Syneto OS
User Guide

Version 4.1.0

December 2017

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

GATHERING REQUIRED INFORMATION
- vSphere license
- IP addresses
- Finding IP addresses
- vCenter and ESX 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
- 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

DISKS AND POOLS
- Introduction
- Replacing a Faulted disk
- Replacing a Faulted OS Disk
- Replacing a Faulted cache disk
- Attaching an expansion shelf
  - Creating a pool
  - Expanding a pool
- Importing a disk pool
- Exporting a disk pool
- Protecting against silent data corruption

DATA PROTECTION AND REPLICATION
- Introduction
- Configuring replication targets
  - Add a replication target
  - Remove a replication target
- Protecting a datastore
  - Configuring the snapshot schedule for the datastore
  - Canceling a running replication
  - Configuring the type of snapshots for the virtual machines
- Protecting a share
- Protecting a volume
- Monitoring

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 disks
  - 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 technical 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 ESX 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 TECHNICAL SUPPORT
- Contacting Syneto Technical 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. If 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 **4.1.0.**

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 Central[1] 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.

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 licences, 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 recommended.

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, and starting from 4.1.0, there is a method to find all 3 from the physical console of ESXi.

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

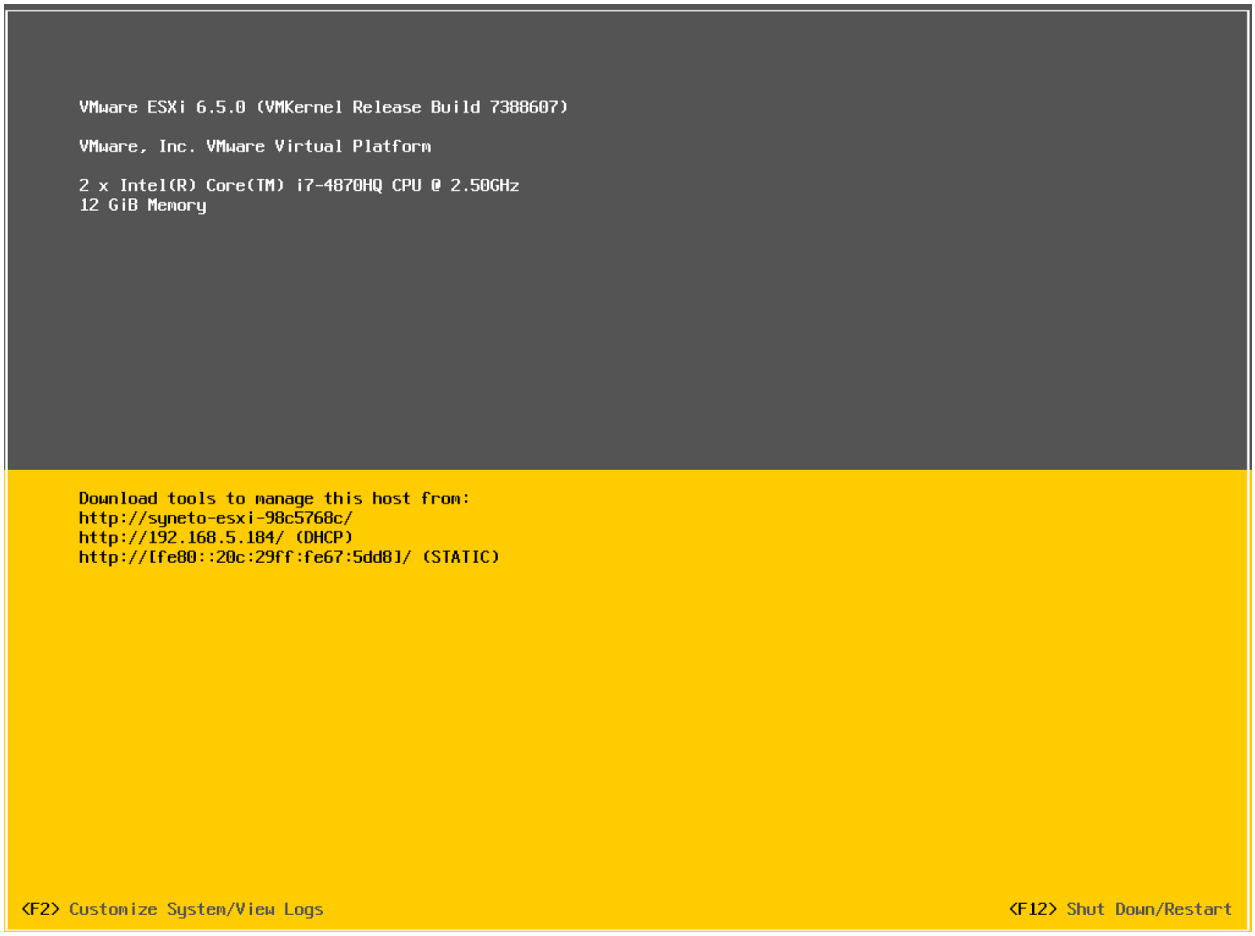

The SynetoOS (and IPMI management - if configured) IP addresses are shown on the ESXi service console shell, accessible using **Alt+F1** on the physical video console:

vCenter and ESX hosts credentials

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

- IP address / hostname
- Username
- Password

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

For more information about connecting to vCenter, can be found in the Join vCenter chapter.

Active Directory credentials

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

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

- Domain name

  NOTE: The domain name is case sensitive.

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

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

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

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

More information about when and how to configure Active Directory integration can be found in the Shares chapter.

ACTIVATE THE WEB MANAGEMENT INTERFACE

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

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

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

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

Access ESXi web console

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

Press `F2` and enter the default ESXi username and password (username: `root`; password: `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).

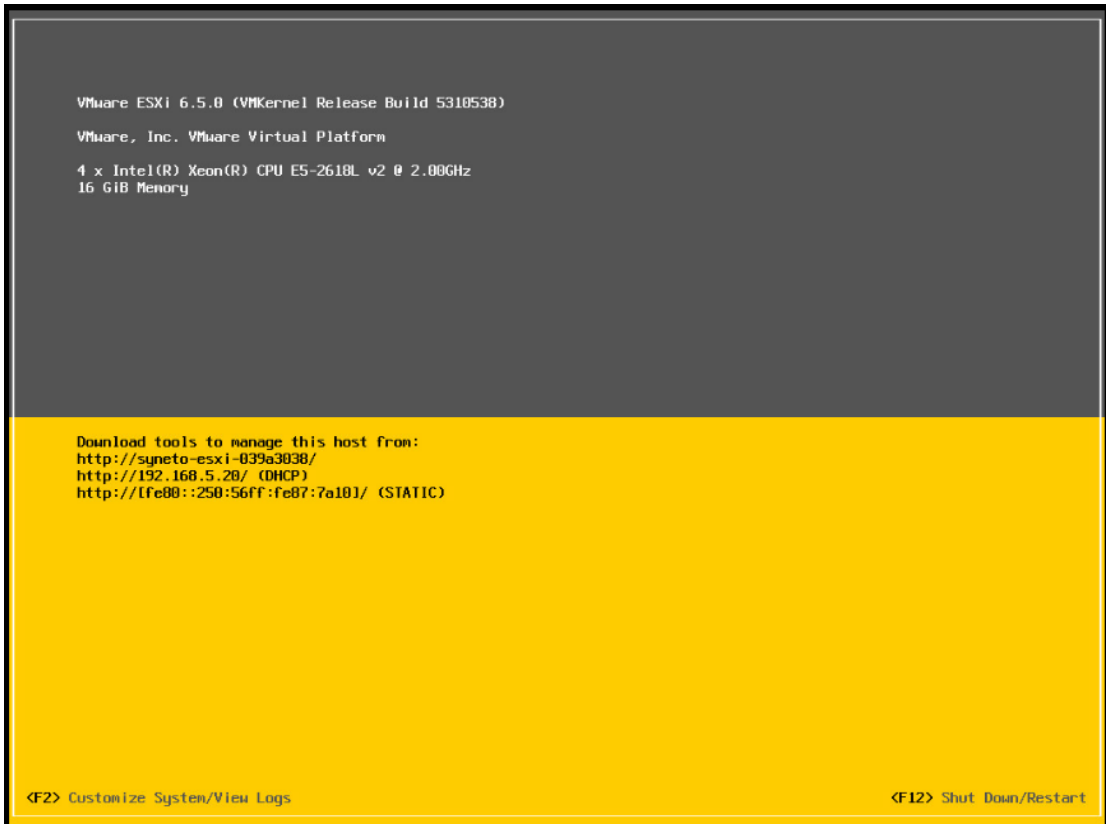

On the ESXi console, you can see the IP address to the ESXi UI. Insert the address in a browser (of your choice) and log in to the UI with the default credentials (default username: `root`; password: `syneto123`).

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`).

