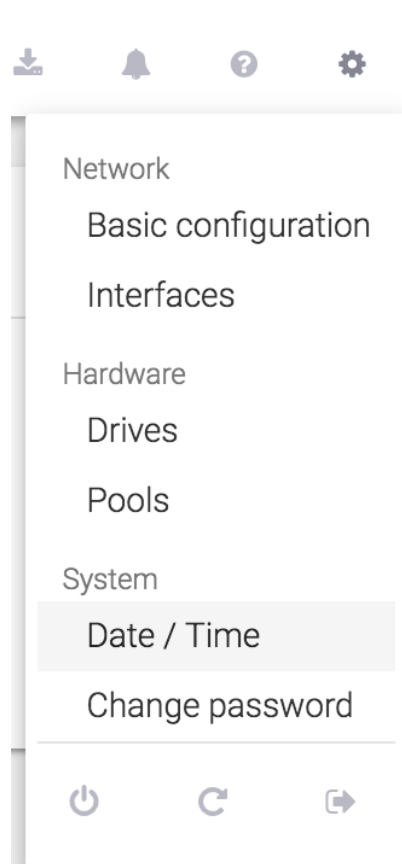
In case one or more of the ethernet interfaces are faulted, the system will display the message: missing: removed or broken!

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

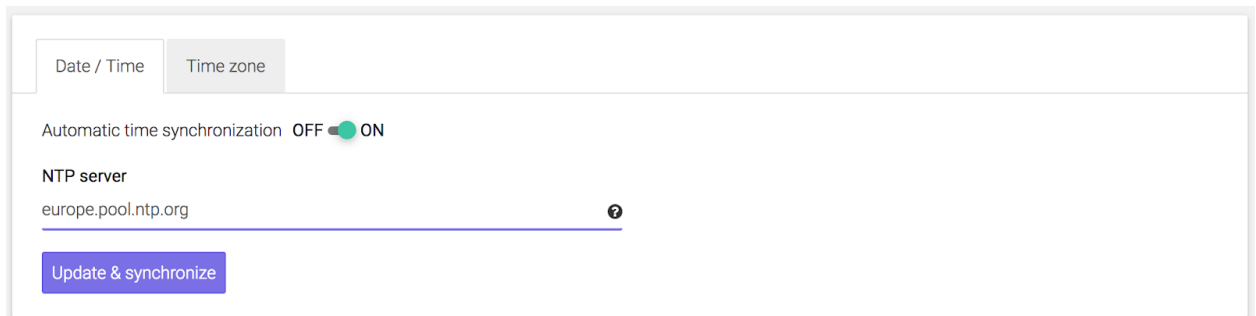
Select Secondary menu → Date / Time, in the System section.



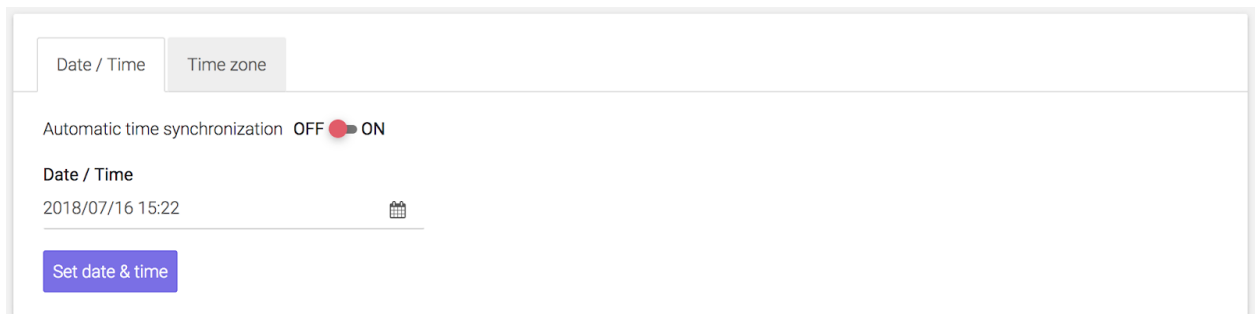
The time page allows you to configure the following:

- **Date / Time:** The date and time can be automatically synchronized with a NTP server or set manually.

Recommended: Enable automatic time synchronization.

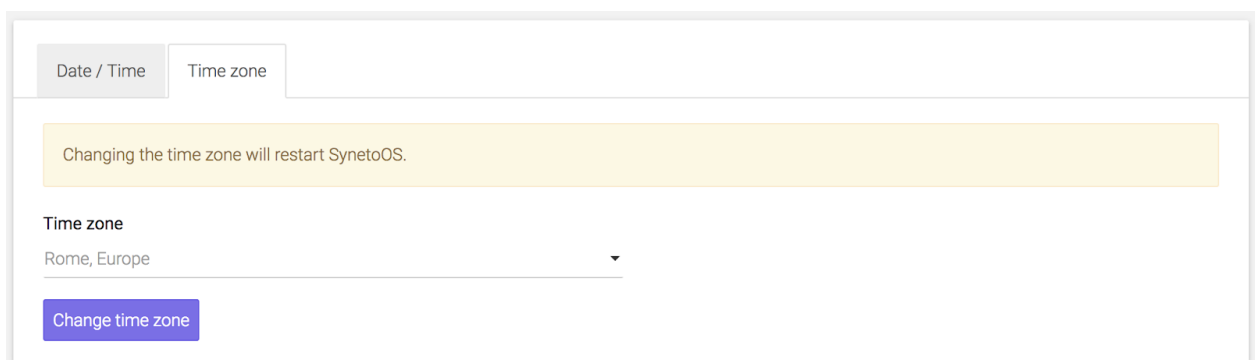


The screenshot shows the 'Date / Time' configuration page. At the top, there are two tabs: 'Date / Time' (selected) and 'Time zone'. Below the tabs, there is a toggle switch for 'Automatic time synchronization' which is currently turned 'ON'. Underneath, the 'NTP server' field contains the text 'europe.pool.ntp.org'. At the bottom of the page, there is a blue button labeled 'Update & synchronize'.