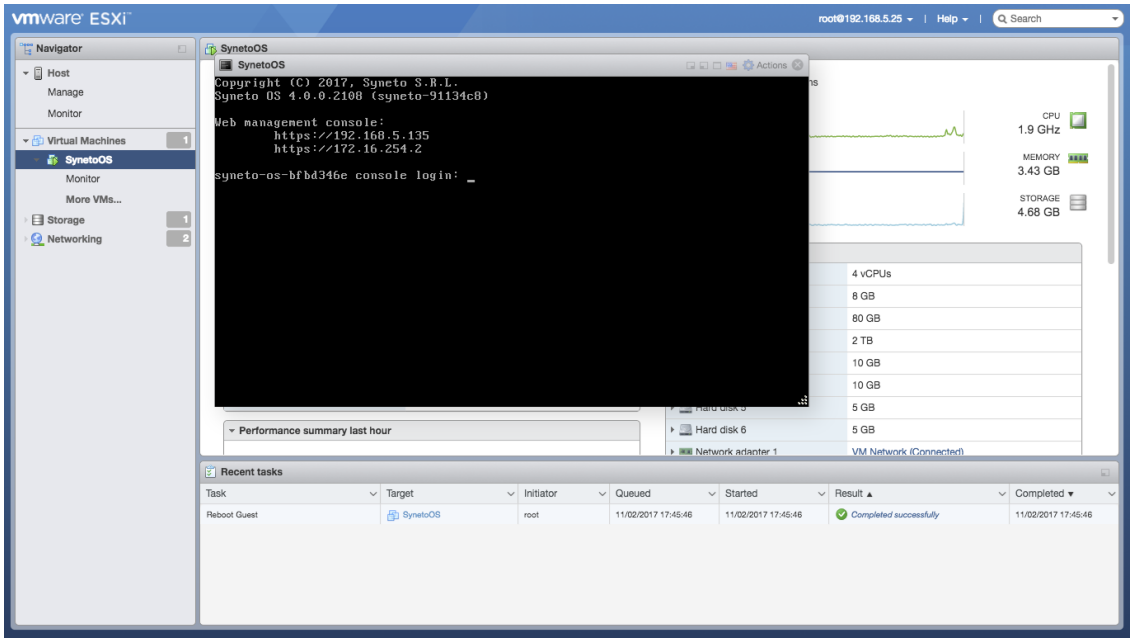

If you do not have DHCP, or want to override the default values for IP, DNS and gateway, run the following commands:

```
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 technical 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.

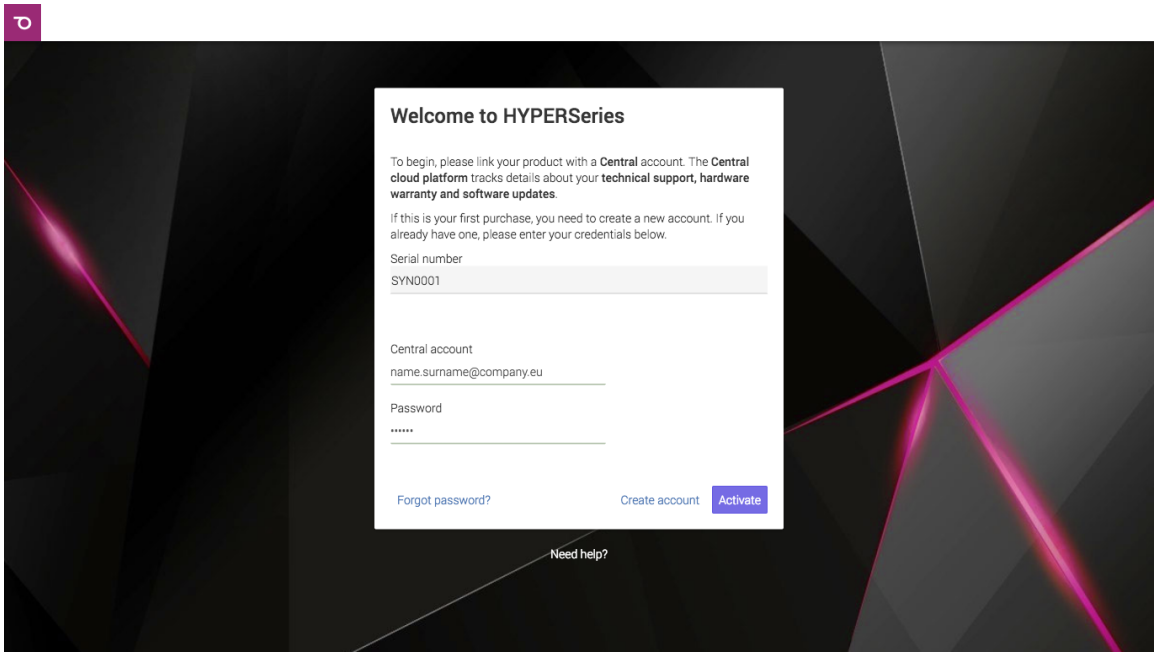

NOTE: Log in to the management UI is not possible without registering the product to Syneto Central first.

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.

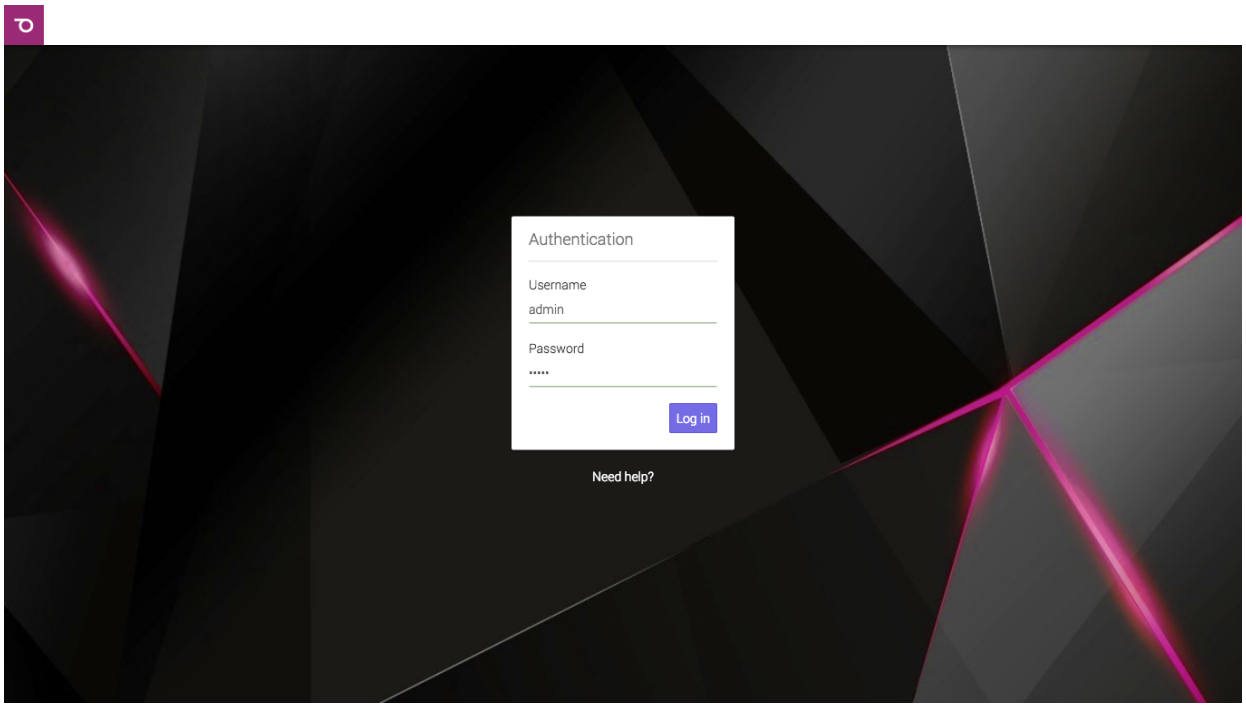

Configure using IPMI

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.

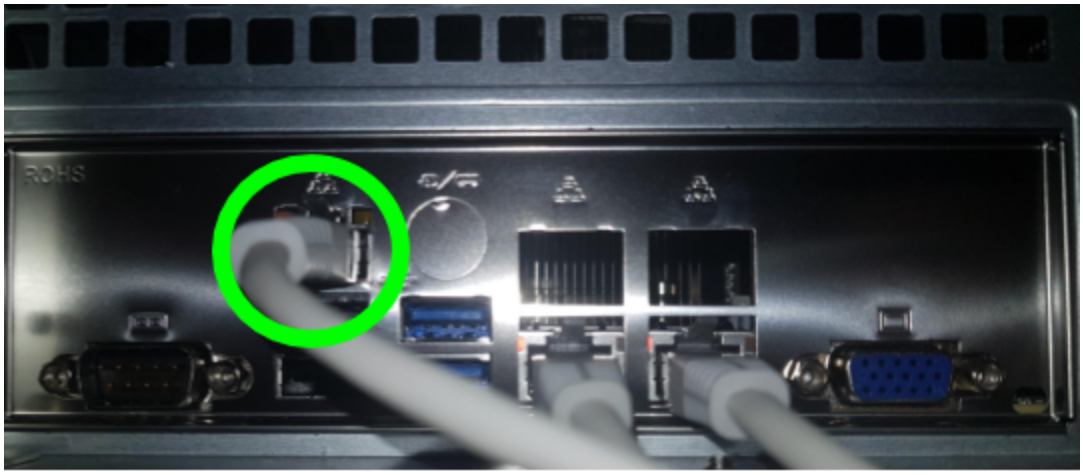

The IPMI interface will acquire a network address over DHCP.

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: `BCM IP:....`

Default credentials

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

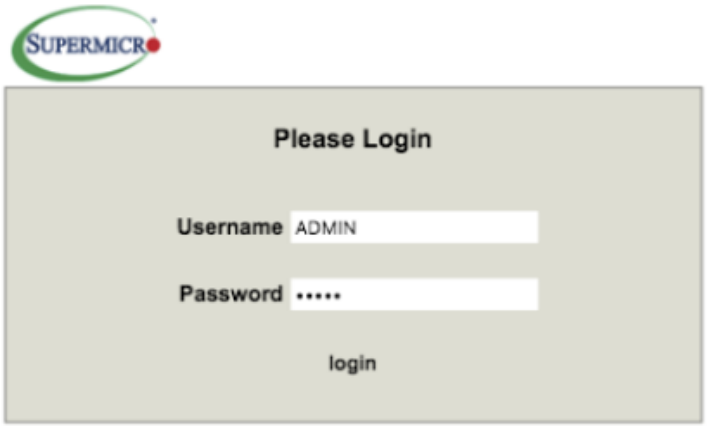

Management interface main components

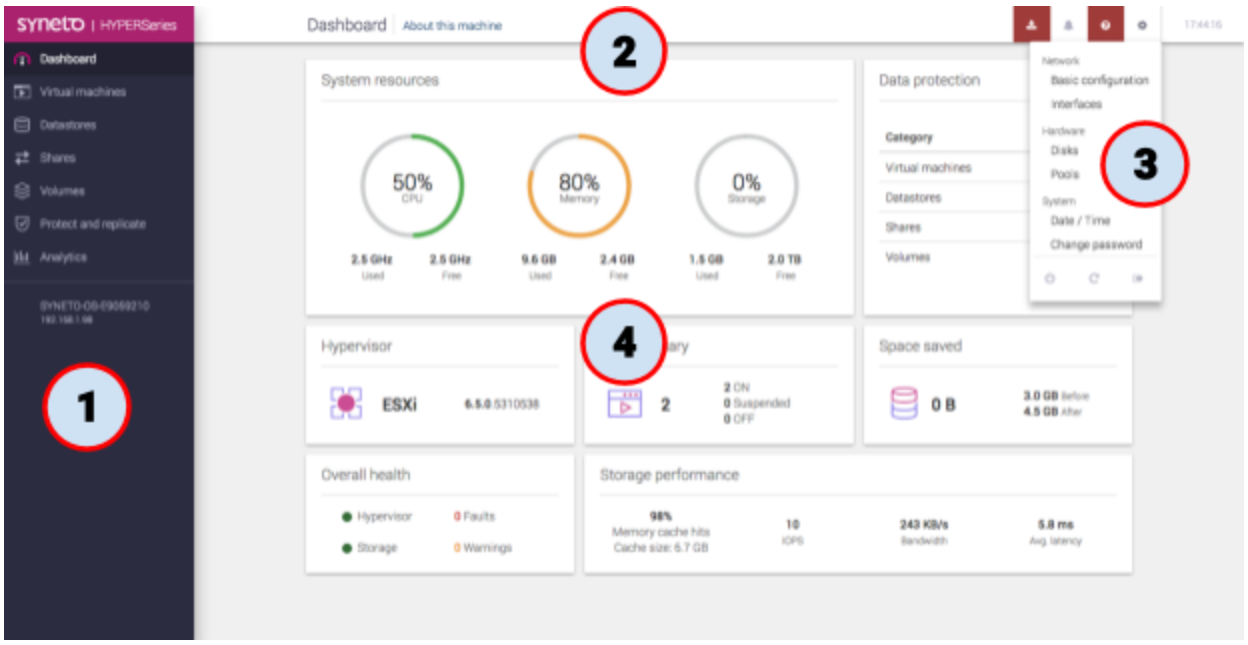

These are the main components of the management interface:

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 Technical Support first.

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

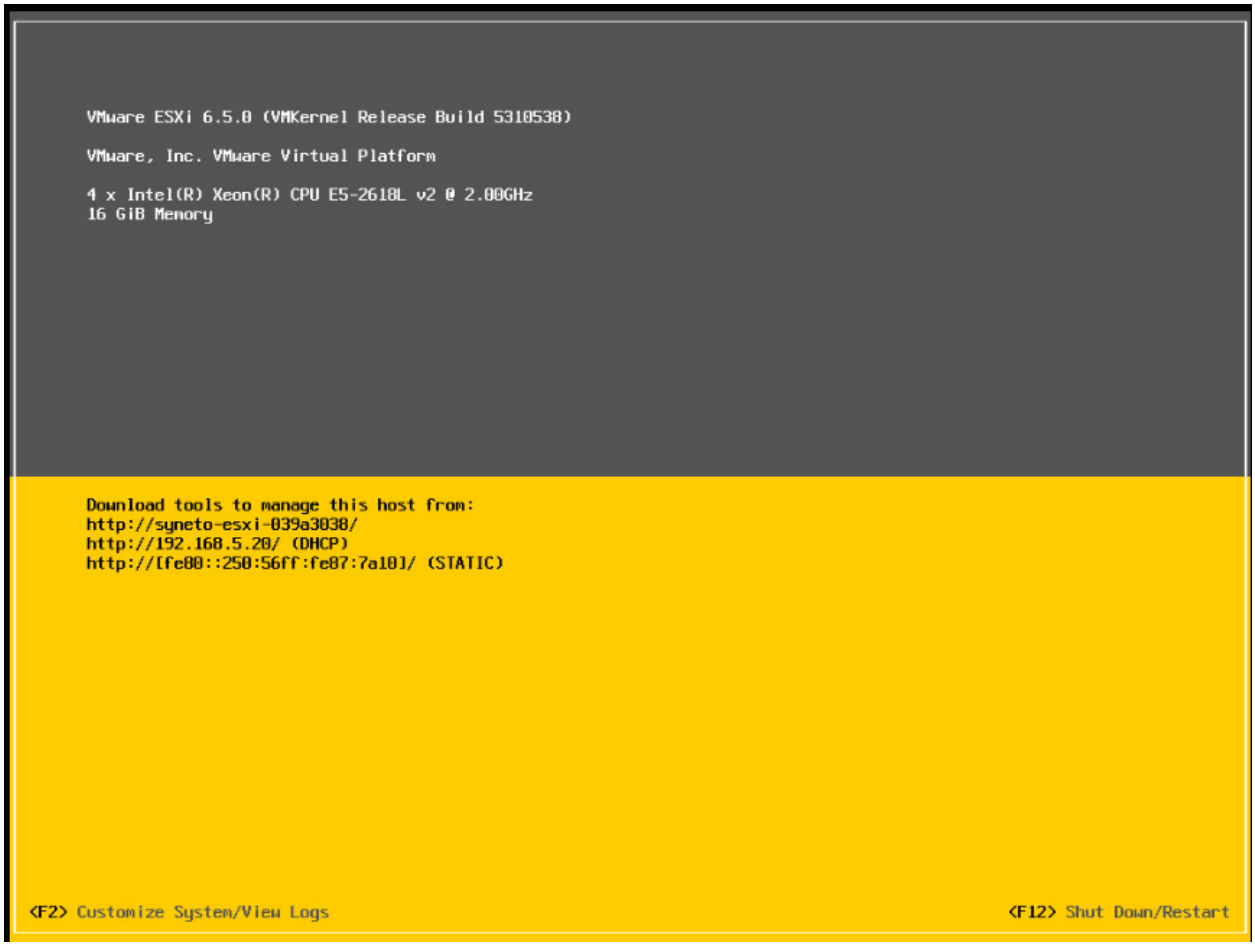

Press *F2* and then enter the *username* and *password* of the ESXI. A menu with different actions will open.

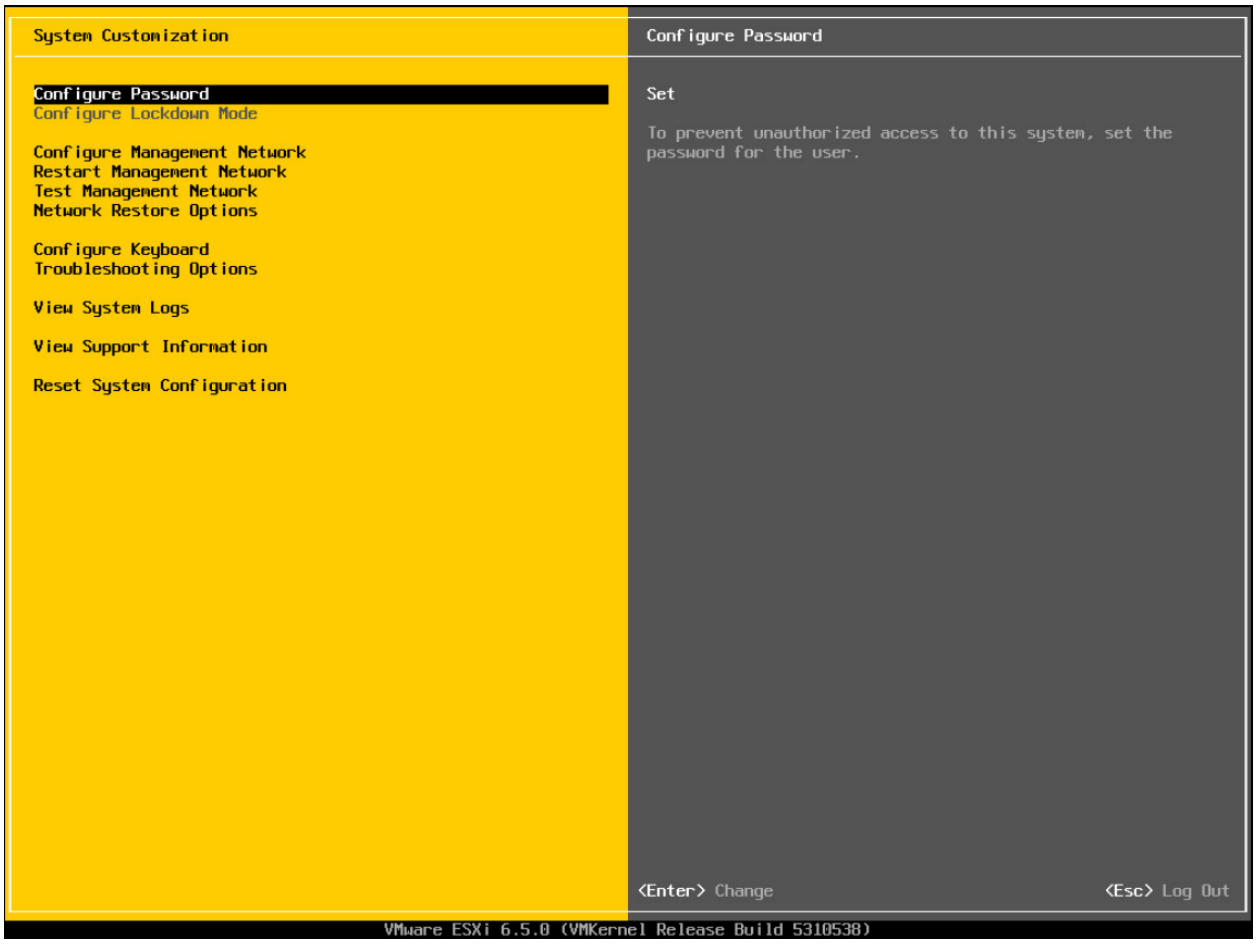

Select *Configure Management Network* and press *Enter*. A menu with network configuration options will open.

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

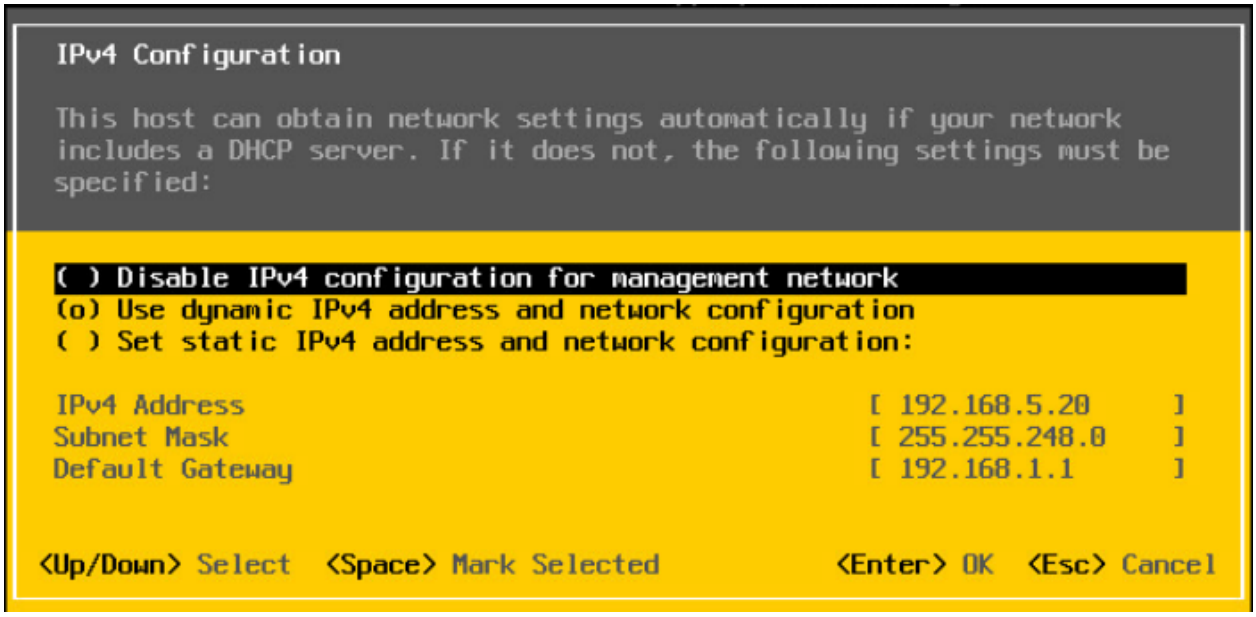

You can select a dynamic or static IP. By default, the dynamic IP option is selected. Also, the *gateway* can be set from this dialog. After you have finished configuring , press *Enter* to save and exit. Select *DNS Configuration* and press *Enter*.

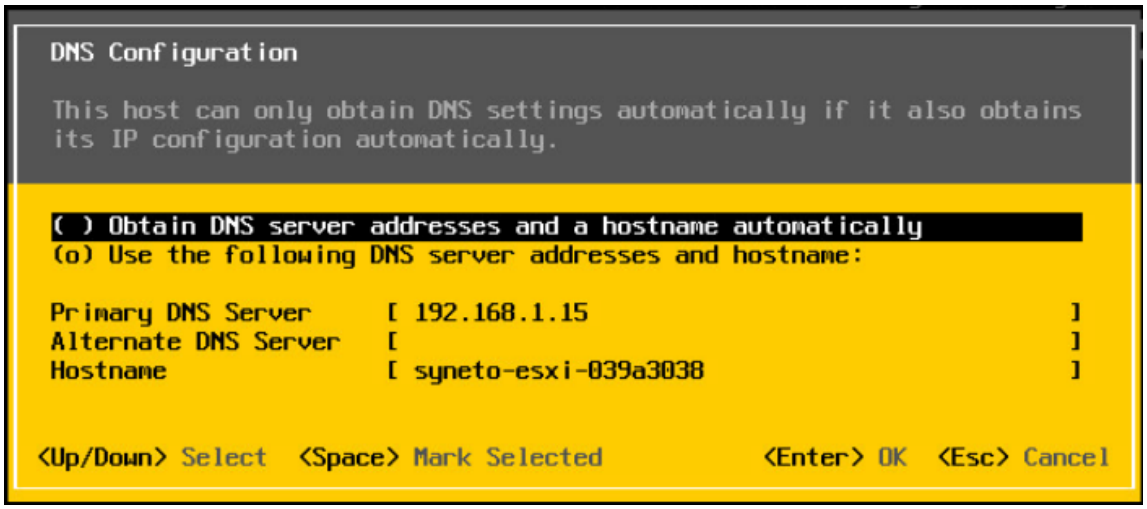

You can select a dynamic or static DNS and hostname. If you select static, fill the *Primary DNS Server* and *Hostname* fields. The *Alternate DNS Server* field is optional. Press *Enter* to save and exit.

Install ESXi license

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

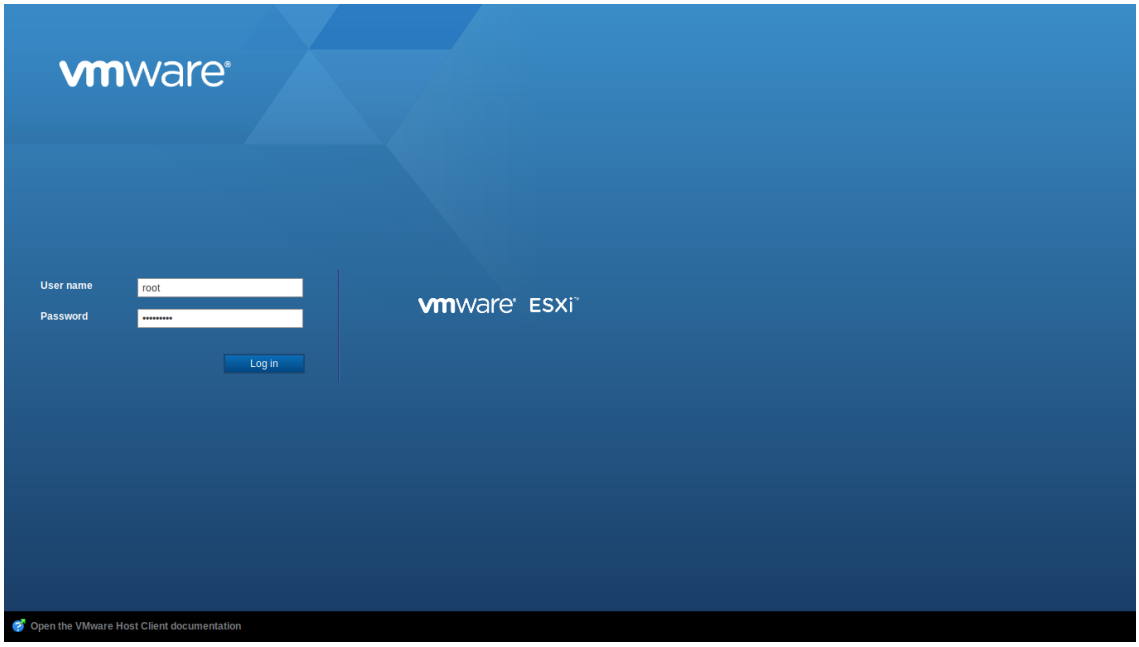

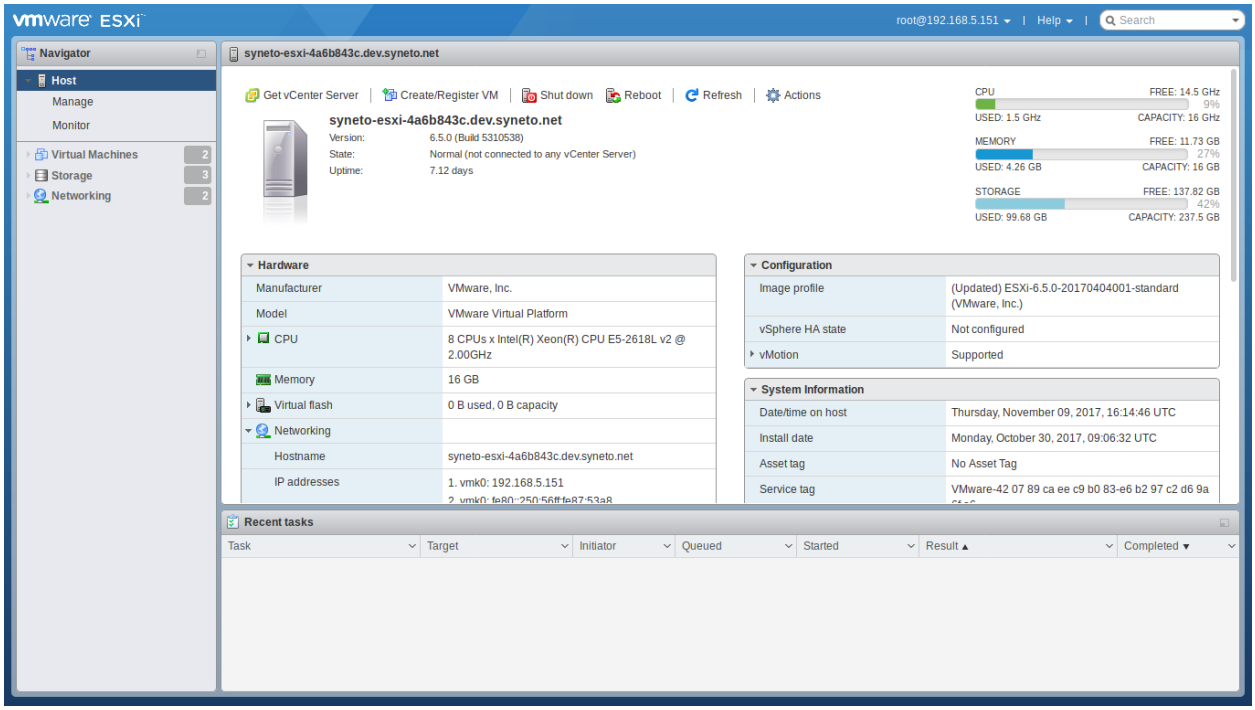

Select *Host* from the *Navigator* on the left and click on *Manage*.

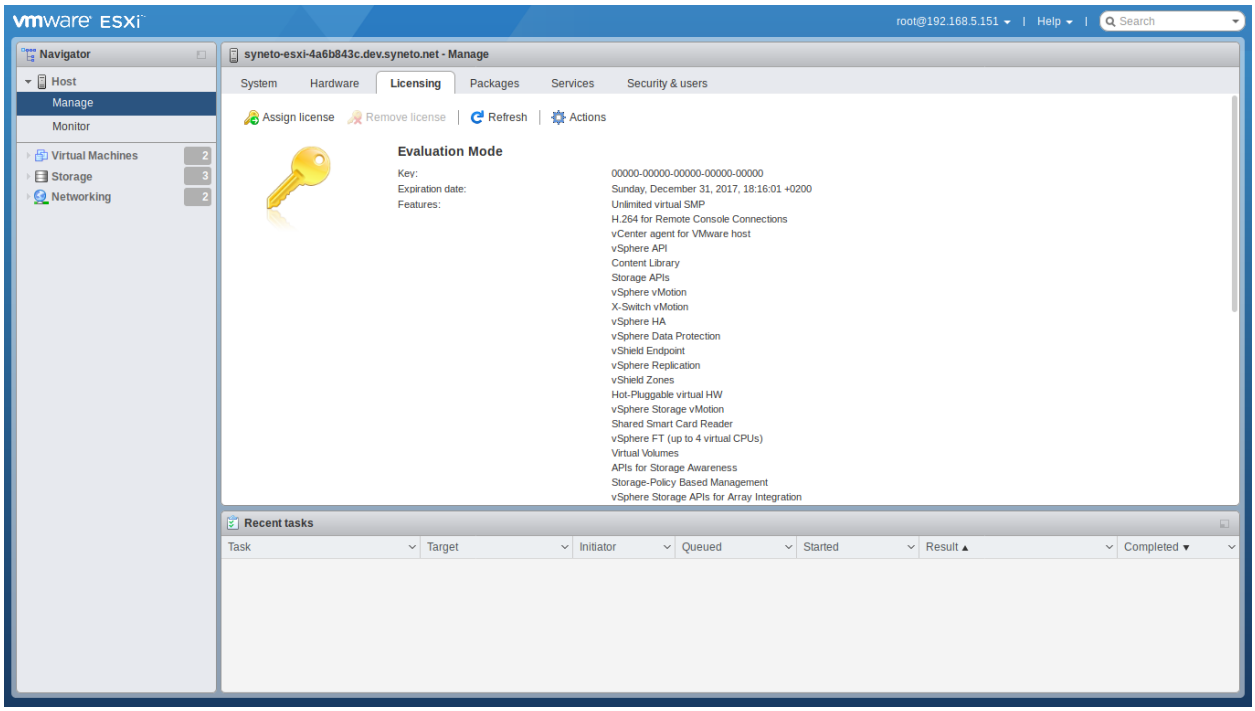

Select *Licensing* tab → *Assign License*.

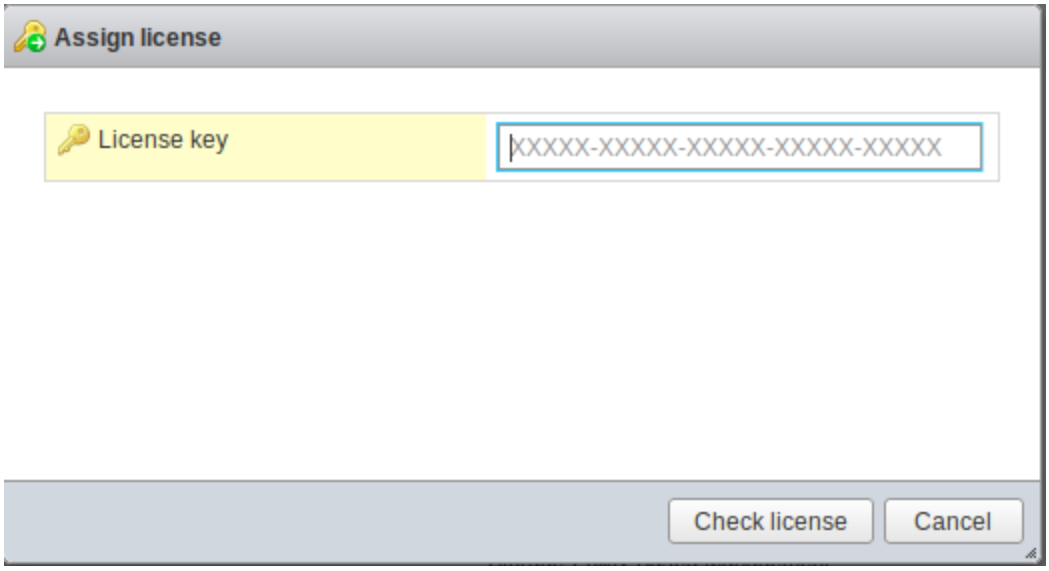

Type or paste the ESXi license key. Click *Check License*.

Configure Management VM networking

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