The screenshot shows the 'Date / Time' configuration page. At the top, there are two tabs: 'Date / Time' (selected) and 'Time zone'. Below the tabs, there is a toggle switch for 'Automatic time synchronization' which is currently turned 'OFF'. Underneath, the 'Date / Time' field contains the text '2018/07/16 15:22'. At the bottom of the page, there is a blue button labeled 'Set date & time'.

- **Timezone**



The screenshot shows the 'Timezone' configuration page. At the top, there are two tabs: 'Date / Time' and 'Time zone' (selected). Below the tabs, there is a yellow warning box that says 'Changing the time zone will restart SynetoOS.'. Underneath, the 'Time zone' field contains the text 'Rome, Europe'. At the bottom of the page, there is a blue button labeled 'Change time zone'.

WARNING: This will require a reboot.

Password change

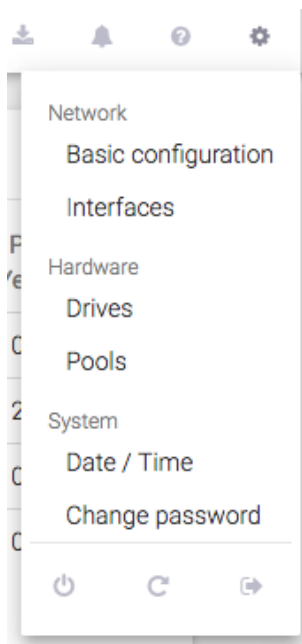
There are two sets of users and passwords that you need to manage regarding Syneto HYPER: local ESXi password, and SynetoOS management interface password.

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.



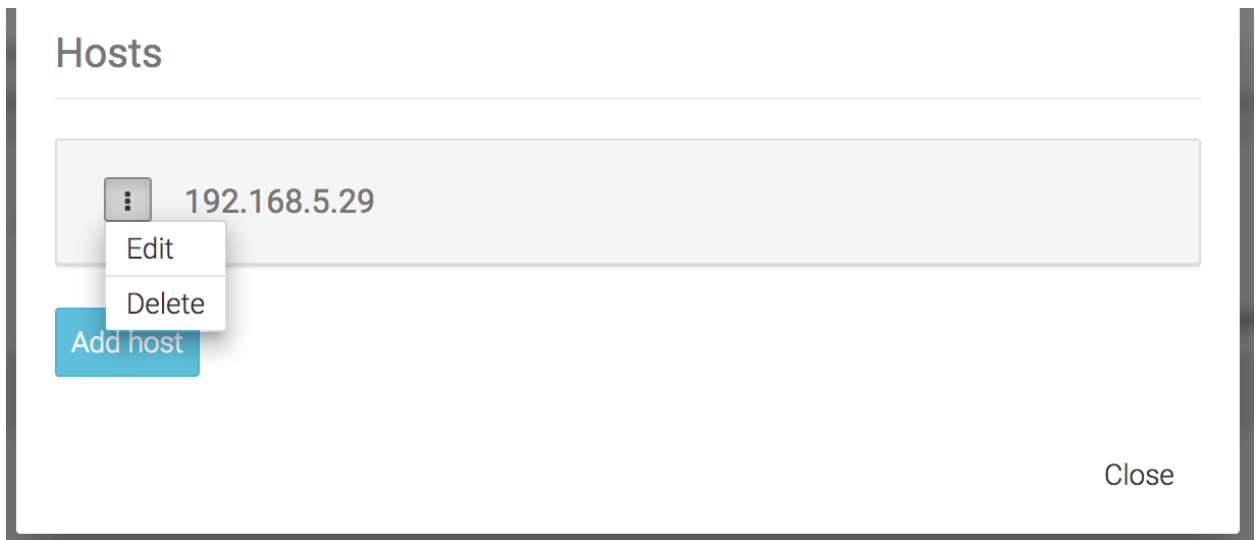
Change passwords

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"/>

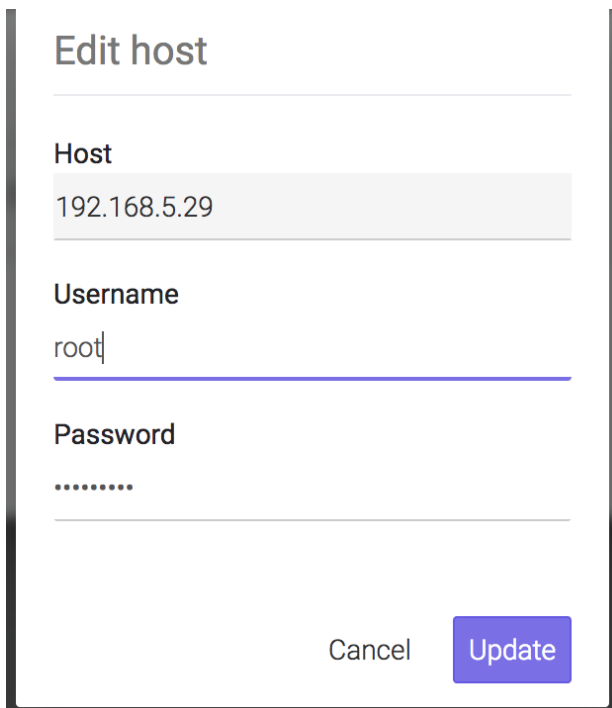
Close

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 shows the Alerts page interface. At the top, there's a navigation bar with 'Alerts', 'Email recipients', and 'Email sender service'. On the right, there are icons for a bell (with a red notification badge), a search icon, and a settings icon, along with the time '14:19:22'. The main content area is titled 'System errors' and includes 'Expand all / Collapse all' links. Two alerts are listed:

- 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 alerts are five 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, Memory health, Drives health, Storage health, Virtual machines health and Sensors health.

The screenshot displays a dashboard with the following components and their status:

- System errors:** Everything works fine.
- CPU health:** Everything works fine. Includes a "Get help" link.
- Memory health:** Everything works fine. Includes a "Get help" link.
- Drives health:** Everything works fine. Includes a "See more" link.
- Storage health:** Everything works fine. Includes a "See more" link.
- Virtual machines health:** Everything works fine. Includes a "See more" link.
- Sensors health:** Everything works fine. Includes 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.

The screenshot displays the 'System errors' section of the SynetoOS interface. At the top, there is a header 'System errors' and a link 'Expand all / Collapse all'. Below this, a yellow alert card is visible with the message 'A system service failed (51447a2e-f22f-6d8a-a90d-bbd9b2764f32)' and a timestamp '2018/07/16, 15:49:34'. A 'Mark as repaired' button is positioned over the alert. Below the alert, there are several health status cards: 'CPU health', 'Memory health', 'Drives health', 'Storage health', 'Virtual machines health', and 'Sensors health'. Each card has a 'Get help' or 'See more' link and displays the status 'Everything works fine' in green text.

Confirm

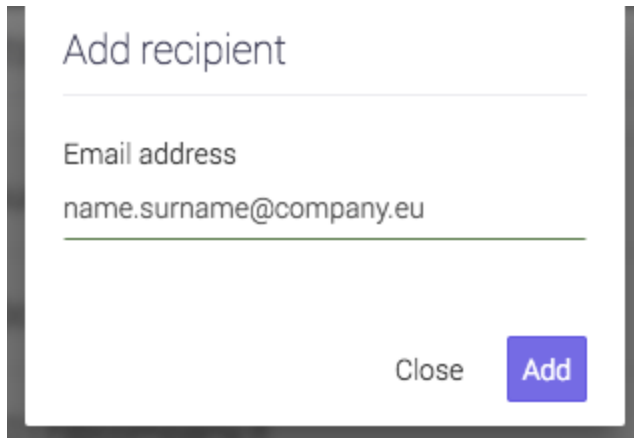
A fault should be marked as repaired only after the repaired procedure was completed.

Cancel

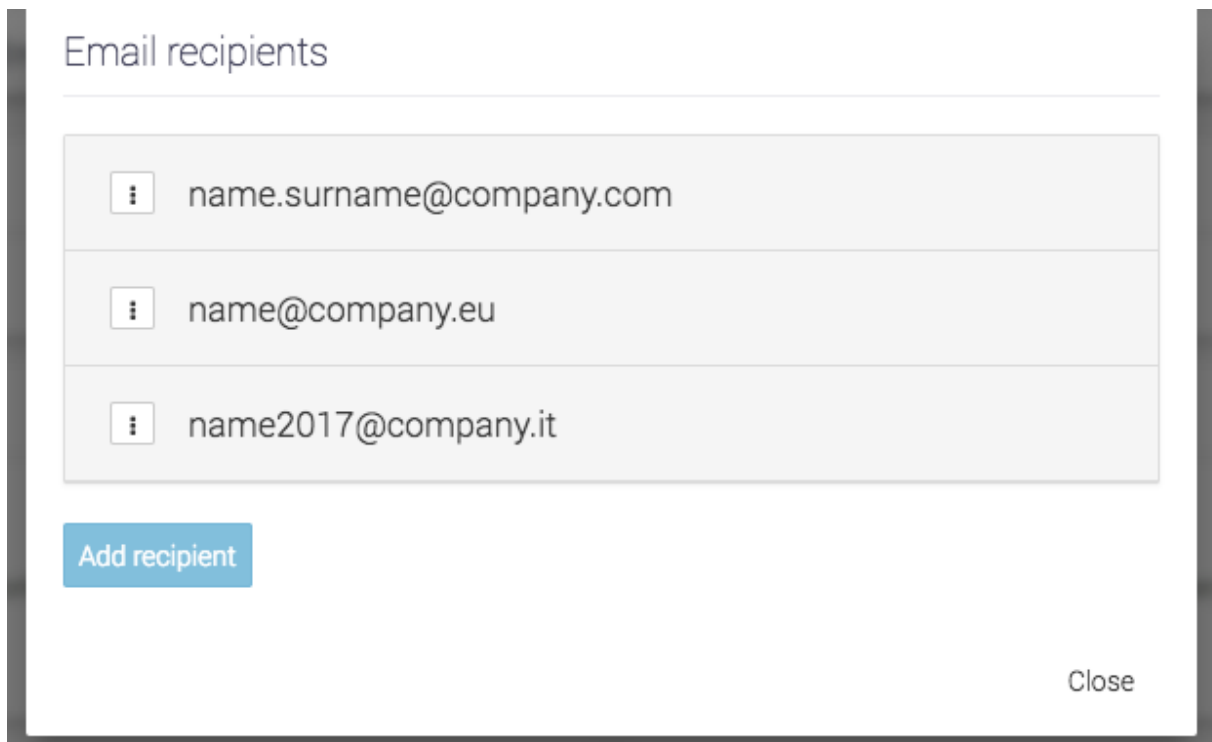
Mark as repaired

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.