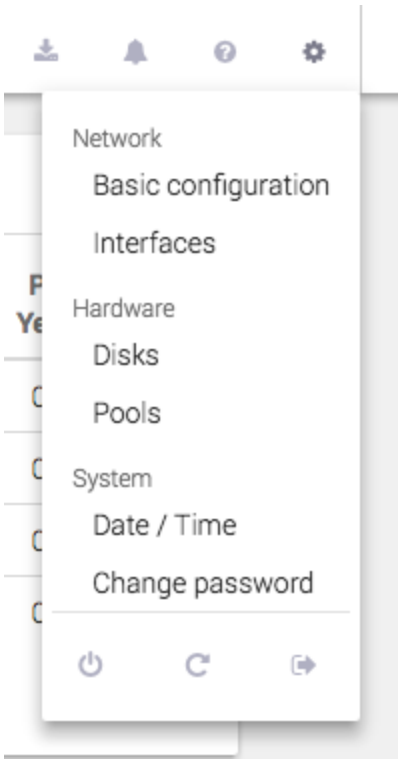

The *Network* page will open. Here you can set the hostname, domain name, DNS, gateway and proxy server.

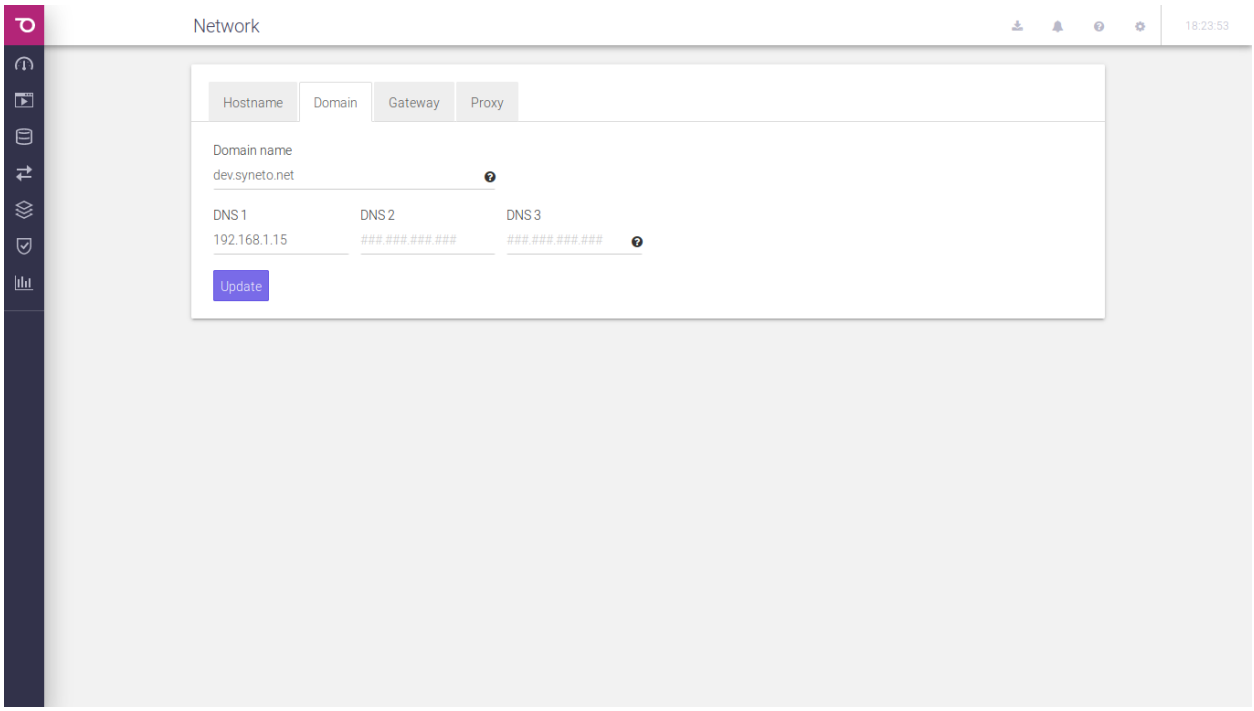

For more information, check the networking section.

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.

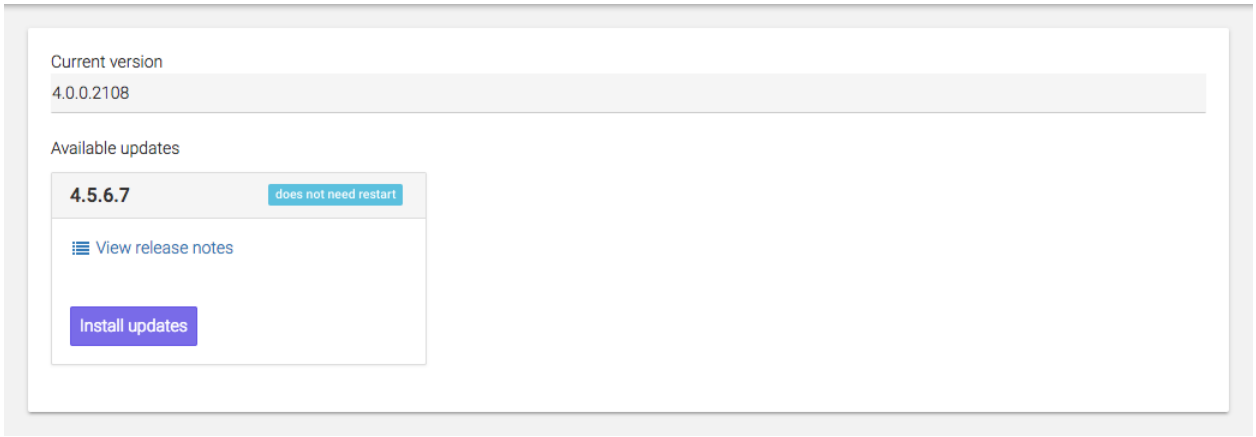

For more information, check the software update section.

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*.

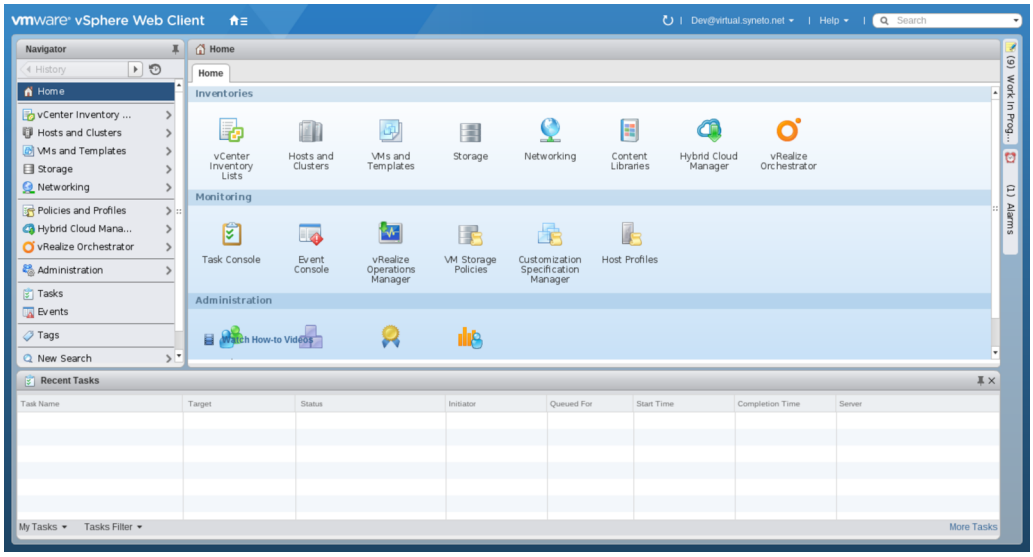

Select *Actions* → *Add Host*.

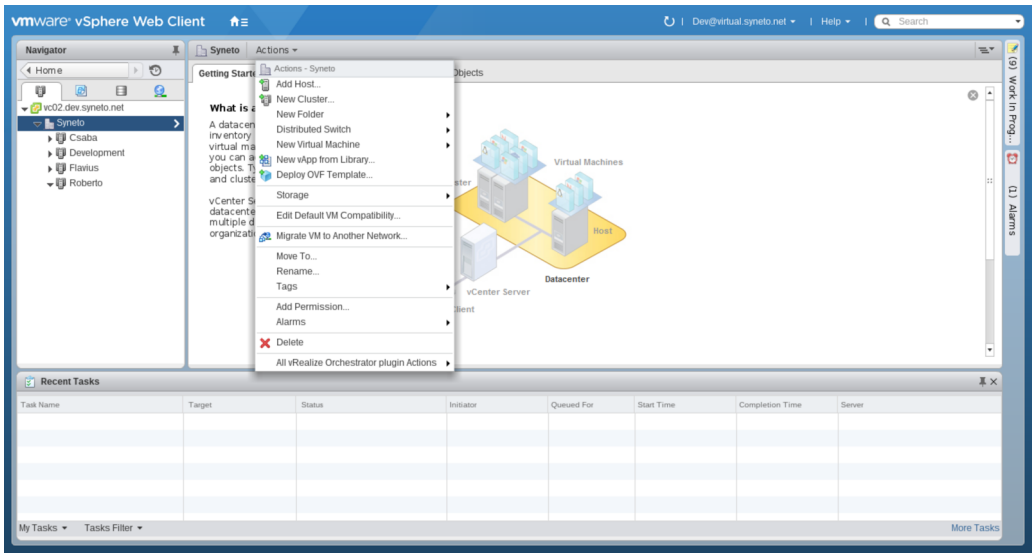

A dialog will open which will guide you through the process of adding a host.

Enter the hostname or IP of the Syneto ESXi. Click *Next*.

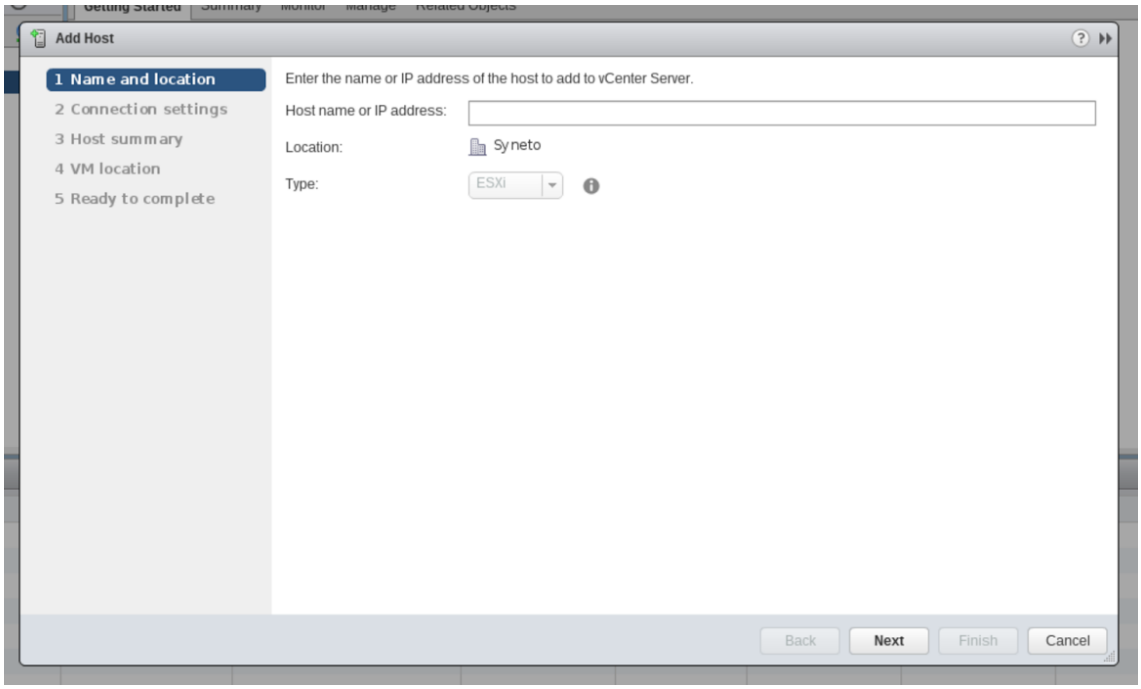

Enter the username and password of the Syneto ESXi. Click Next. If a security alert opens, click OK.

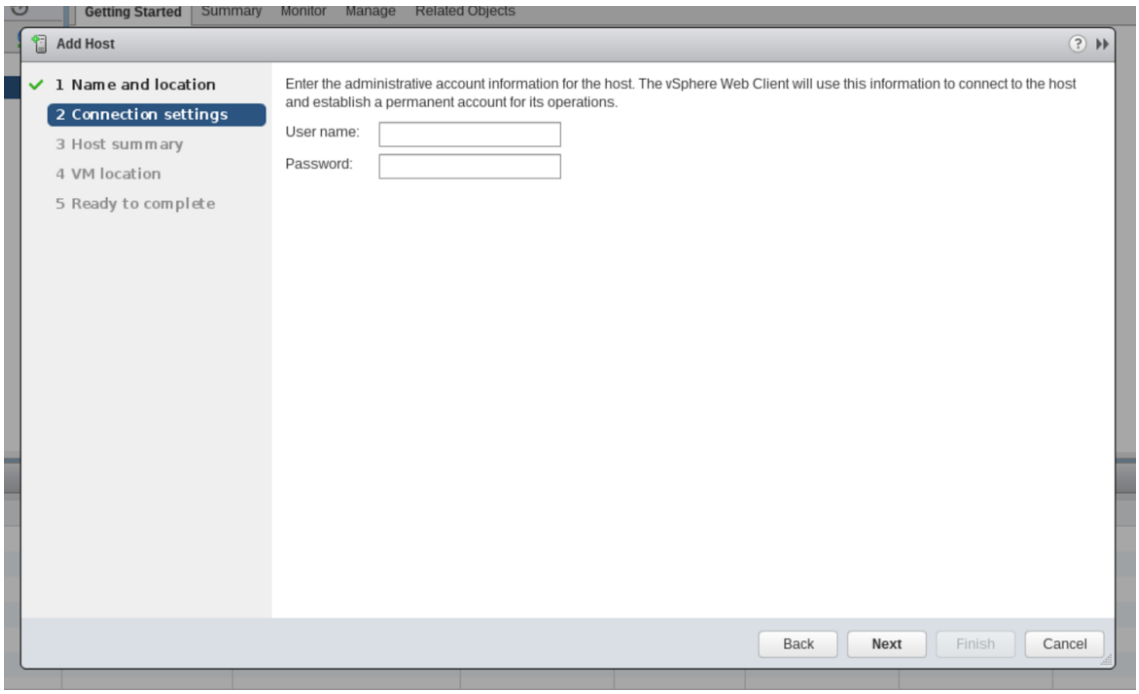

In this view, you can review the host configuration. After that, click *Next*.

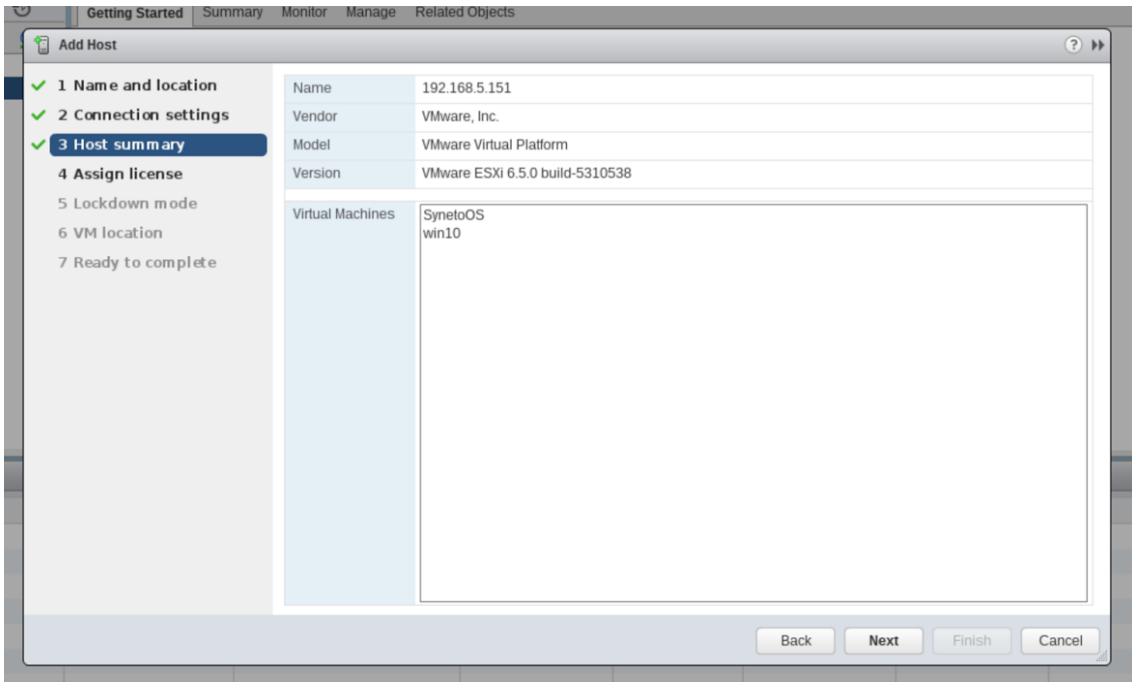

You will see a list of available licenses. You can assign one of those or add a new one by clicking on the *green plus*.

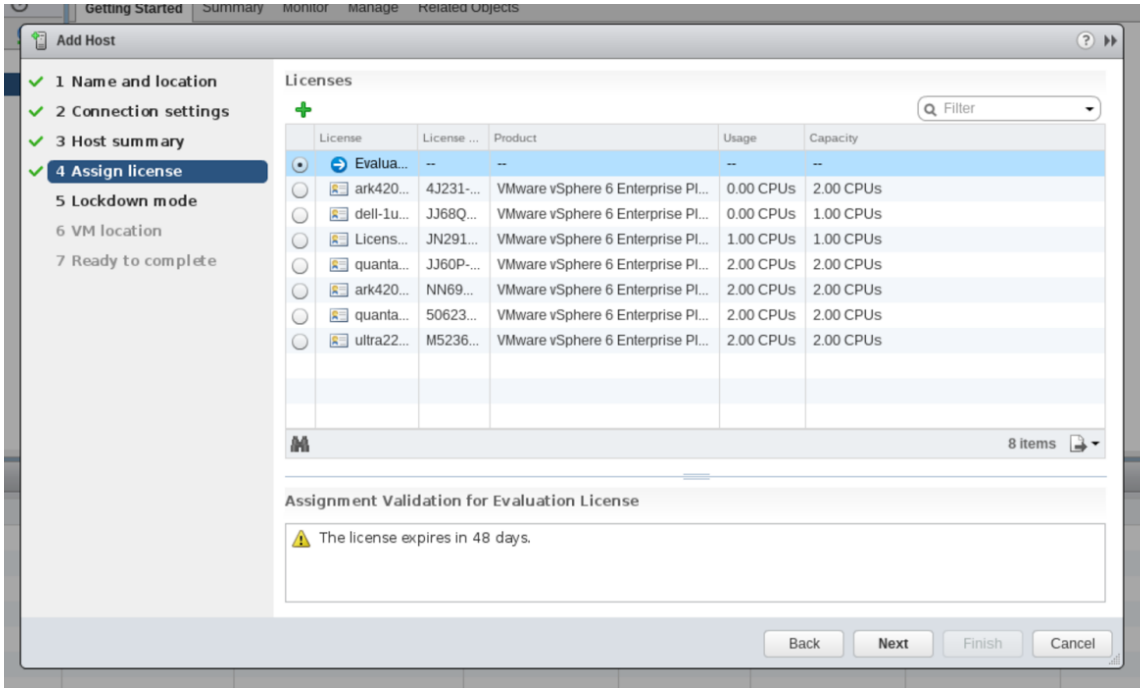

This window lets you configure the lockdown mode.

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*.

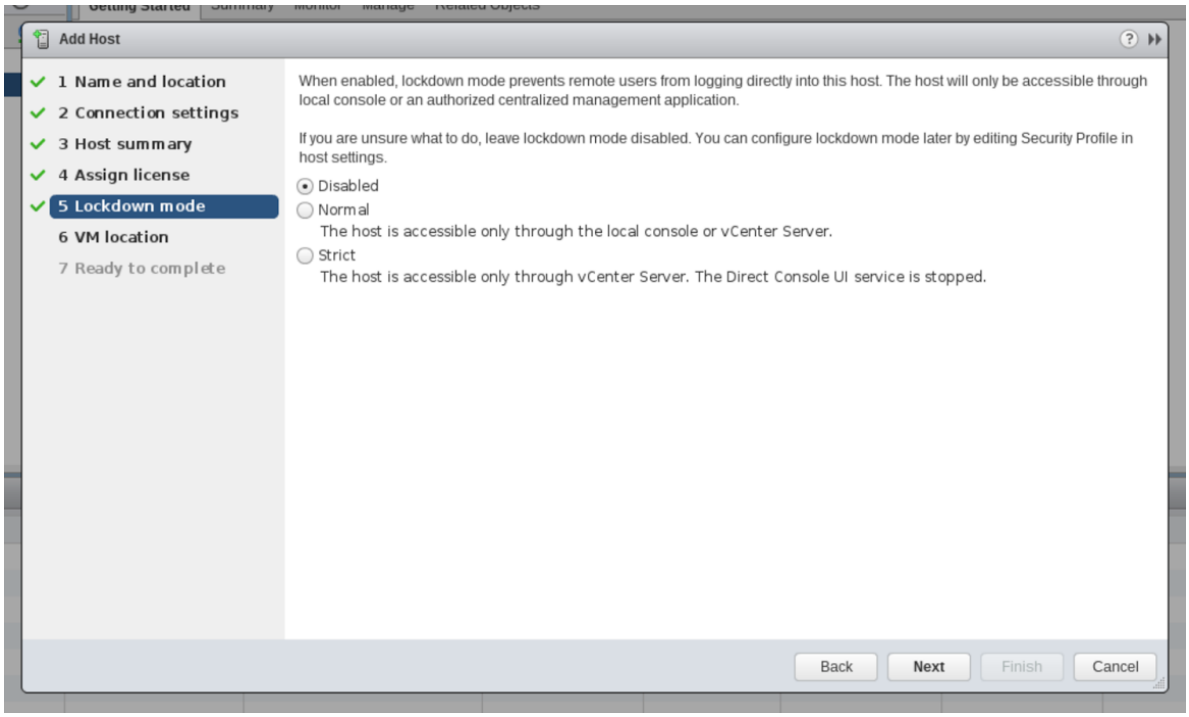

Select the virtual machine location. Click *Next*.

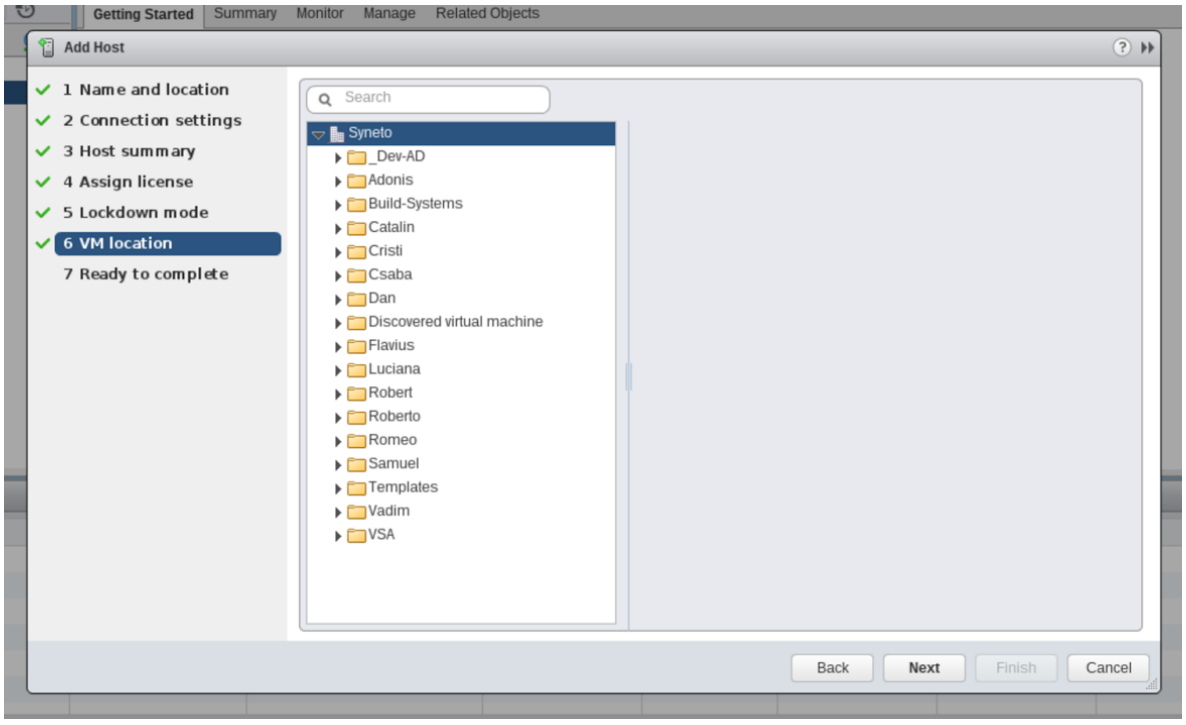

Review your selections and if you do not want to modify anything, click *Finish*.

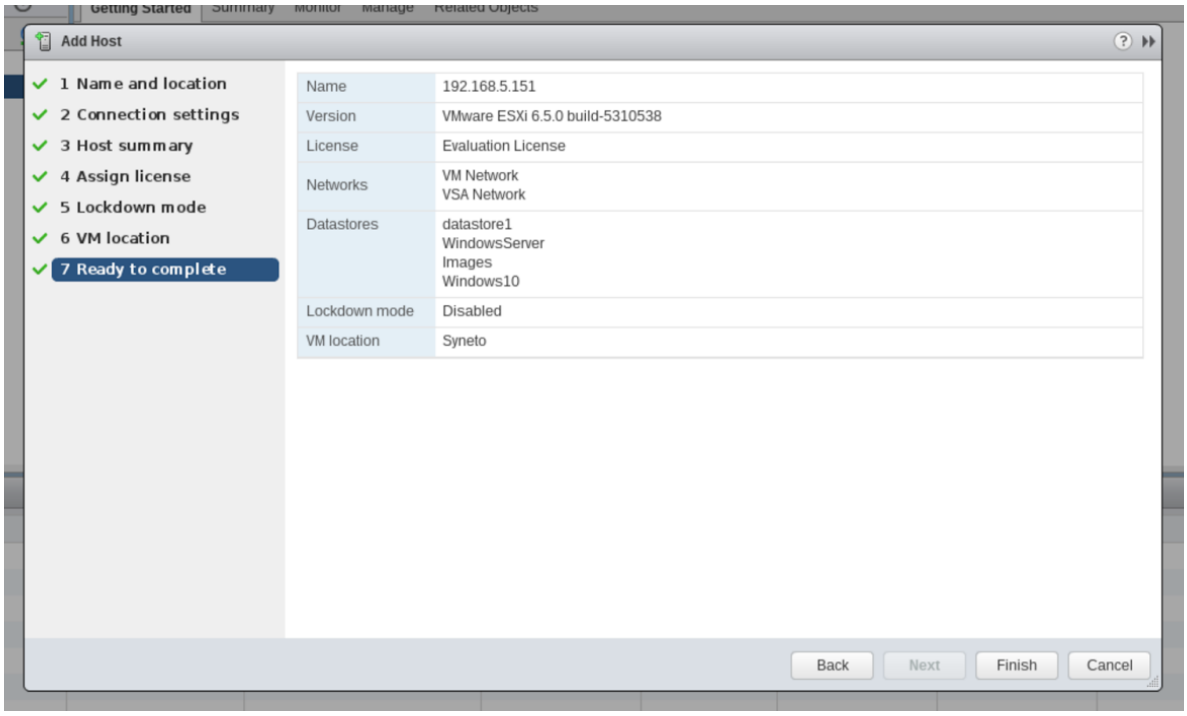

Now, we need to add vCenter to SynetoOS. Go to SynetoOS Management Interface, on *Datastores* page.

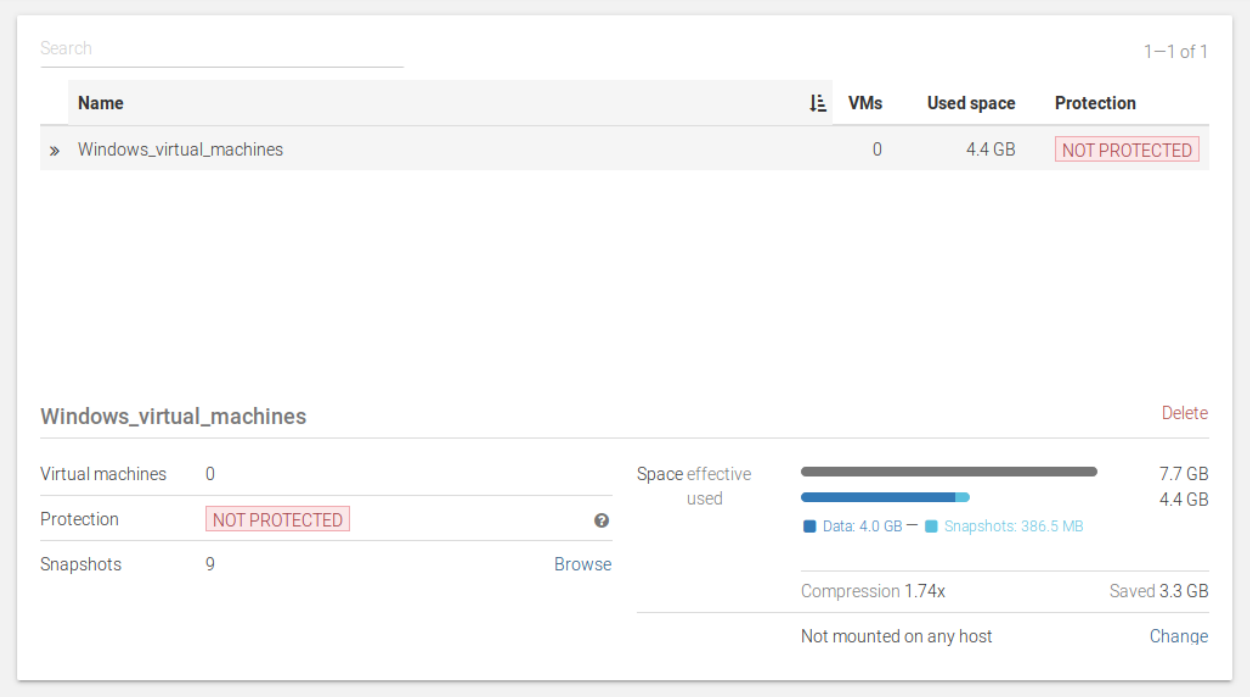

Click on *Hosts* in the action bar, top of the page. A dialog listing the registered host will open.

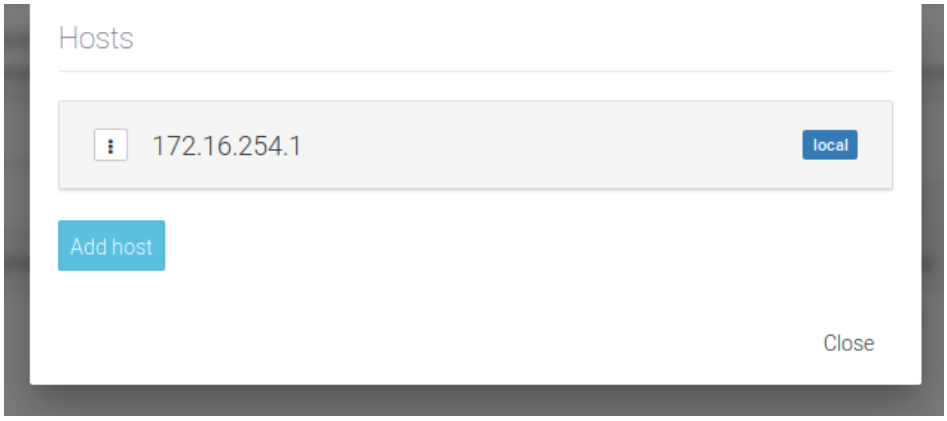

Click on *Add host*. A form will open. In the *host* field, enter the hostname or IP of the vcenter. Enter a *username* and its *password*. Click *Add*.

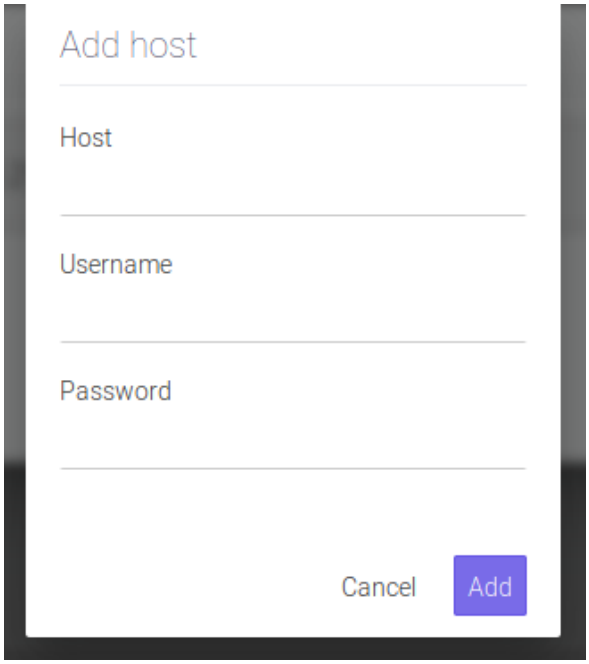

This will add the vCenter to SynetoOS. Now, you will be able to mount datastores on any of the vCenter's hosts.

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*.

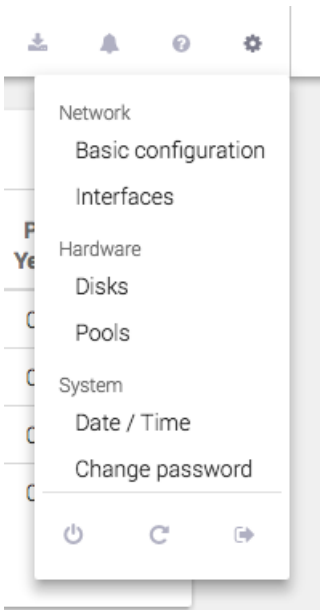

Click on *Date / Time* option. The *Time* page will open. Here you can set the date, time, and timezone. We recommend setting *Automatic time synchronization* on.

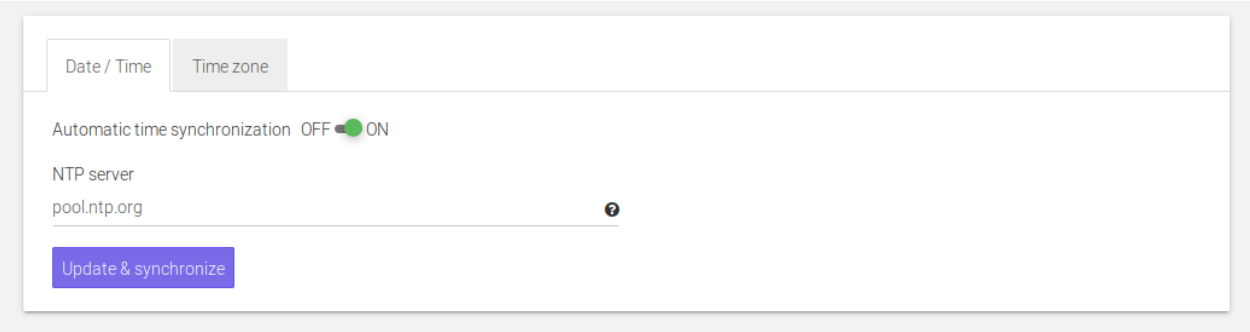

For more information, check the date/time section.

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 disk 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.