The screenshot shows a modal window titled "Add recipient". It contains a text input field with the label "Email address" and the placeholder text "name.surname@company.eu". At the bottom right of the modal, there are two buttons: "Close" and "Add".



The screenshot shows a modal window titled "Email recipients". It displays a list of three email addresses, each with a small square icon containing three vertical dots to its left. The email addresses are "name.surname@company.com", "name@company.eu", and "name2017@company.it". At the bottom left of the modal, there is a blue button labeled "Add recipient". At the bottom right, there is a "Close" button.

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.

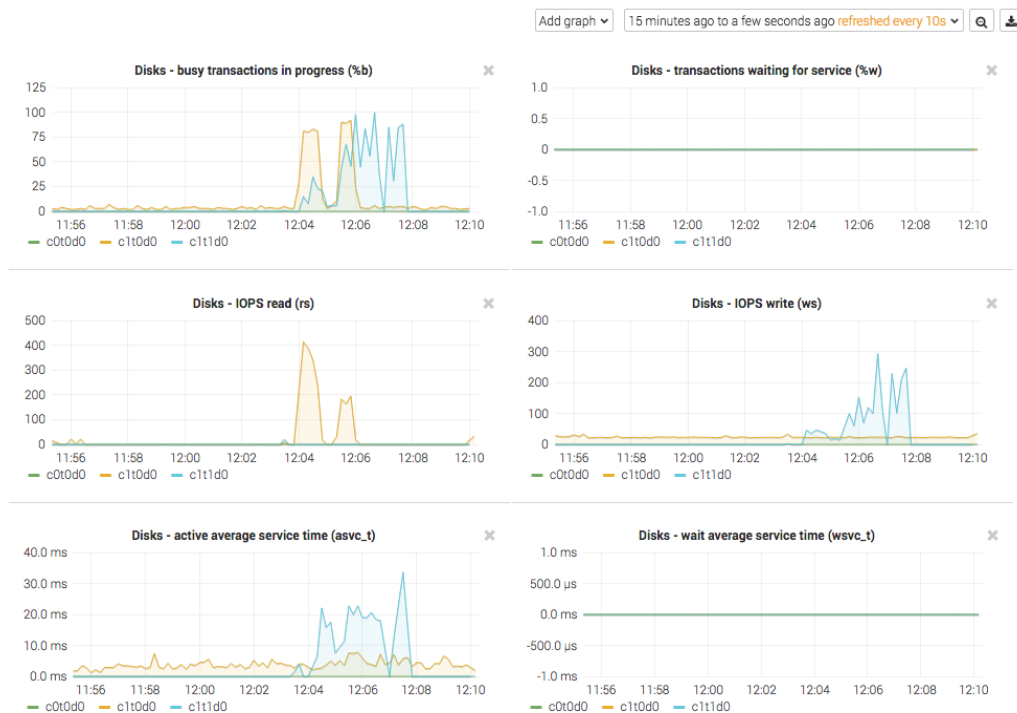
Low level analytics

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.

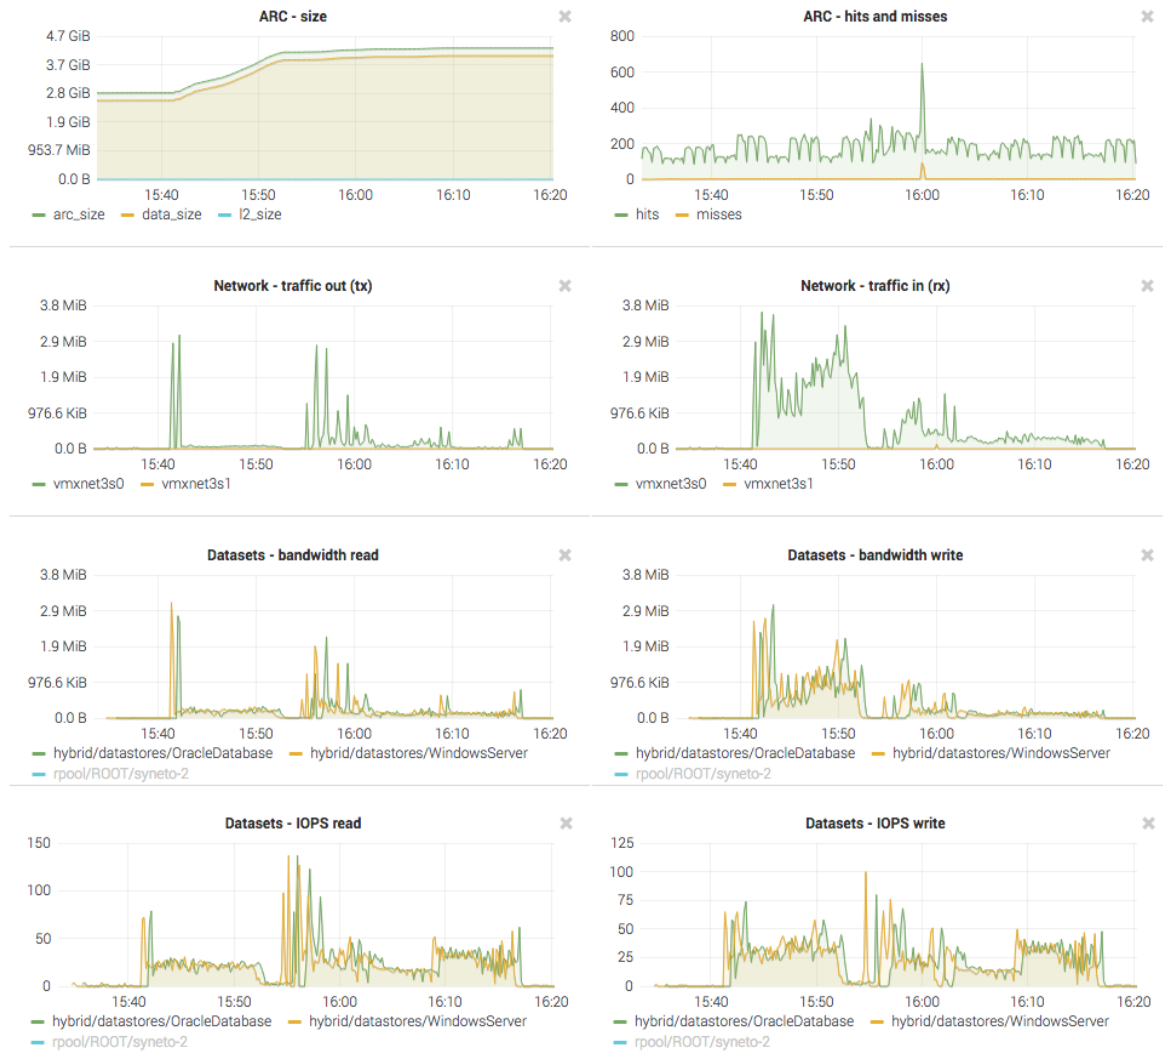
For example one virtual machine can consume all available resources and make other virtual machines starve for resources as well, leading to general slowness of the whole infrastructure.

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

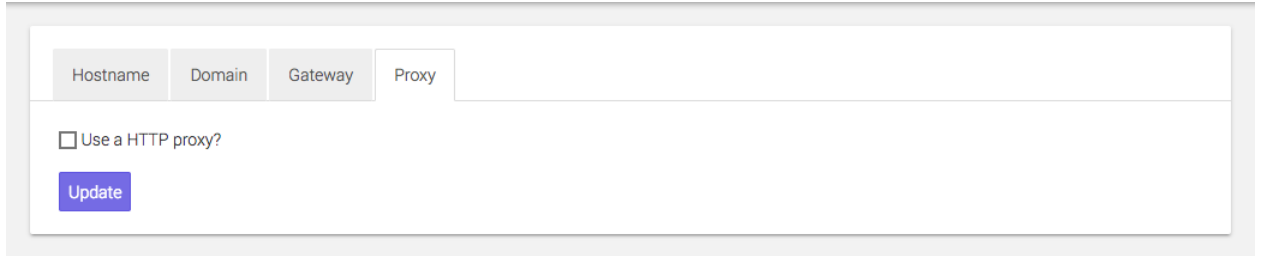
ID	7a5a3d15390e1ce03347193bb6f6027a
Serial number	SYN0001
Processor	4x GenuineIntel Intel(R) Xeon(R) CPU E5-1620 v4 @ 3.50GHz
Memory	8.0 GB
Data disks	5
Data storage (raw)	7.3 TB
Network	2 x 10 Gb/s

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](#)

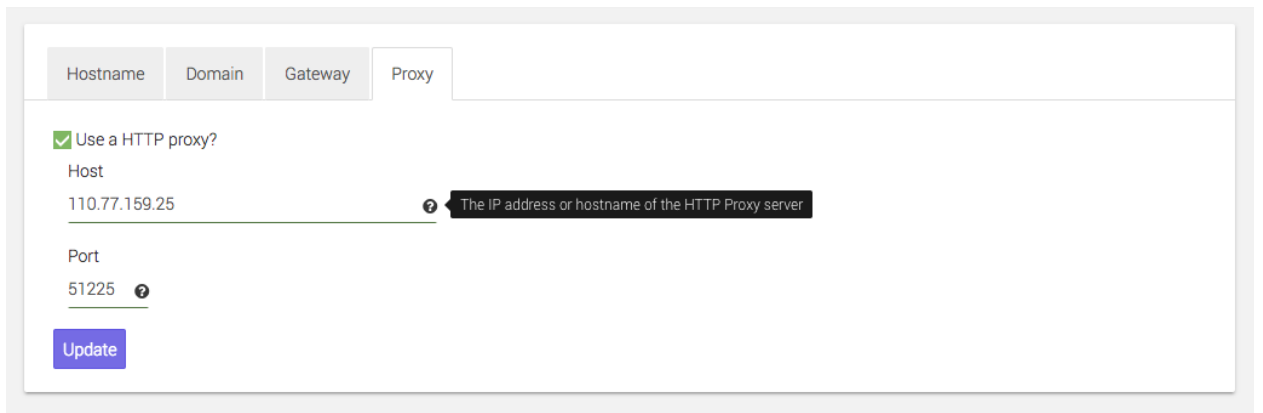
Internet access is required to check for updates. If the storage needs a HTTP proxy to access the internet, configure the proxy on the Network → Basic configuration page.