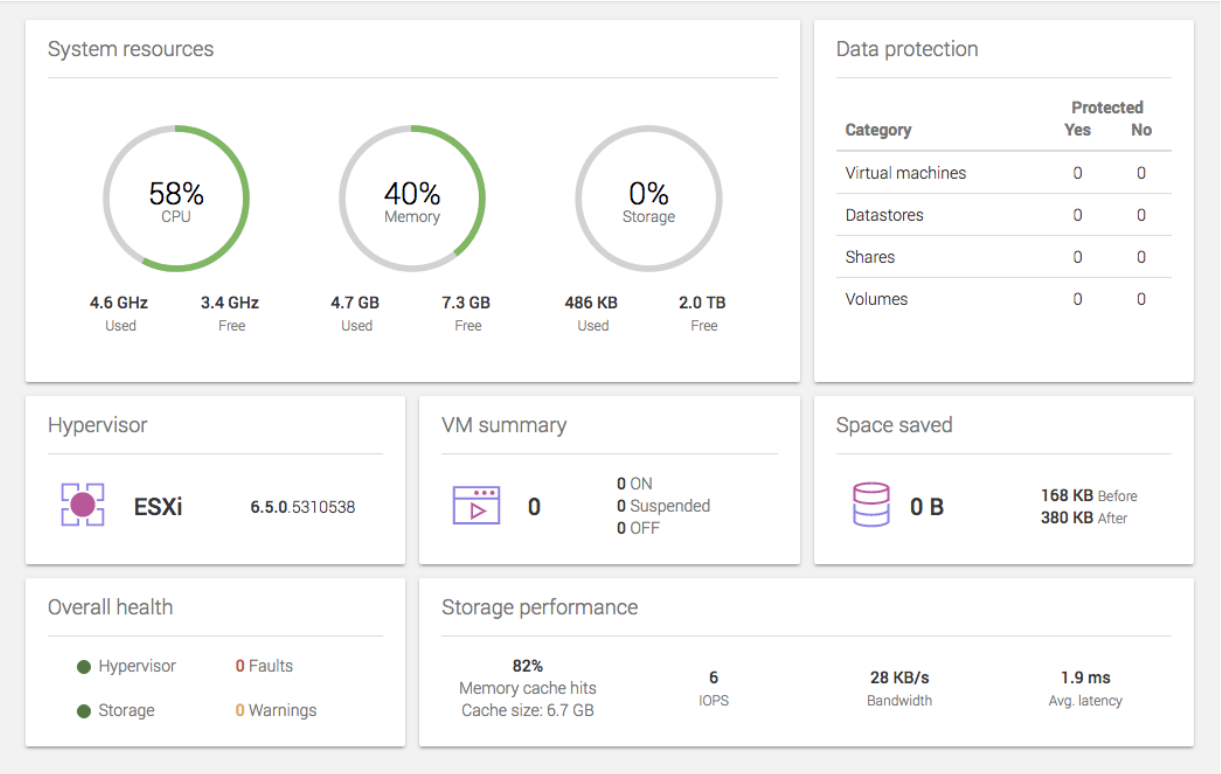

NOTE: The data protection widget does count clones.

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.

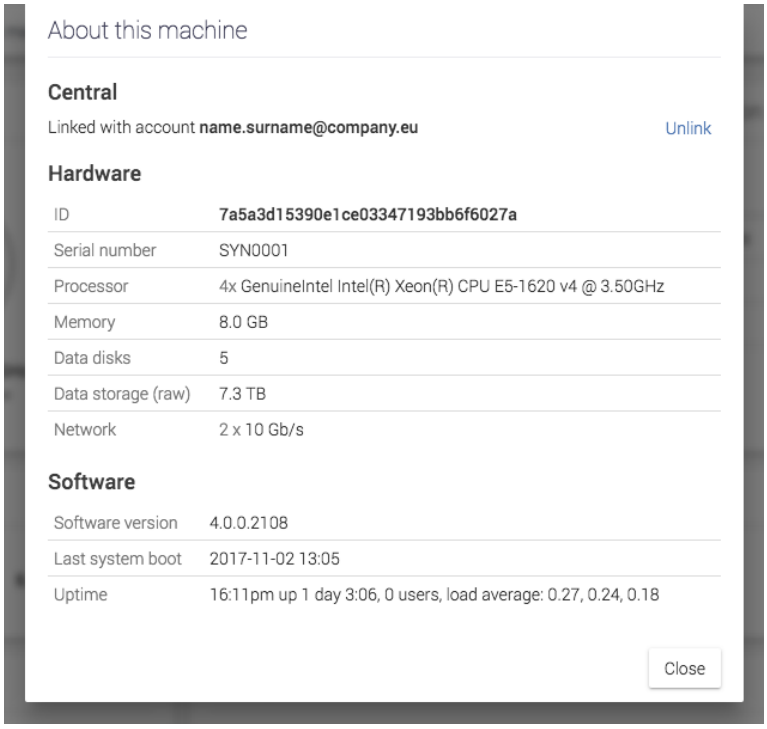

From the *About this machine* dialog, you can also *Unlink* from Central by clicking *Unlink*. A confirmation dialog will pop out saying: *"You will be logged out and need to activate this machine with a new Central account in order to access the UI again."* Should you agree with the terms, the product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

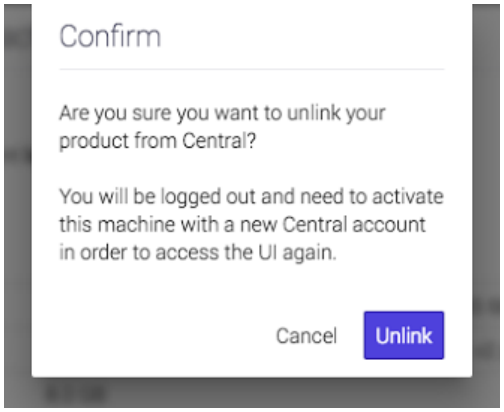

You will be redirected to a form in order to insert the new central account you want the product to be linked to. Complete the new email address and the password. Press *Activate*. You can now log in to your product again.

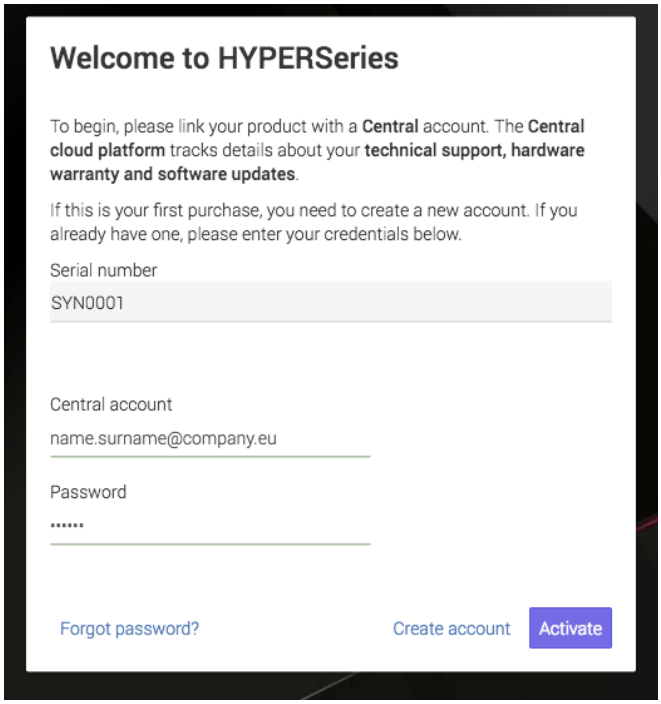

WARNING: If there is no gateway server set, a working DNS configured, or network connection, the Central portal will not be reachable.

DATASTORES

Introduction

Datastores 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 datastores. When you go to *Main menu → Datastores* you will see an empty list.

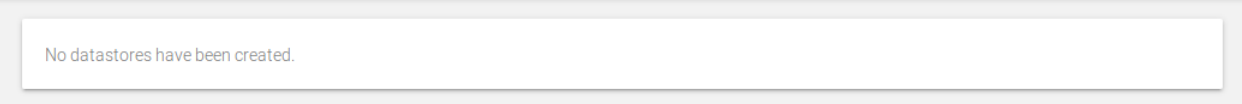

To create a new datastore click *New datastore* in the action bar.

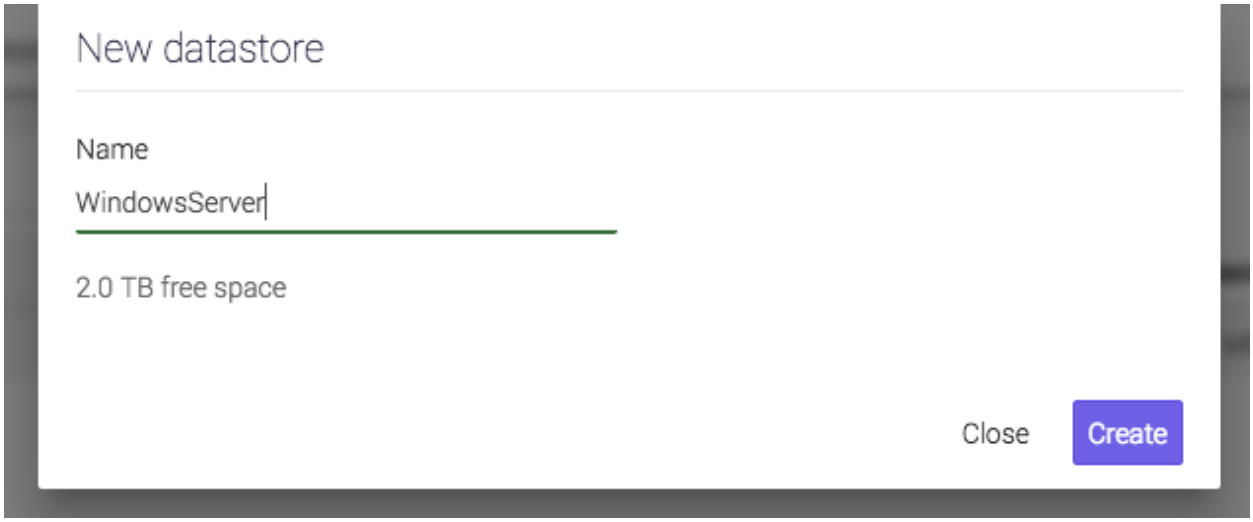

Provide a name for the datastore and click *Create*.

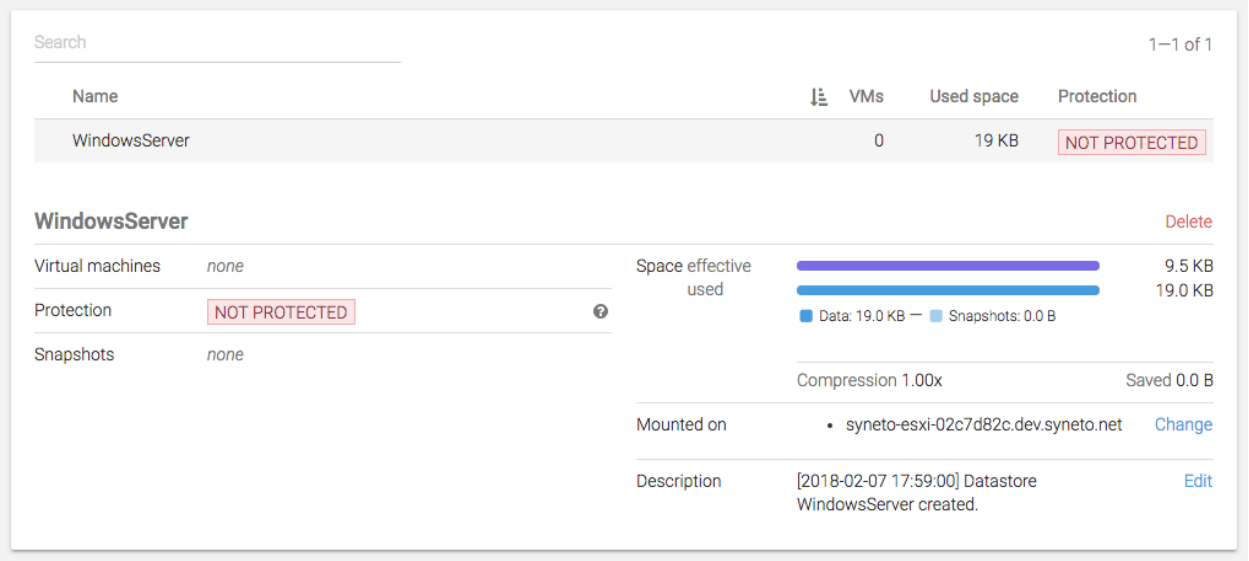

The new datastore will be added to the list, and preselected for you. Under the lists you will see the details about it. In the details section you will be able to perform a set of actions:

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

Provisioning new datastore when multiple pools are available

Your device may come preconfigured with two data pools: hybrid and flash.

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

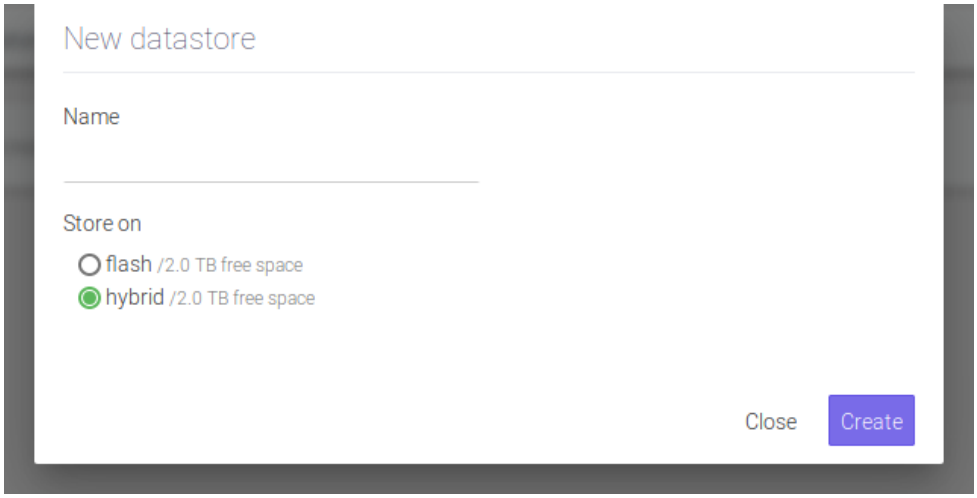

Depending on which pool do you select, you will see the available space. Usually, flash pools are very fast but small. Hybrid pools are slower but provide a large amount of space.

Here are some tips what to 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.

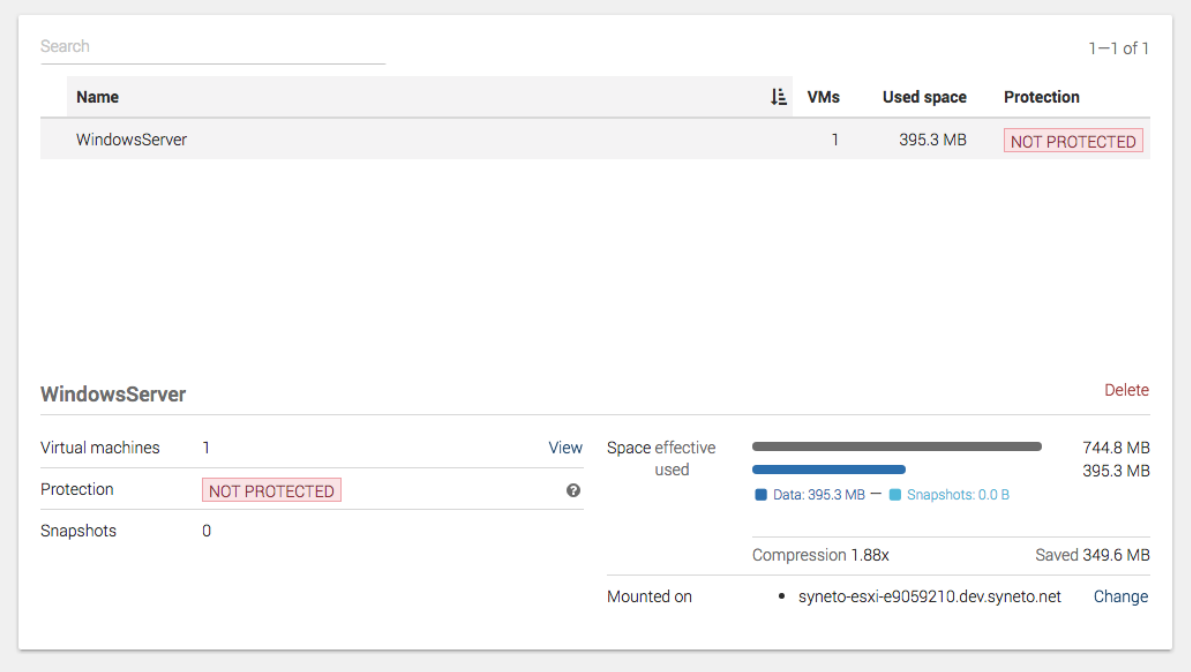

After you click *Delete* a confirmation dialog will explain you that snapshots and virtual machines related to this datastore will be also deleted. You will need to confirm it in order to proceed further.

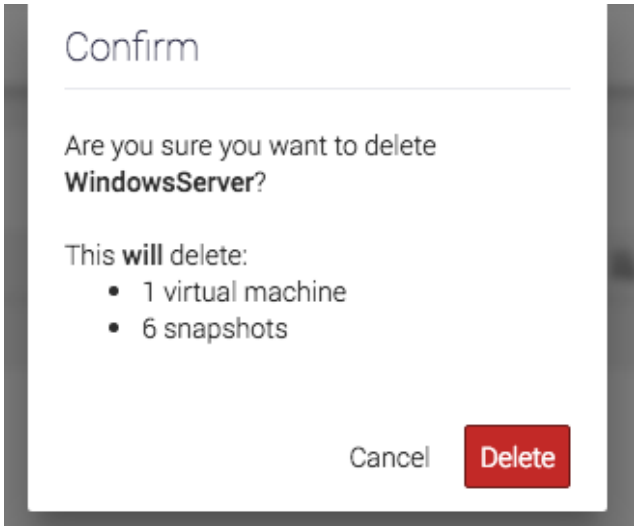

NOTE: Only datastores with no clones can be deleted