Hostname Domain Gateway Proxy

Use a HTTP proxy?

Update




Hostname Domain Gateway Proxy

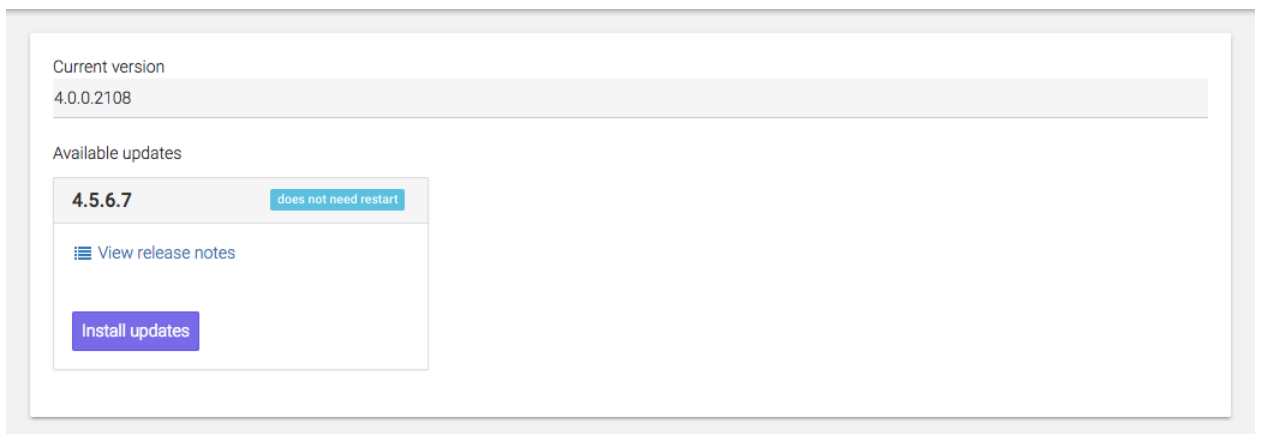
Use a HTTP proxy?

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

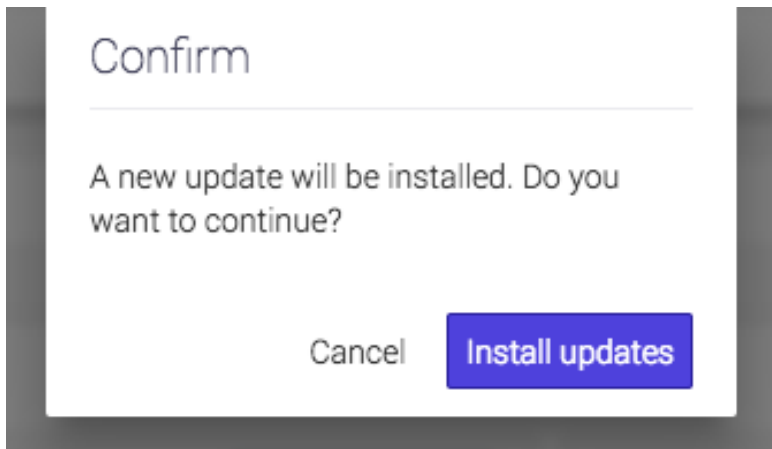
[View release notes](#)

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

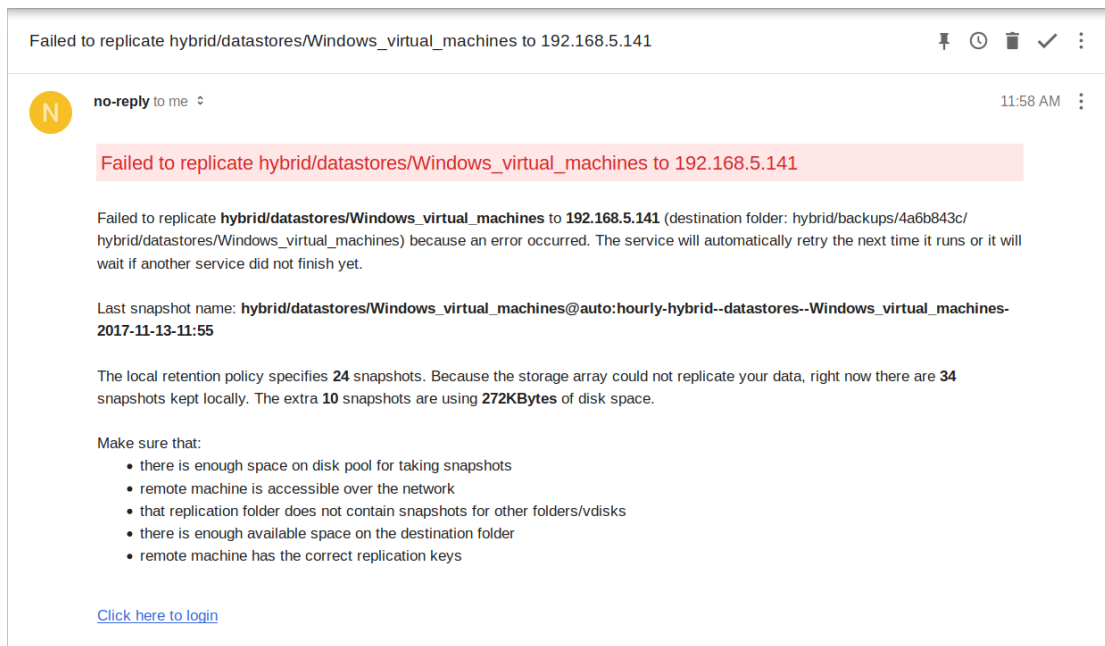
New versions that contain changes which touch the core of the operating system will require a reboot. Versions requiring reboot are marked with a red warning label. We recommend the scheduling of this type of update outside working hours.