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.

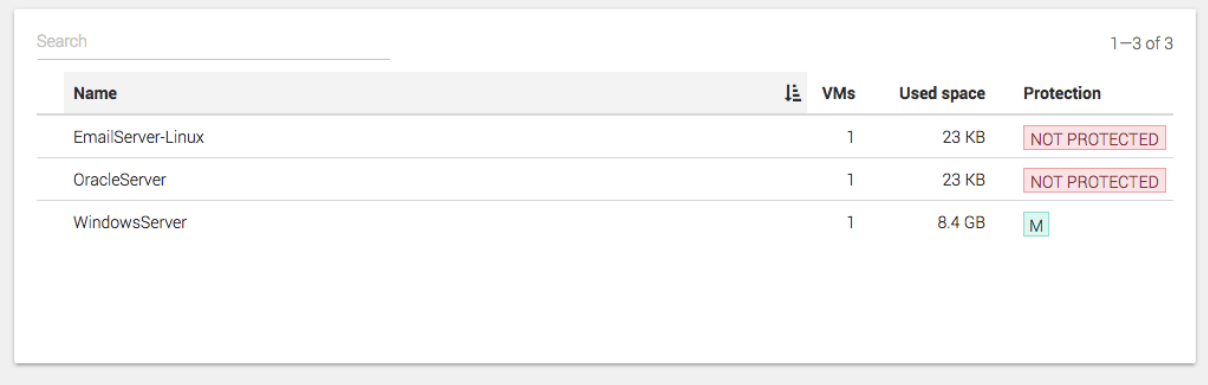

The list contains several columns highlighting the most important information about each datastore: number of virtual machines, used space, protection.

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.

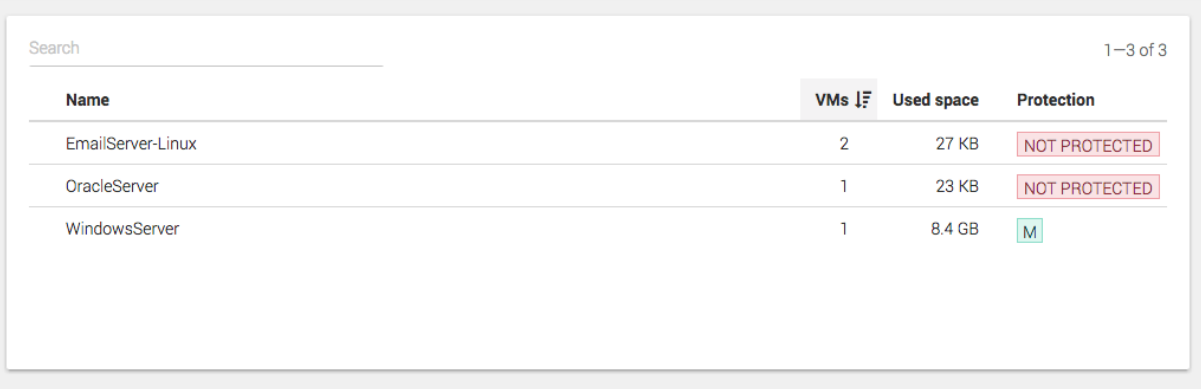

NOTE: We recommend one virtual machine per datastore.

Used space

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

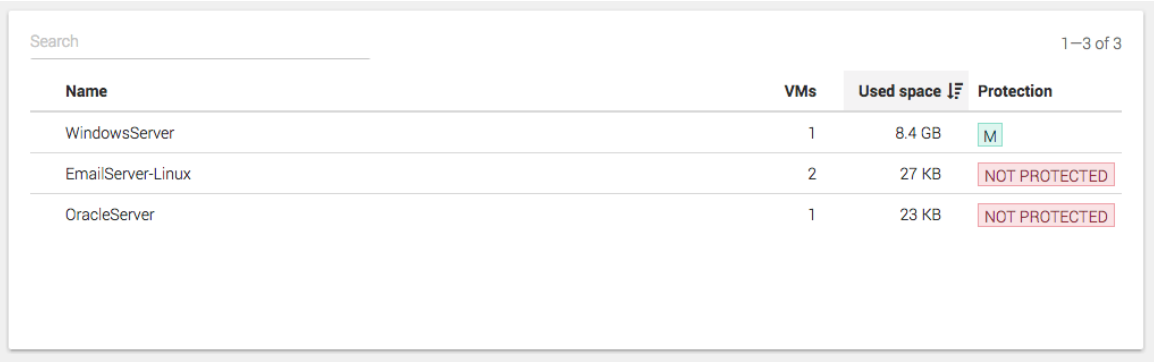

NOTE: This helps you spot virtual machines, or groups of virtual machines that are using up a lot of space.

Recovering data from a snapshot

Get access to a previous version of the datastore

Sometimes, it happens that you lose some data. A virus penetrated you LAN, a user deleted a file by accident, your virtual machine crashed and won't start anymore or a database upgrade went 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.

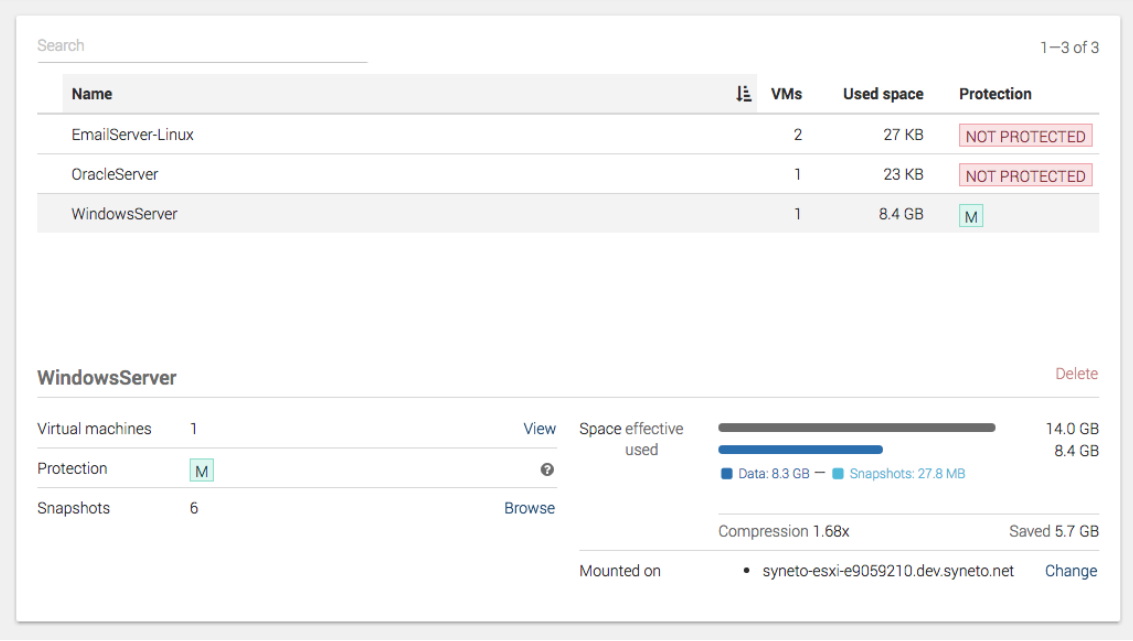

For all these, and many more cases, HYPER clones come in very handy. You can recreate the data exactly as it was at the time of the snapshot using the *Clone* function. To create a clone, start by clicking browse in the snapshots section of the detailed view of a datastore.

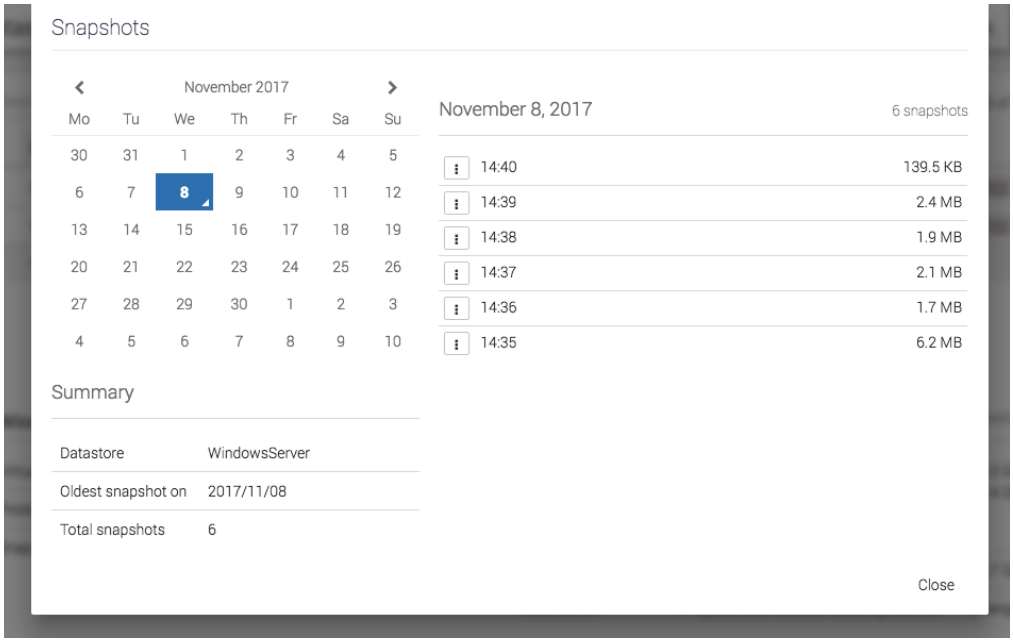

Pick a date from the calendar, and click on the contextual menu for a snapshot.

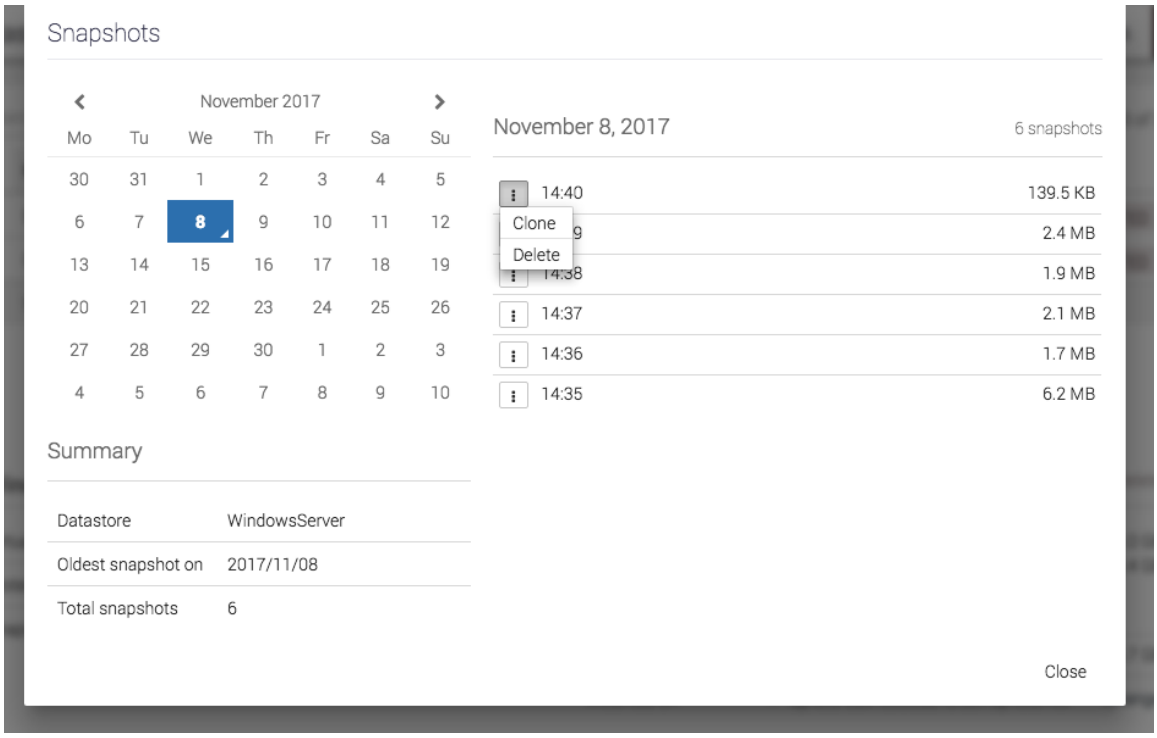

Click *clone* and select the virtual machines you want to recover.

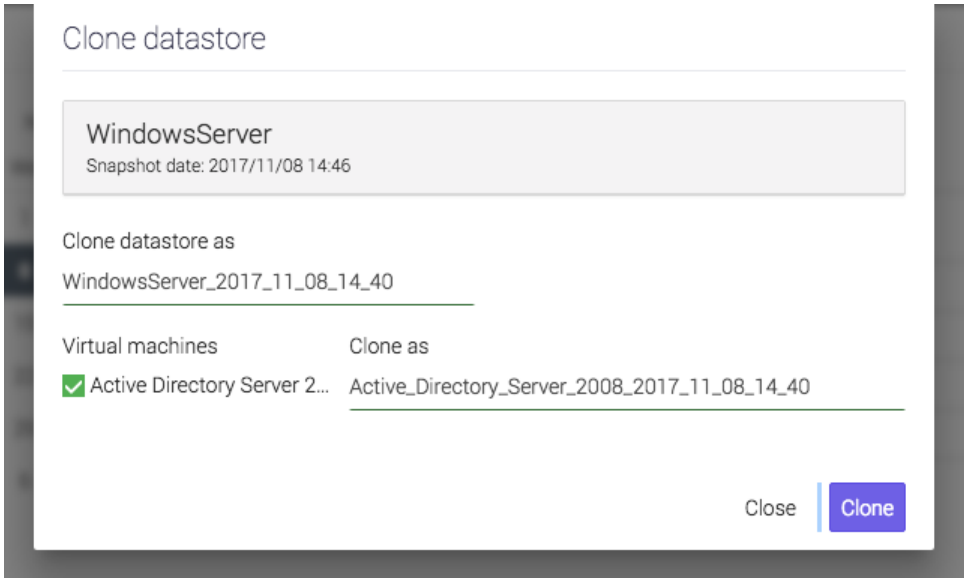

If the snapshot stored more than 1 virtual machine, you can select all of them or just the virtual machines you need. Only the selected VMs will be registered with the hypervisor.

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

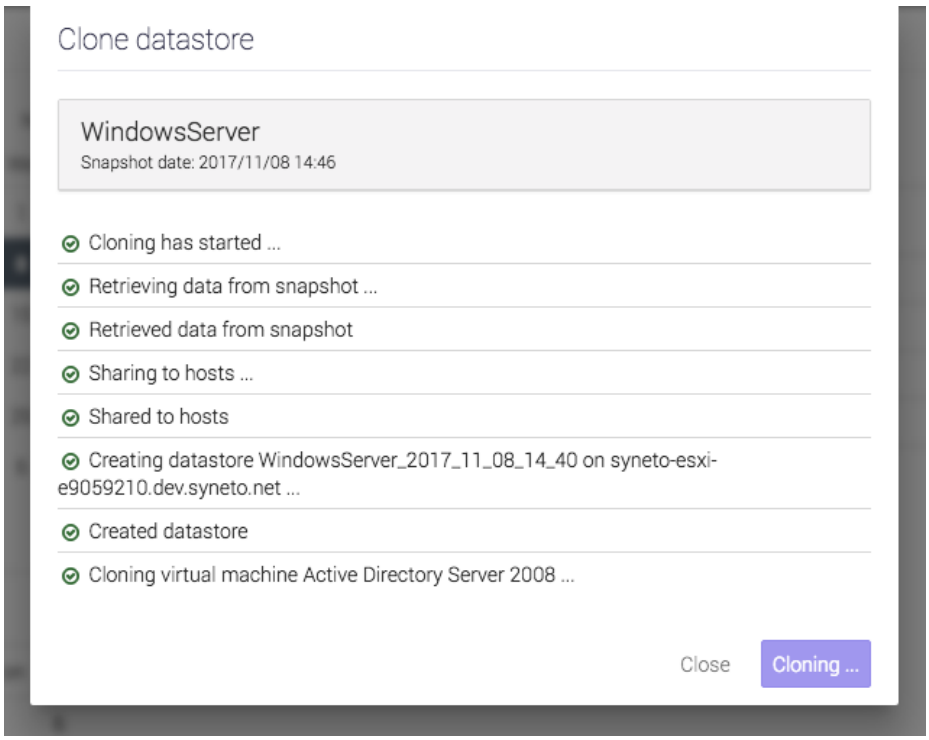

Connect to your vSphere configuration interface and note that a new virtual machine was created for you.

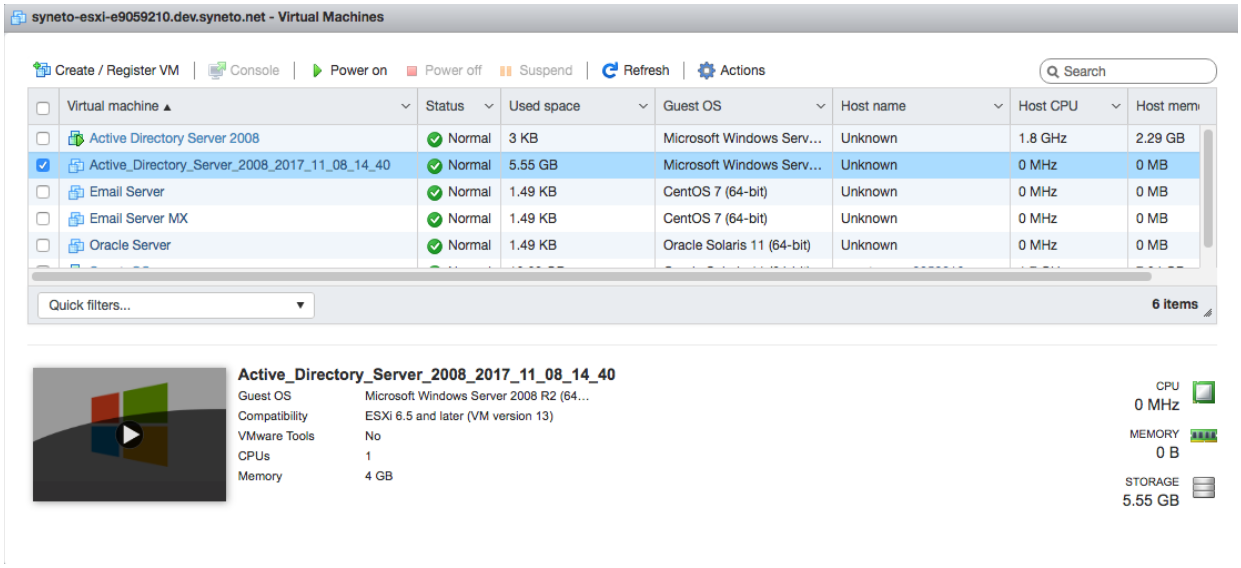

From here you can start the virtual machine and recover/copy the data you need.

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.

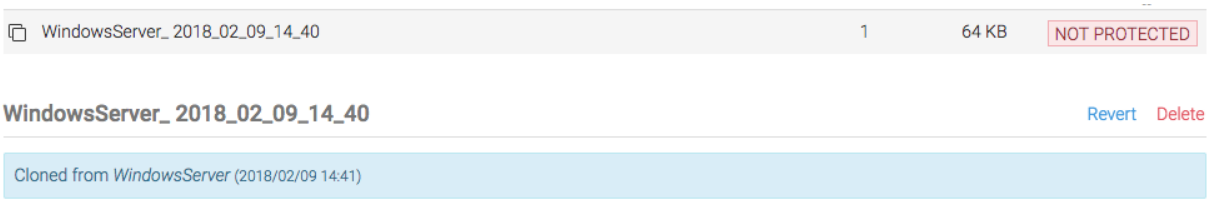

It is also possible to see all the clones of a datastore by selecting the original datastore and in the details view there will be a list of all the clones that it has.

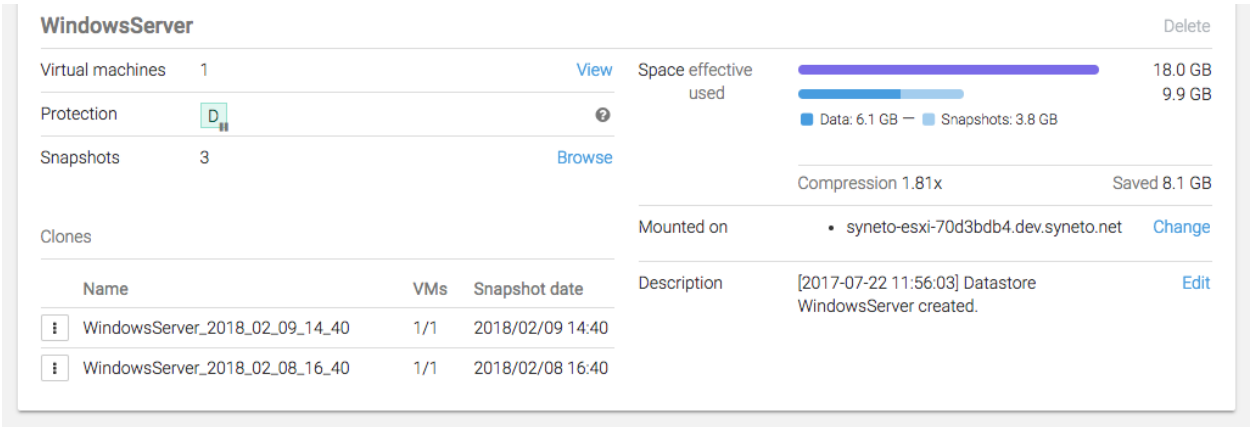

Clean up when you are done

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*.

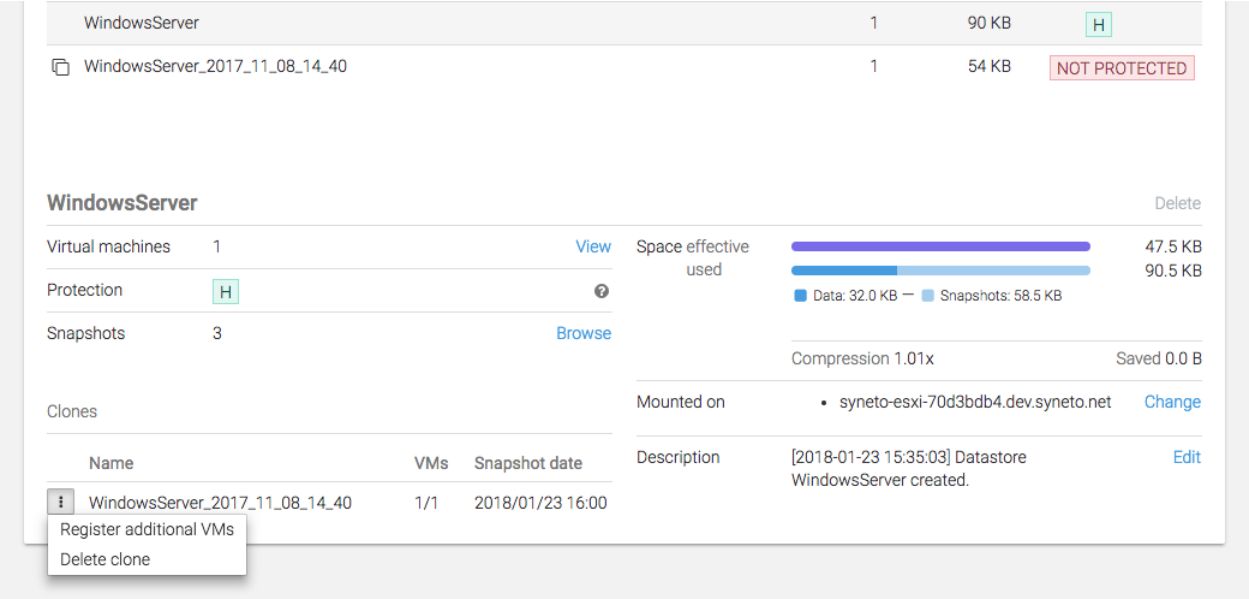

Or select the cloned datastore and click *Delete*.

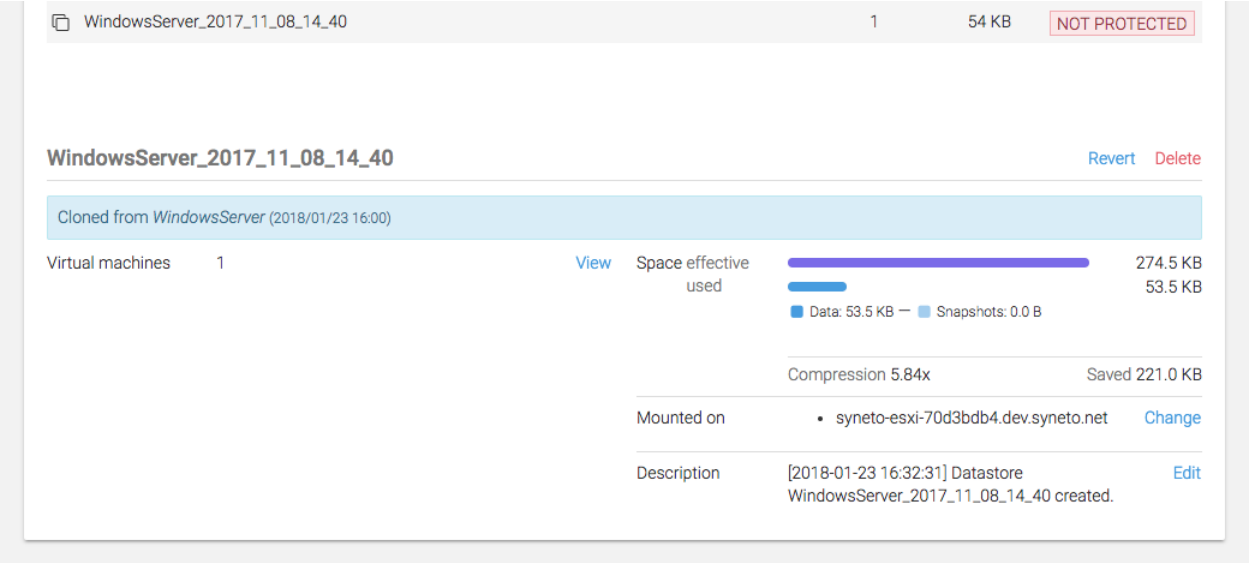

It will ask you to confirm the process and it will explain all the things that will happen, such as virtual machines will be removed, datastore unmounted from VMware, etc.

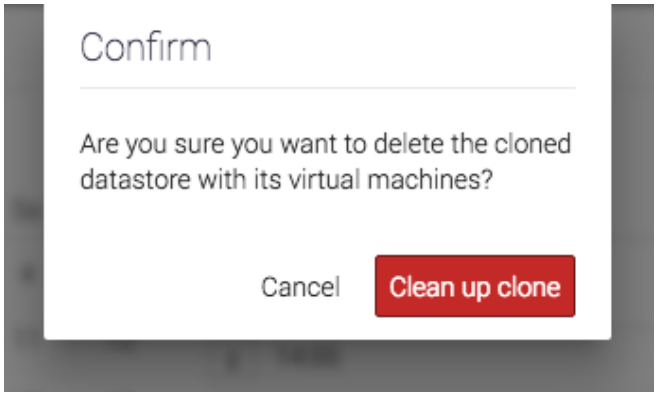

The process takes a little time, you will see a step by step progress.

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

Reverting a clone

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

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

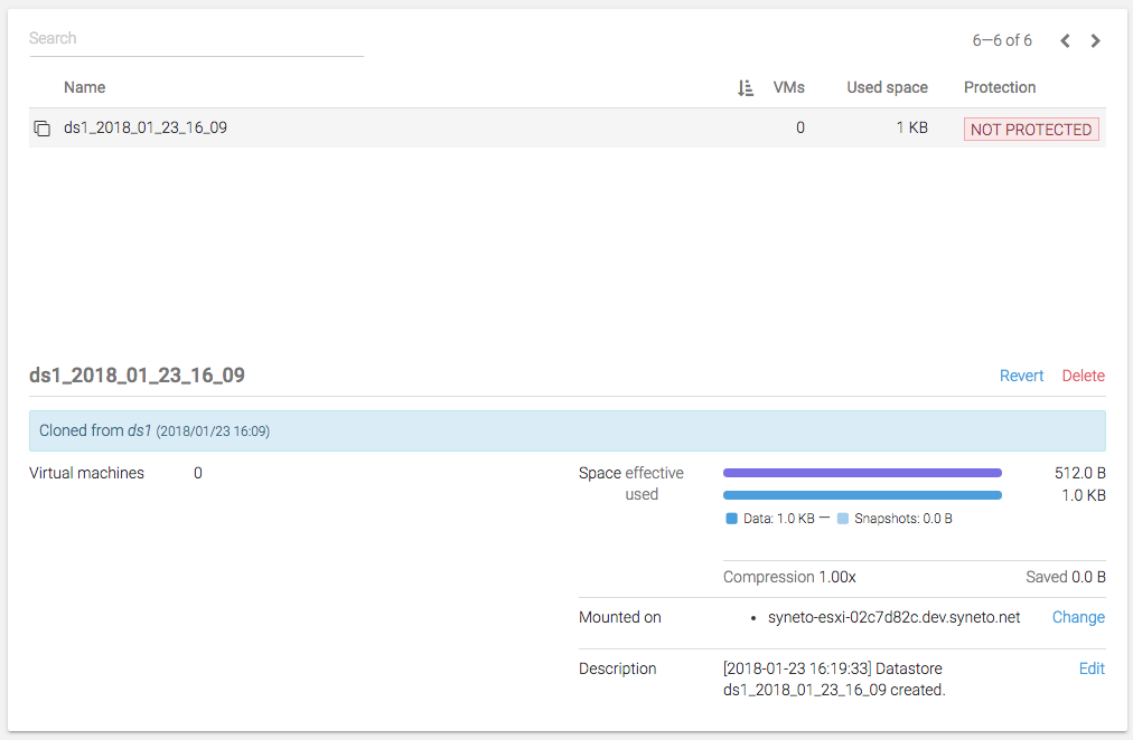

A clone revert process will do the following steps:

- Force stop and unregister all virtual machines from the datastore
- Unmount the original datastore
- Delete all snapshots created after the snapshot from which the clone was created
- Replace the original datastore with the clone both on SynetoOS and ESXi
- 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.

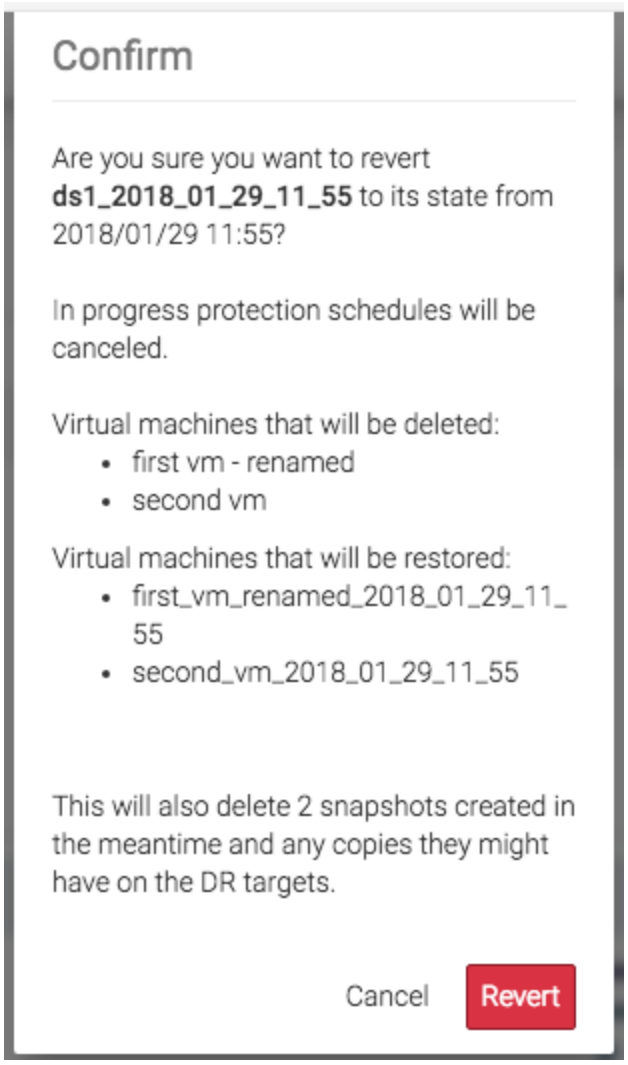

During the revert process you will see a step-by-step progress. If something goes wrong, an appropriate message will be shown.

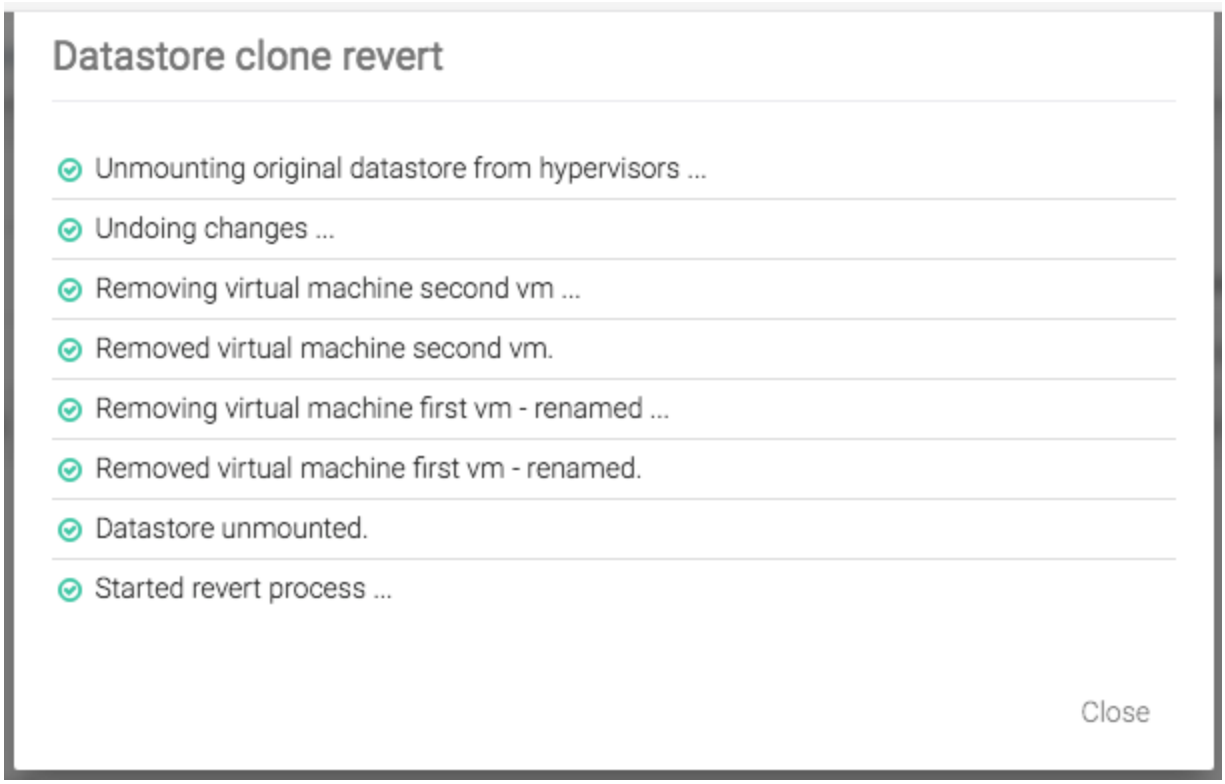

WARNING: On DR unit, on a received datastore, you can create clones but you cannot revert the clone.

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

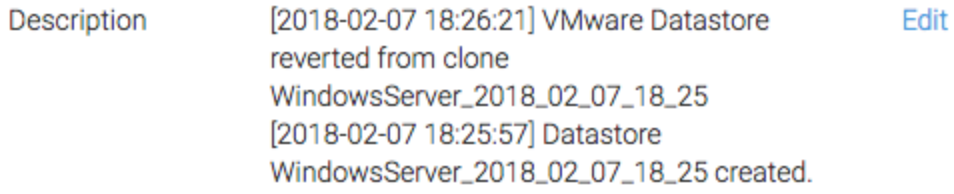

Malware / Ransomware use case

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 forth 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 ESXi 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 (not recommended). 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*.

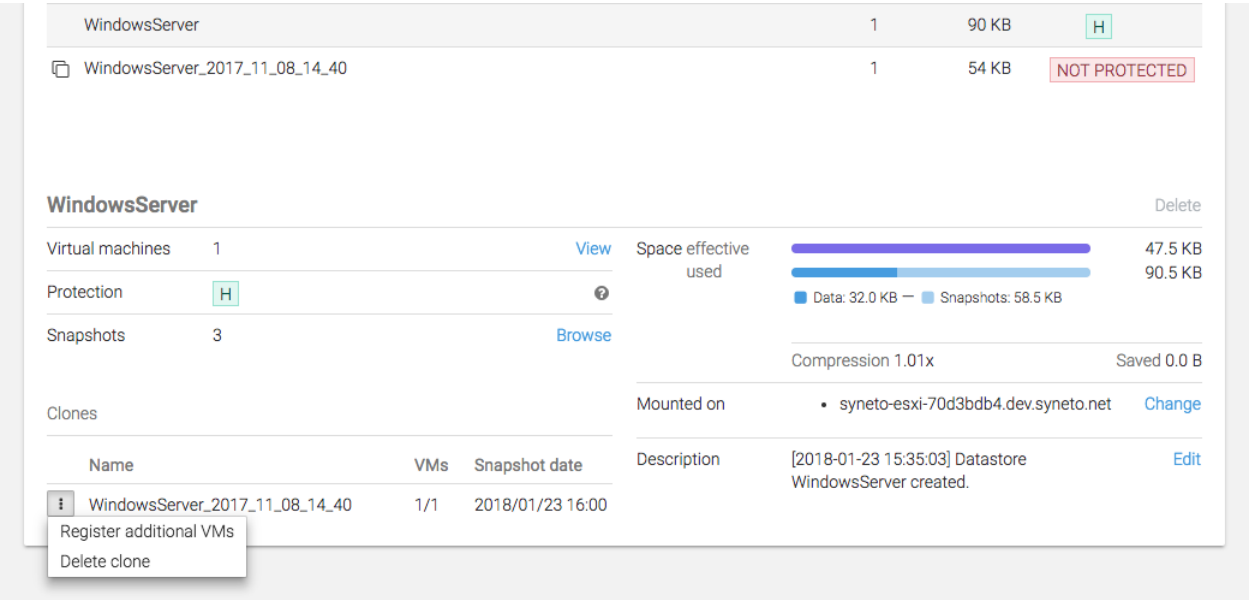

Serving non-Syneto hosts

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 (ESX 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 *Hosts*.

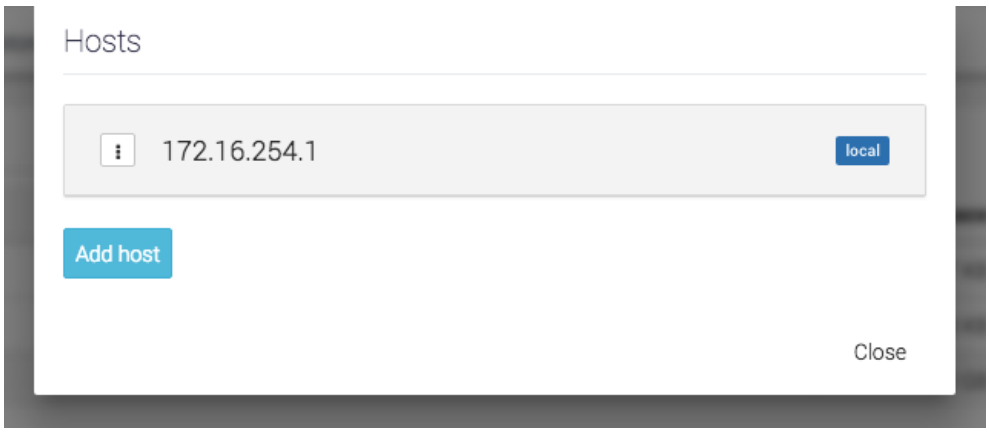

By default the dialog will show you only the local hypervisor.

NOTE: Local hypervisor cannot be removed

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

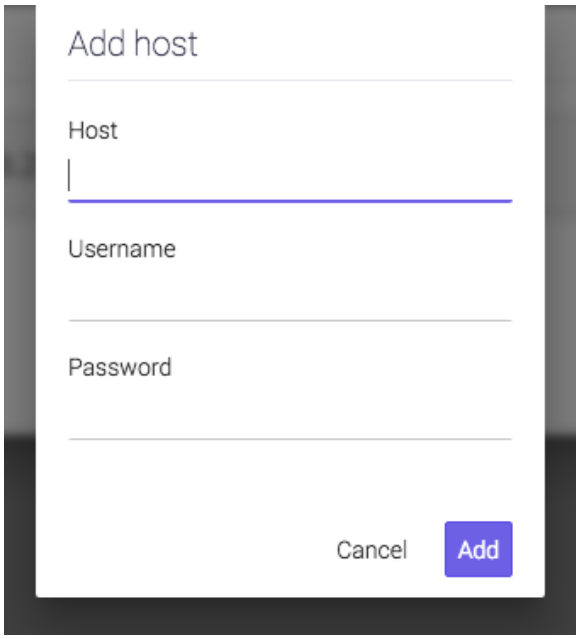

The *Host* field specifies an IP address or a DNS resolvable FQDN (host name).

The *Username* field must be filled with user heaving full VMware administrative privileges.

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

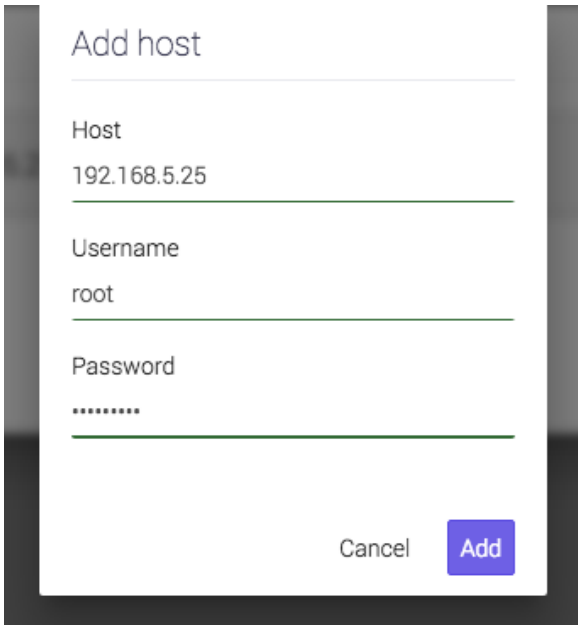

After clicking *Add* the host will appear in the list

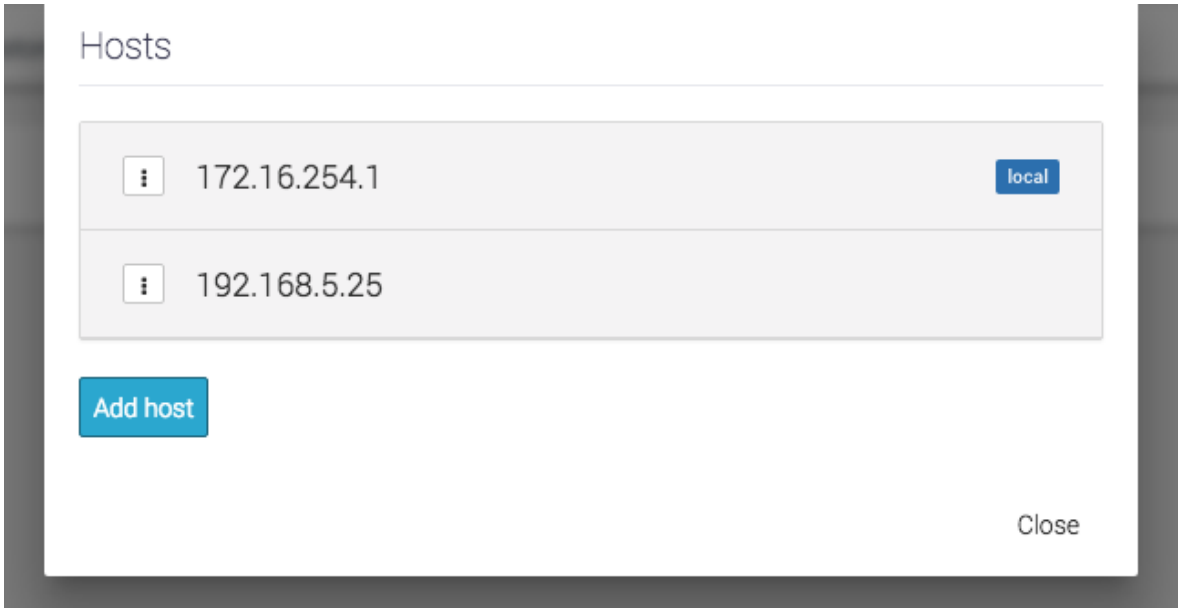

Mounting datastore to additional hosts

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

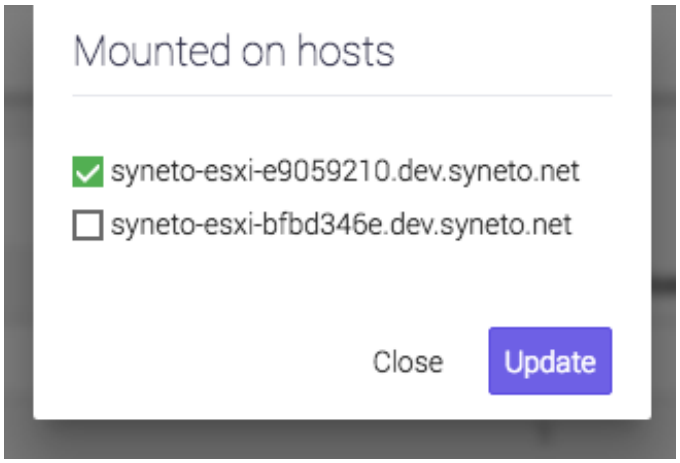

If you added a standalone ESXi to *Hosts*, it will be on the list. You will be able to mount the datastore on this host by checking it and clicking *Update*.