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; text-align: center;"> <p style="color: green; font-weight: bold;">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.

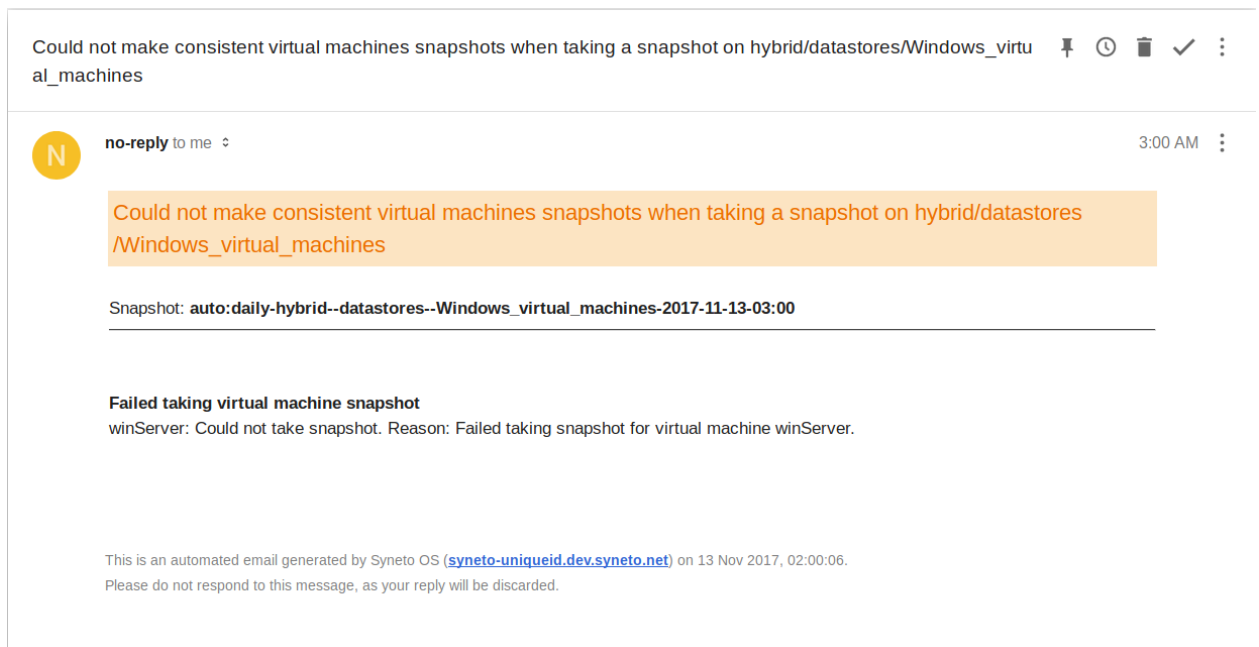
Can't connect to Active Directory



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:



⚠ Could not create filesystem with path ✖
/shares/smb/smb-share-test.
17:45:56 - 6 Nov, 2017

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

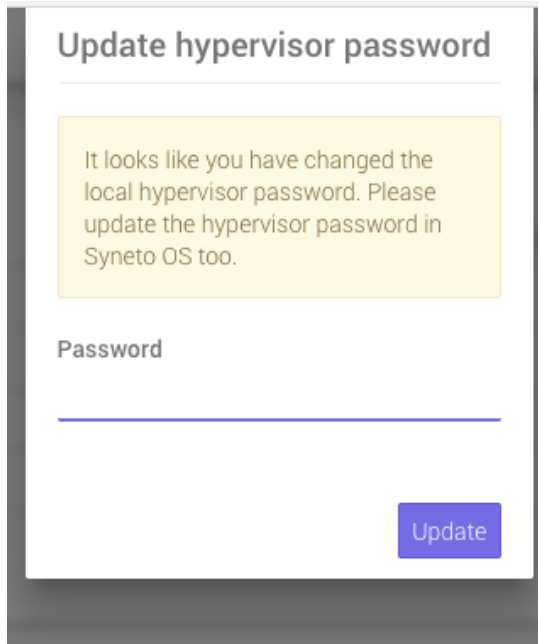
No connection with Central
Linked with account **user@my.domain**

[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.