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

NOTE: All ESX 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 ESX 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.

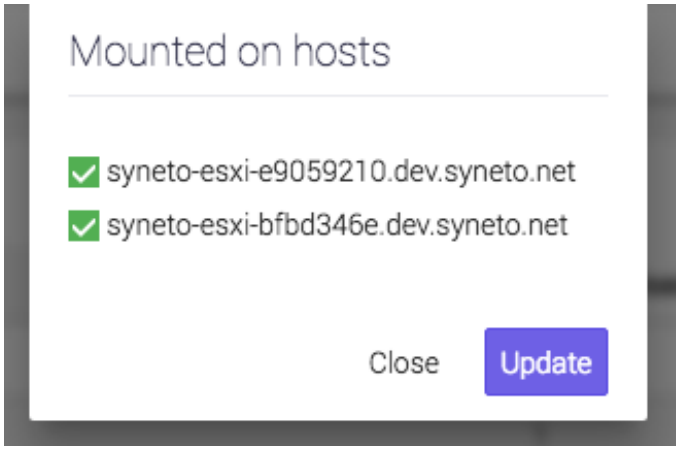

Simply uncheck the hosts you want the datastore to be unmounted from.

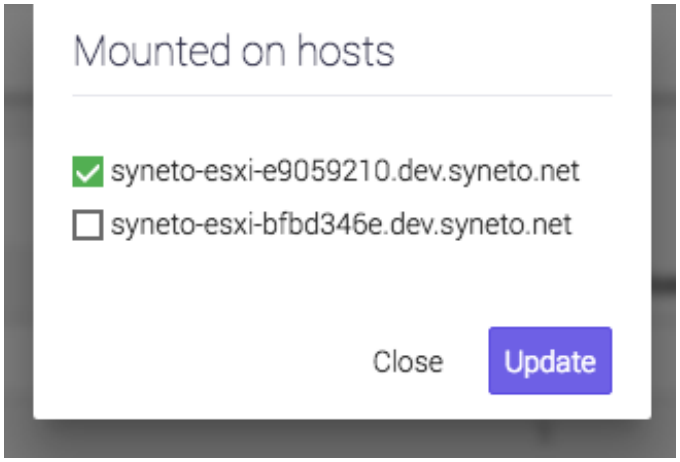

Click *Update*.

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 ESX 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.

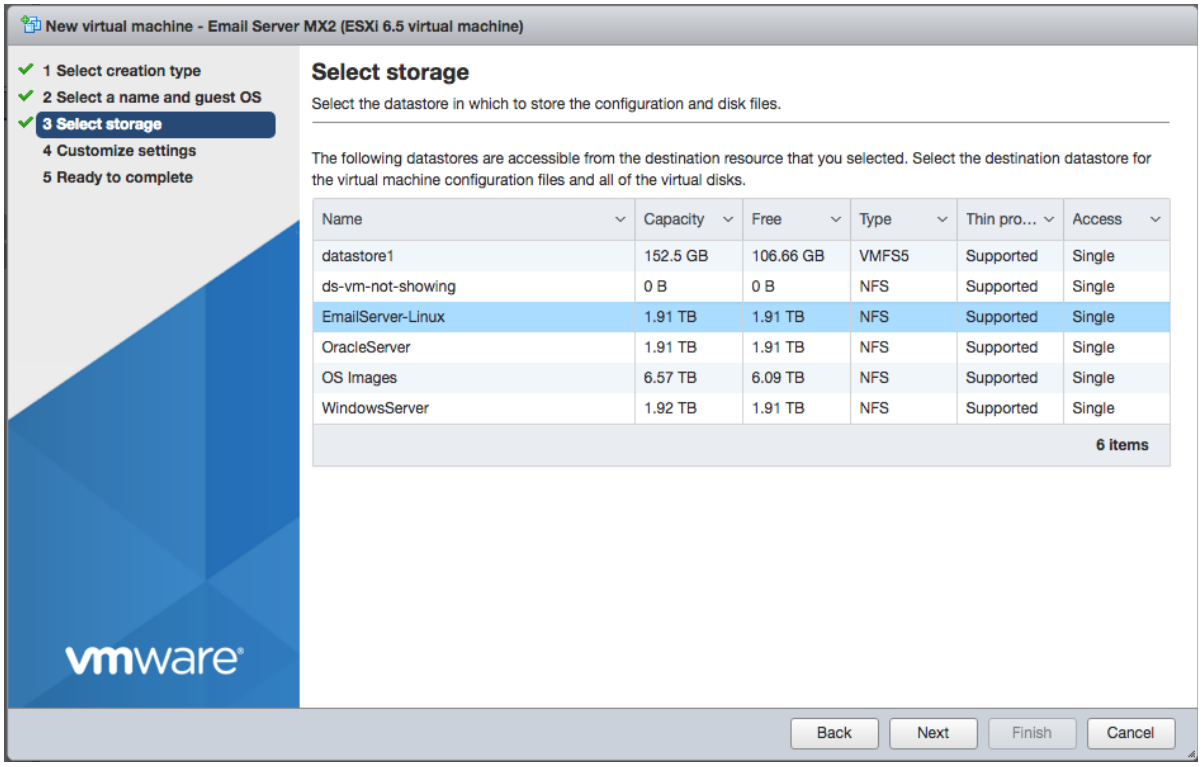

NOTE: For granular protection and performance analyses we recommend creating **one virtual machine per datastore.**

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

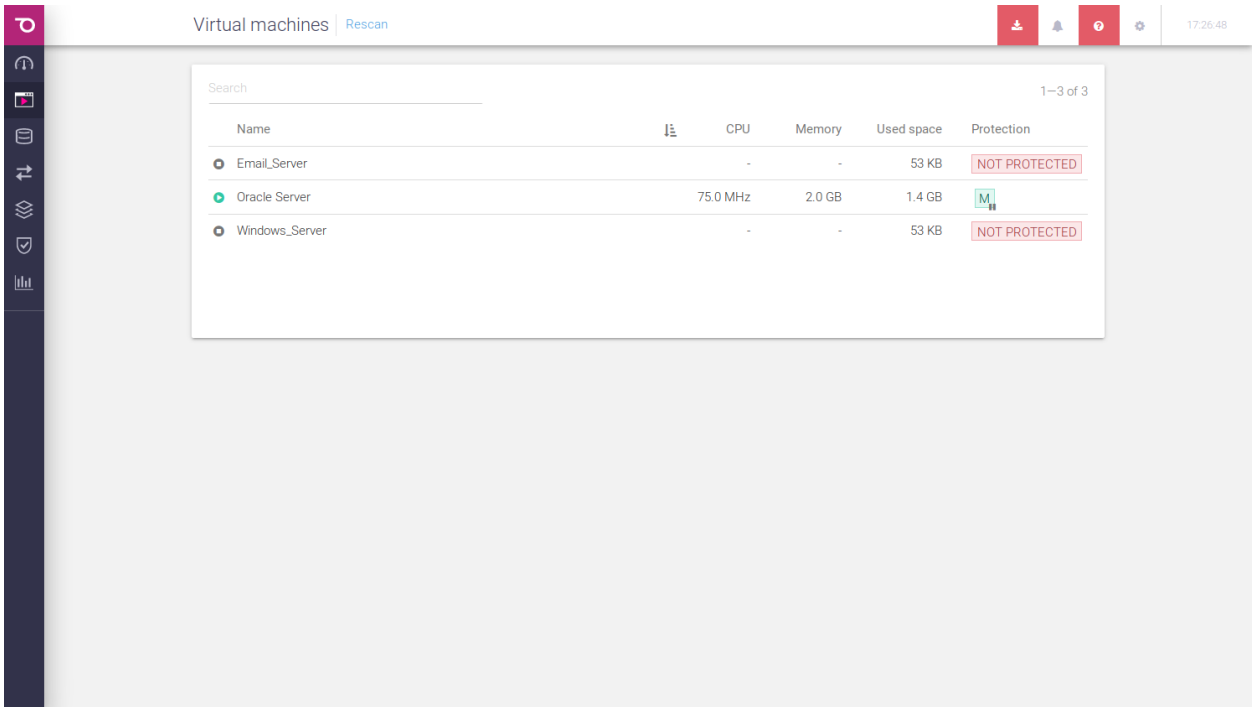

Let's start with a virtual machine that is powered off. After clicking it in the VMs list, its details pane opens. In the top right corner of the details pane, click *Power On*. This kickstarts the boot sequence of the operating system.

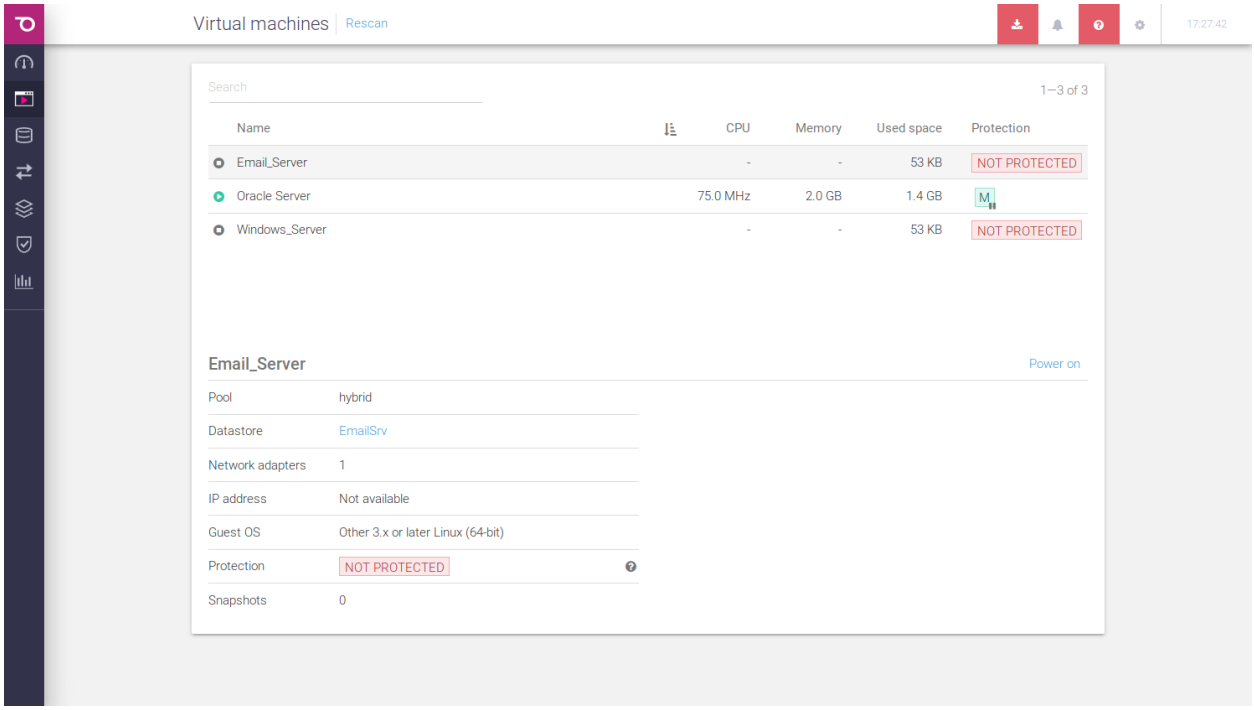

The transition state *Powering on...* will be displayed instead of the button while the machine is starting.

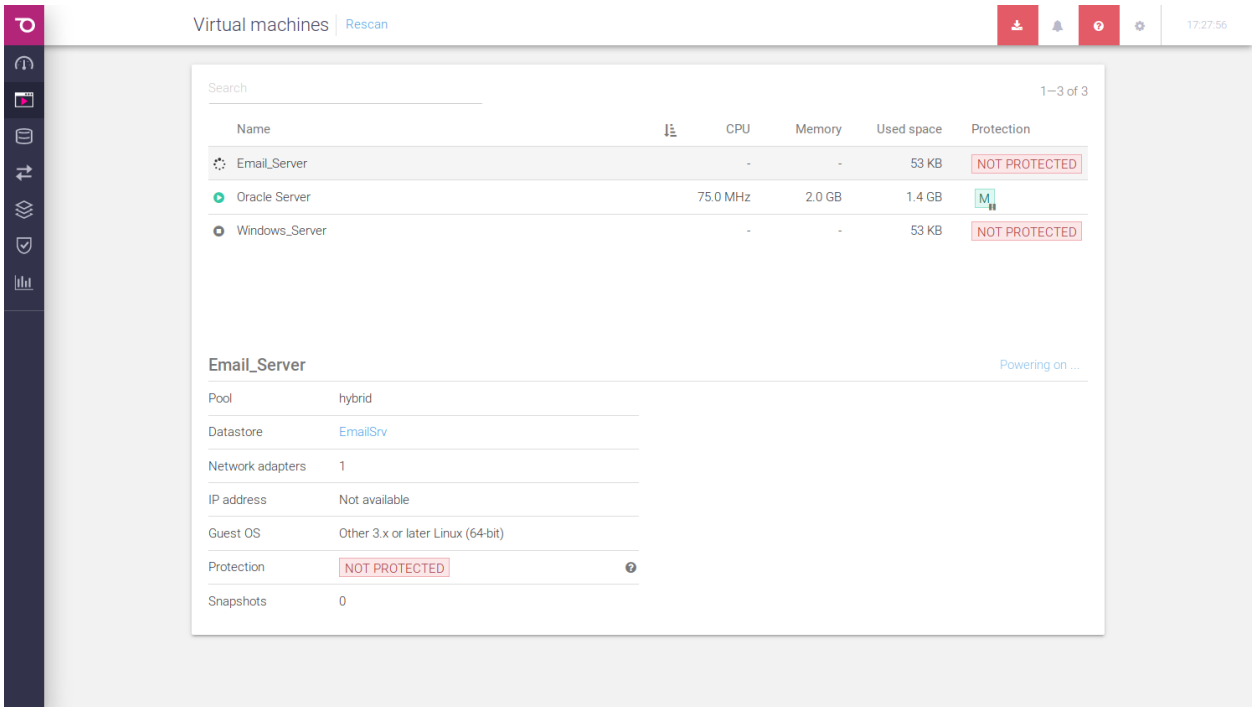

After a VM was powered on, new actions become available: *Power off*, *Suspend* and *Restart*. Also, the virtual machine's console is displayed on screen.

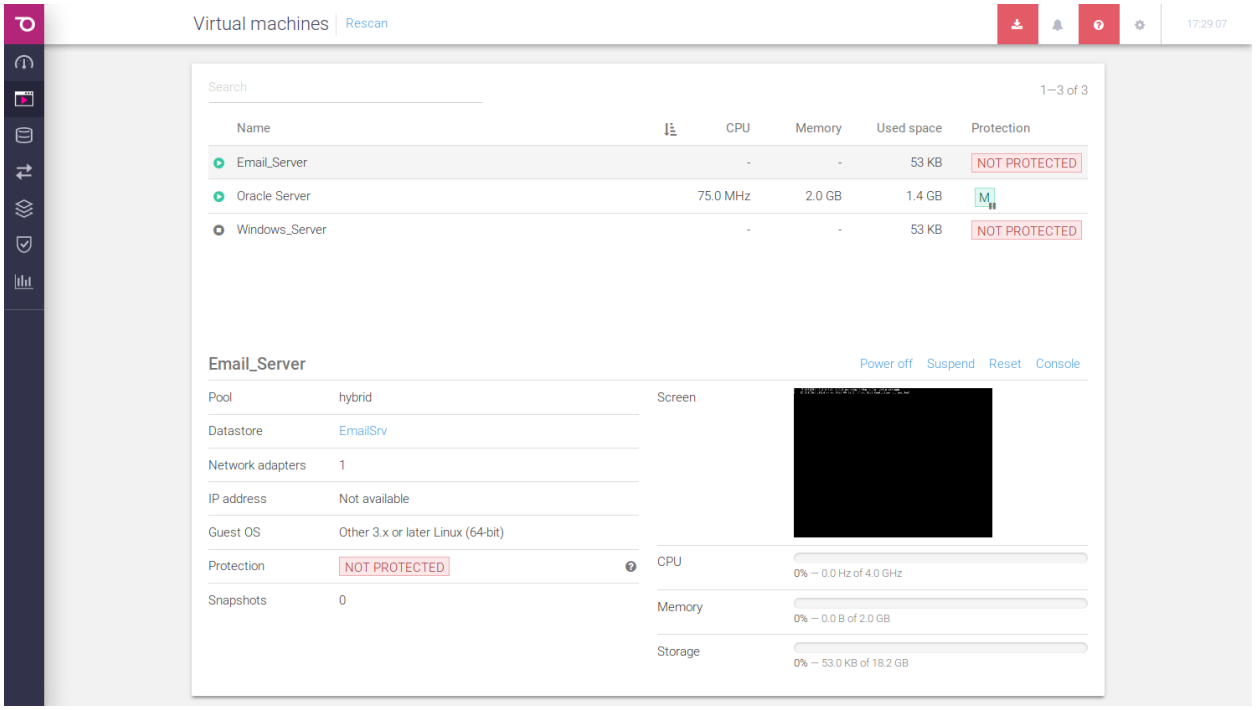

WARNING: *Power off* and *Reset* will perform hard operations, it will behave exactly as if you switched off a physical machine by turning off the power.

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

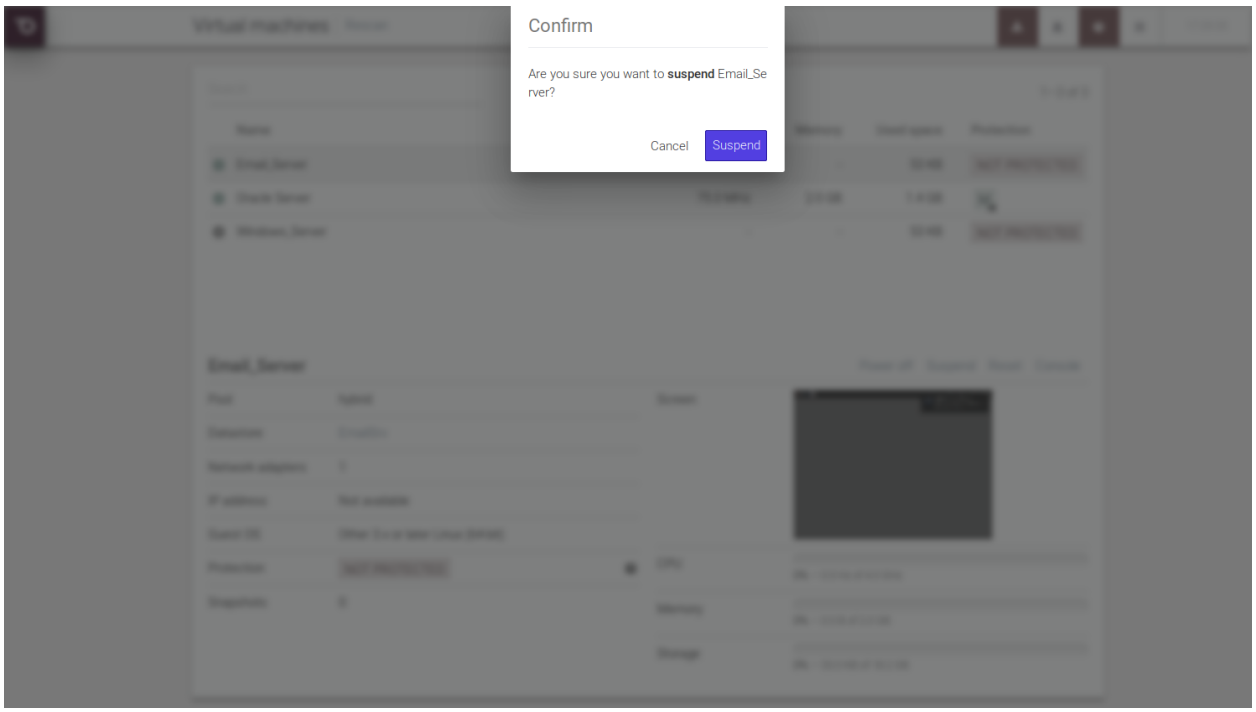

A confirmation dialog will open. Click on *Suspend* if you want to continue.

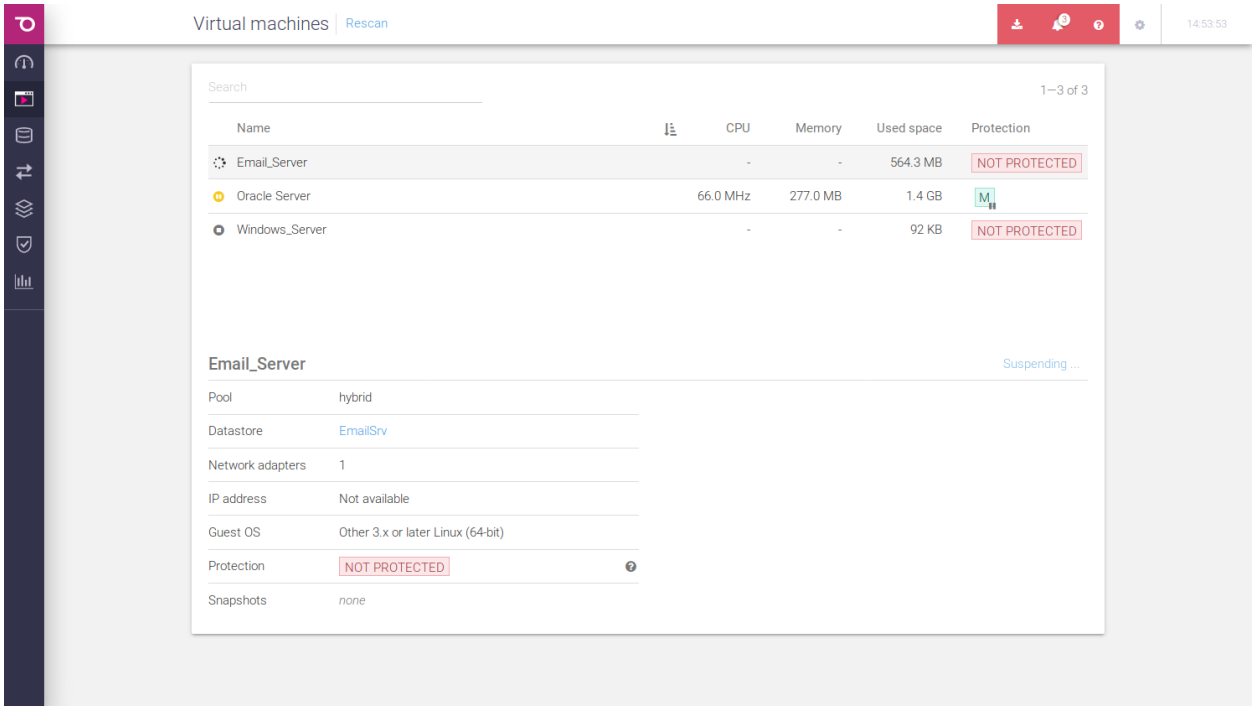

The transition message *Suspending...* will be displayed while the memory is being written to disk, and the other power buttons will disappear. *Power off* is still an option if the virtual machine is suspended or is being suspended.

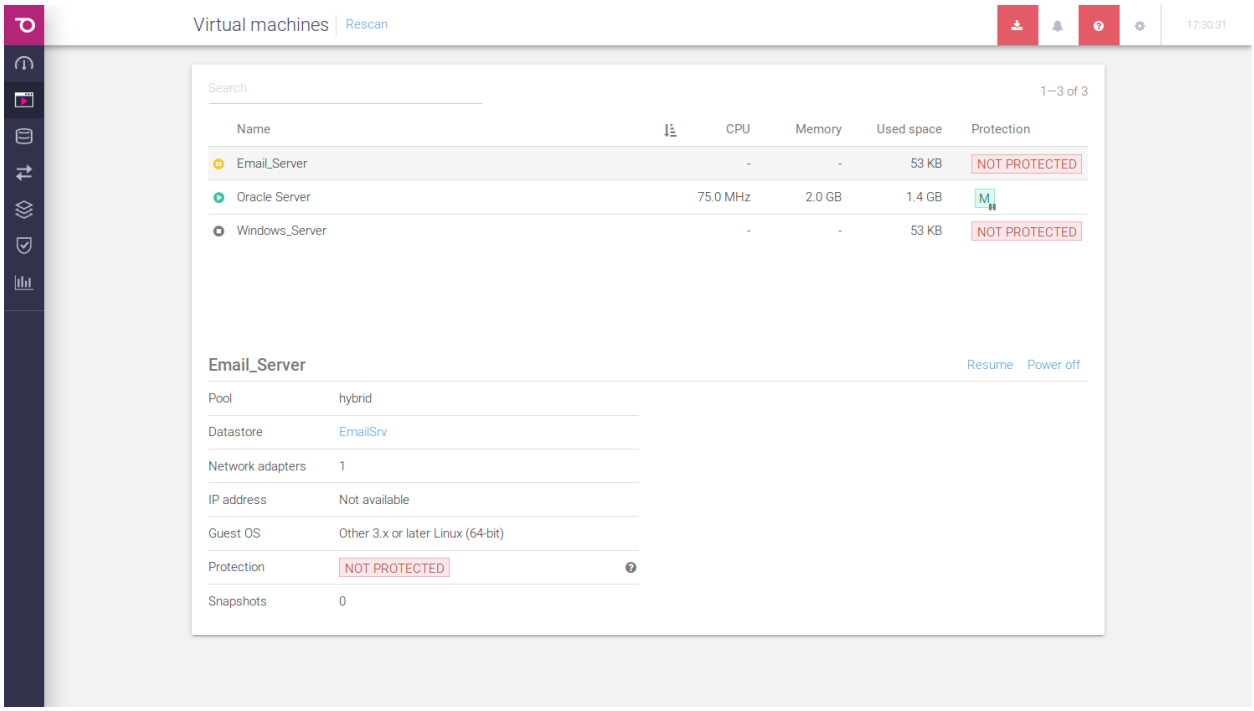

You can *Resume* or *Power off* the suspended virtual machine. Resuming a suspending VM will load it from disk in the state it was before the suspend. No boot sequence is performed.

Console

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

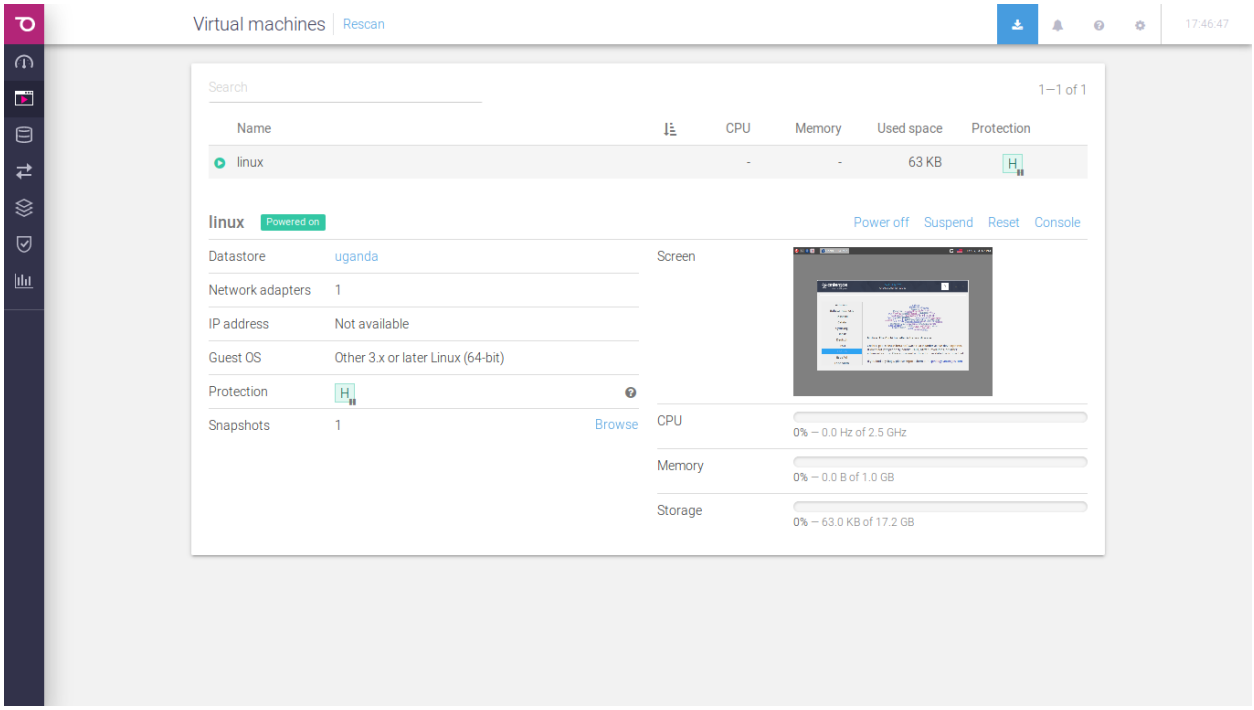

If you want to interact with the virtual machine, click on the image of the console. This will open a new tab where you are able to do actions on your virtual machine.

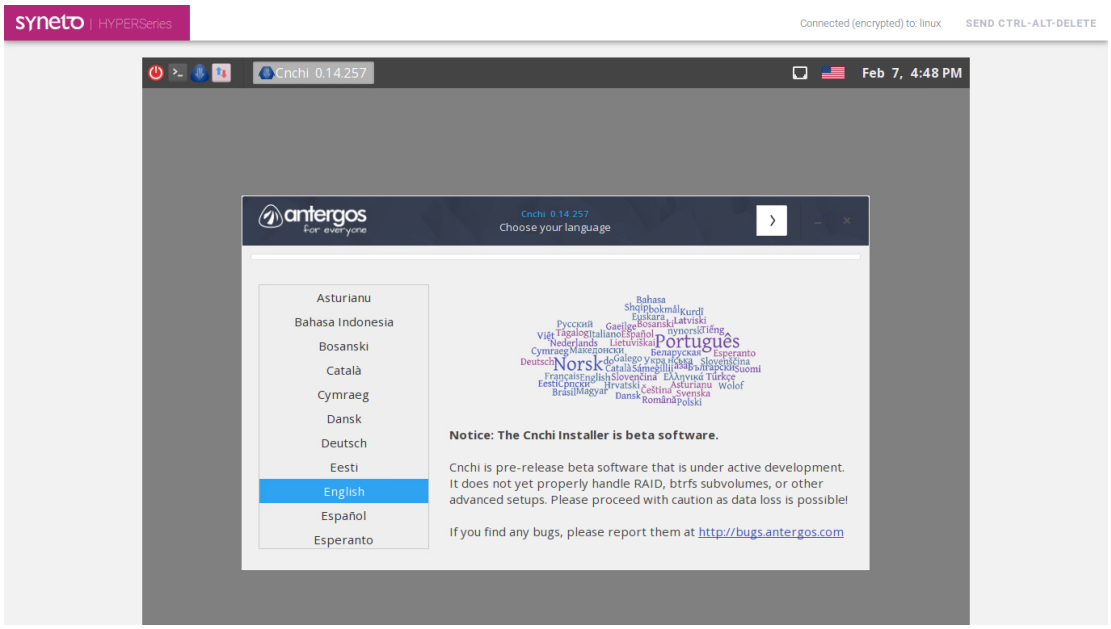