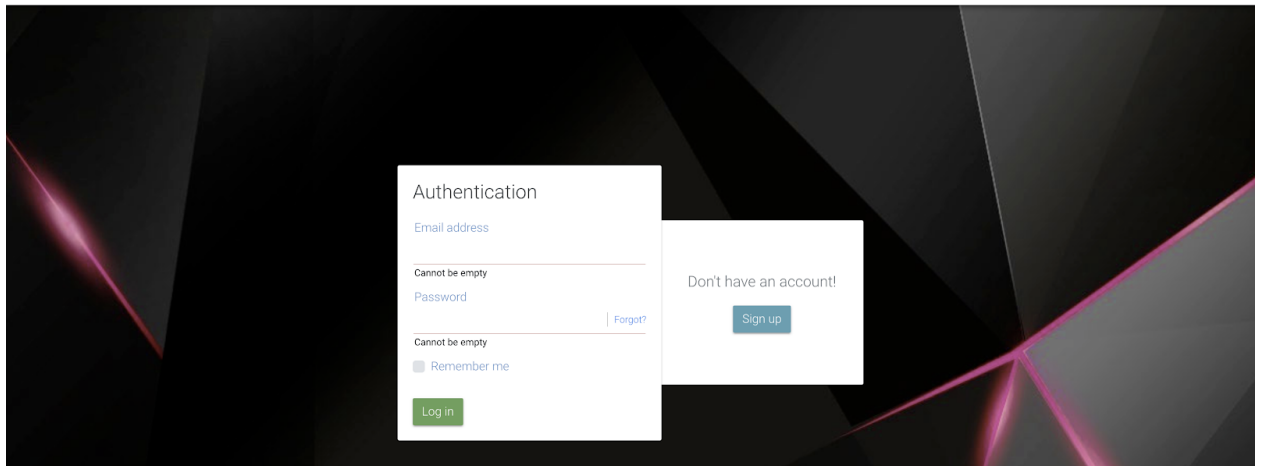
The screenshot shows a dialog box titled "Update hypervisor password". Inside the dialog, there is a yellow message box that reads: "It looks like you have changed the local hypervisor password. Please update the hypervisor password in Syneto OS too." Below the message box, there is a label "Password" 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
2. Contact details	Password	Re-type password

3. Company	First name	Last name
	Name	Surname
	Phone number	
	00303203556179	
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
.....

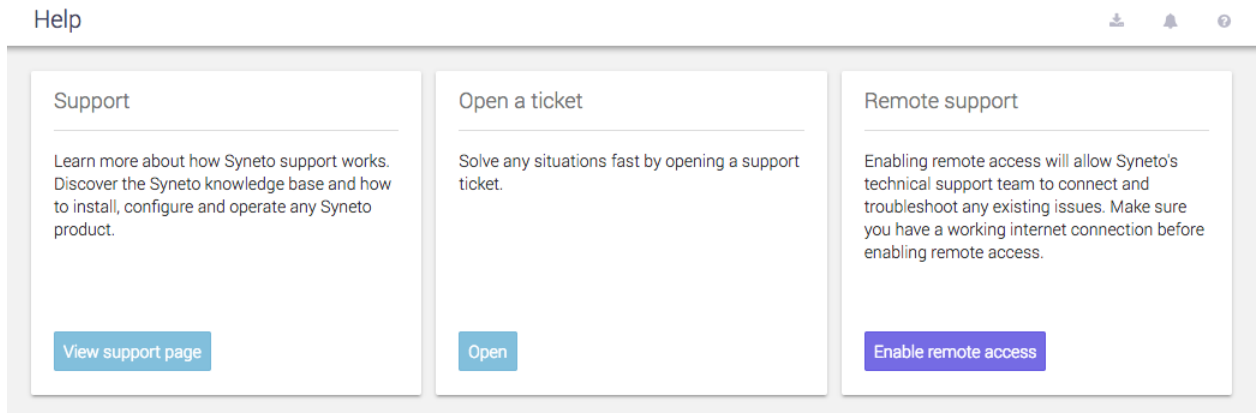
Re-type new password
.....

CHANGE PASSWORD

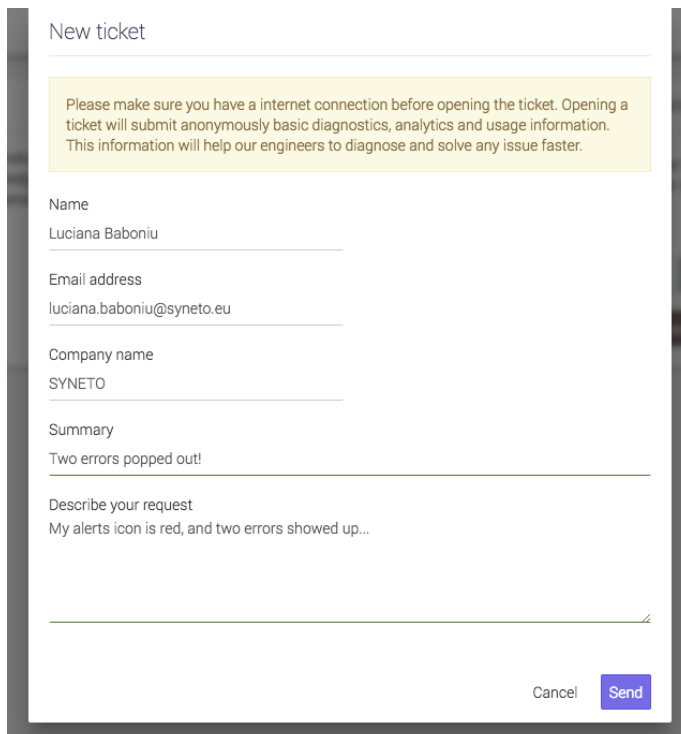
APPENDIX B: ACCESSING CUSTOMER SUPPORT

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.



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.





The 'New ticket' form is pre-filled with user information:

- 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...

At the bottom right, there are 'Cancel' and 'Send' buttons.

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.

After enabling remote access, just click the  button in order to copy to clipboard the numbers provided, and paste them (ctrl + v) to your open conversation with our support engineers.



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.