Protecting

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 disk. The snapshot will be taken after the flush operation finishes. Disk consistency for the application supporting the protocol will be consistent. When restoring, the virtual machine will be powered off.
   Recommended for Microsoft servers (Active Directory, MS-SQL, etc)
3. *Live snapshot* - Before taking a snapshot all operations on the virtual machine will be suspended and saved to the disk. This includes waiting for fisk 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 ESX hosts. You will see that only the virtual machines that are on datastores serverved by the Syneto HYPER will be included in the list.

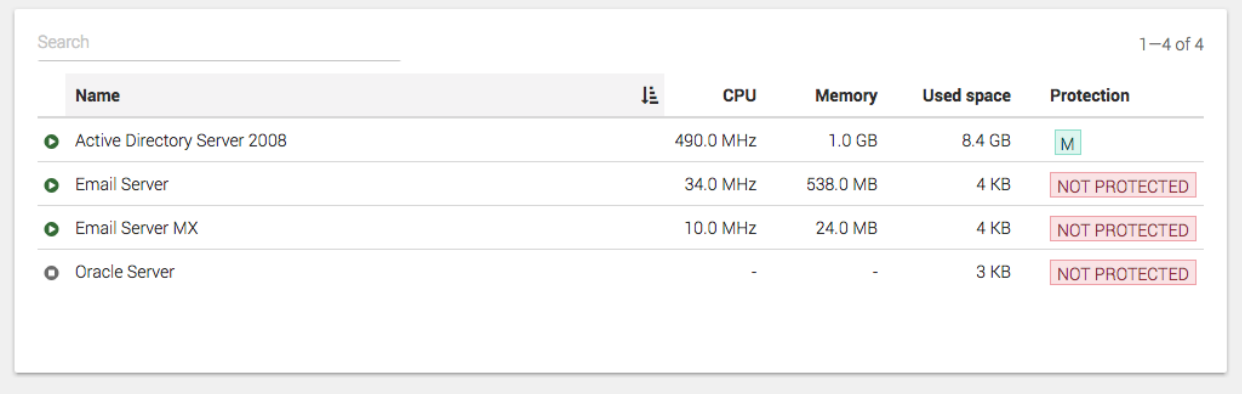

The list has several columns with information about the virtual machine: CPU usage, memory usage, used space, and protection status.

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.

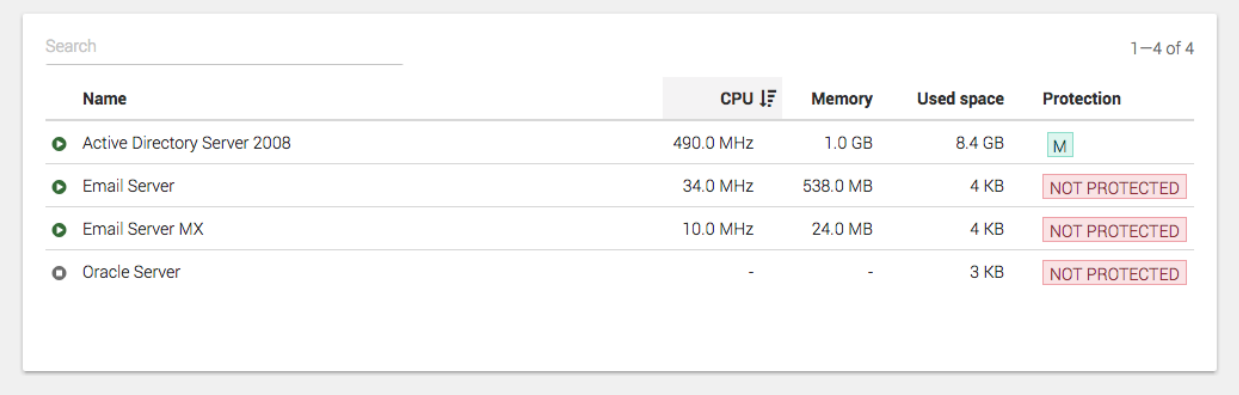

NOTE: it is useful to sort the list descending by *CPU* when you want to spot virtual machines using too much CPU. For example when you observe a virtual machine is very slow, it is a good practice to come here and quickly check its CPU usage.

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.

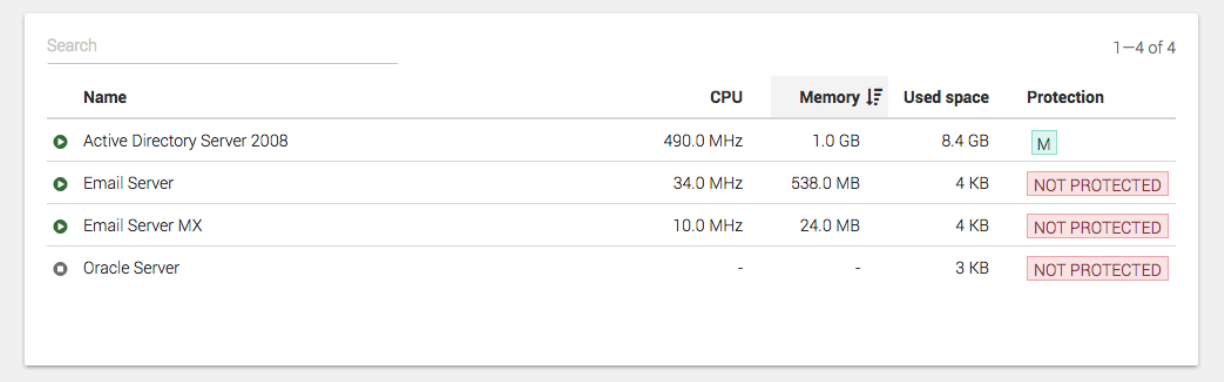

NOTE: it is useful to sort the list descending by *Memory* when you want to spot virtual machines using too much memory. For example when you observe a virtual machine is very slow, it is a good practice to come here and quickly check its memory usage.

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

Used space

This column show how much disk 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 disk usage in GB, maximum disk space available for the virtual machine, and a percentage of the current value.

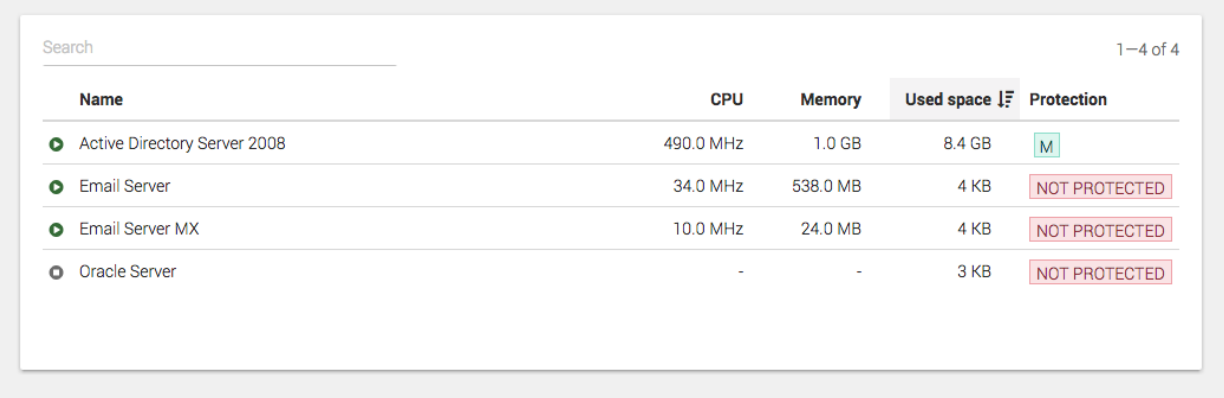

NOTE: it is useful to sort the list descending by *Used space* when you want to spot virtual machines using too much disk space. For example when you observe a virtual machine cannot write new data or it is very slow with disk operation you may want to check this column.

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 Syneto OS 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.

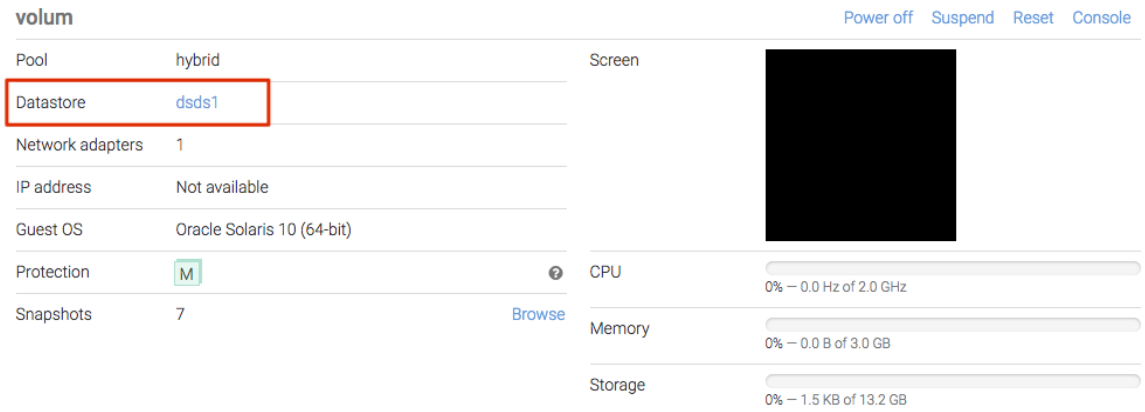

Analyzing performance

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

Dataset IOPS

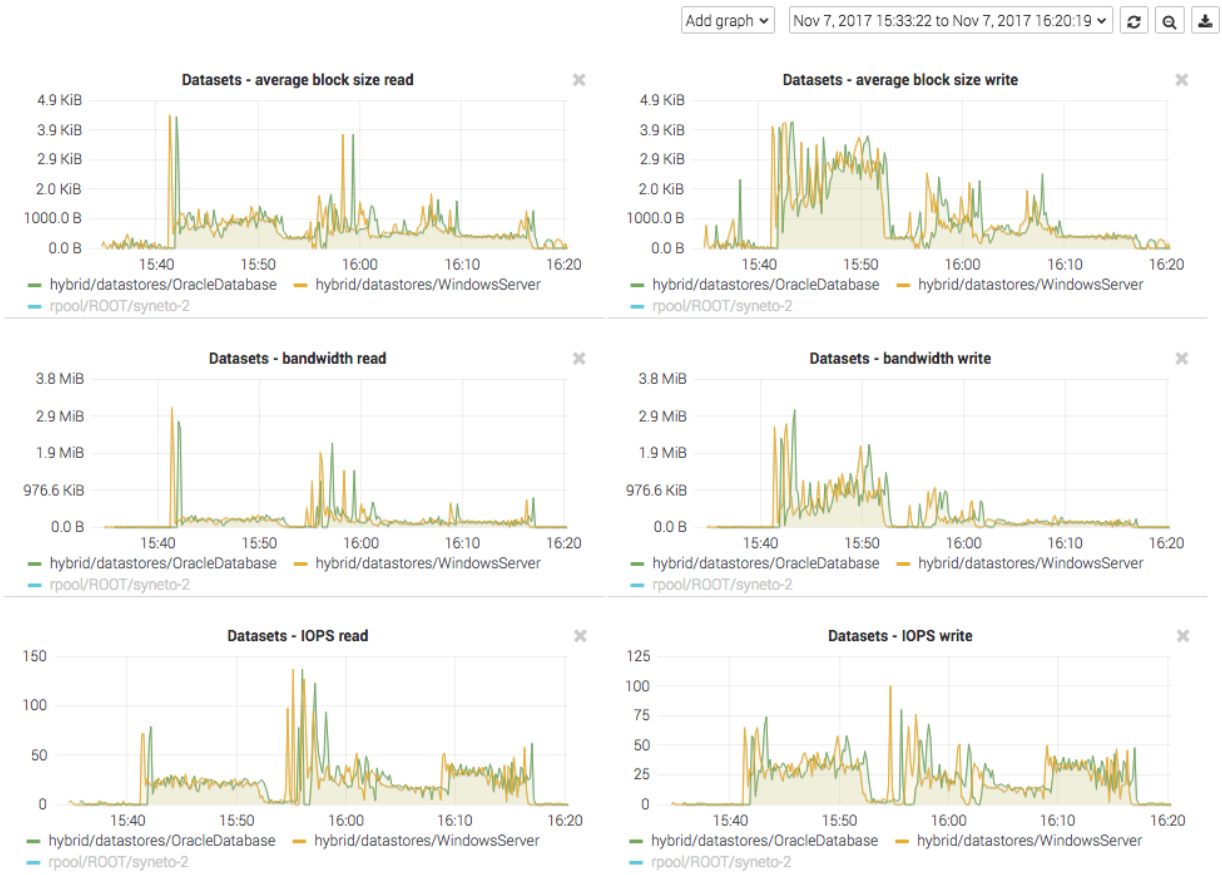

Go to *Main menu* → *Analytics*, click *Add graph*, select *Dataset* → *IOPS*.

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*.

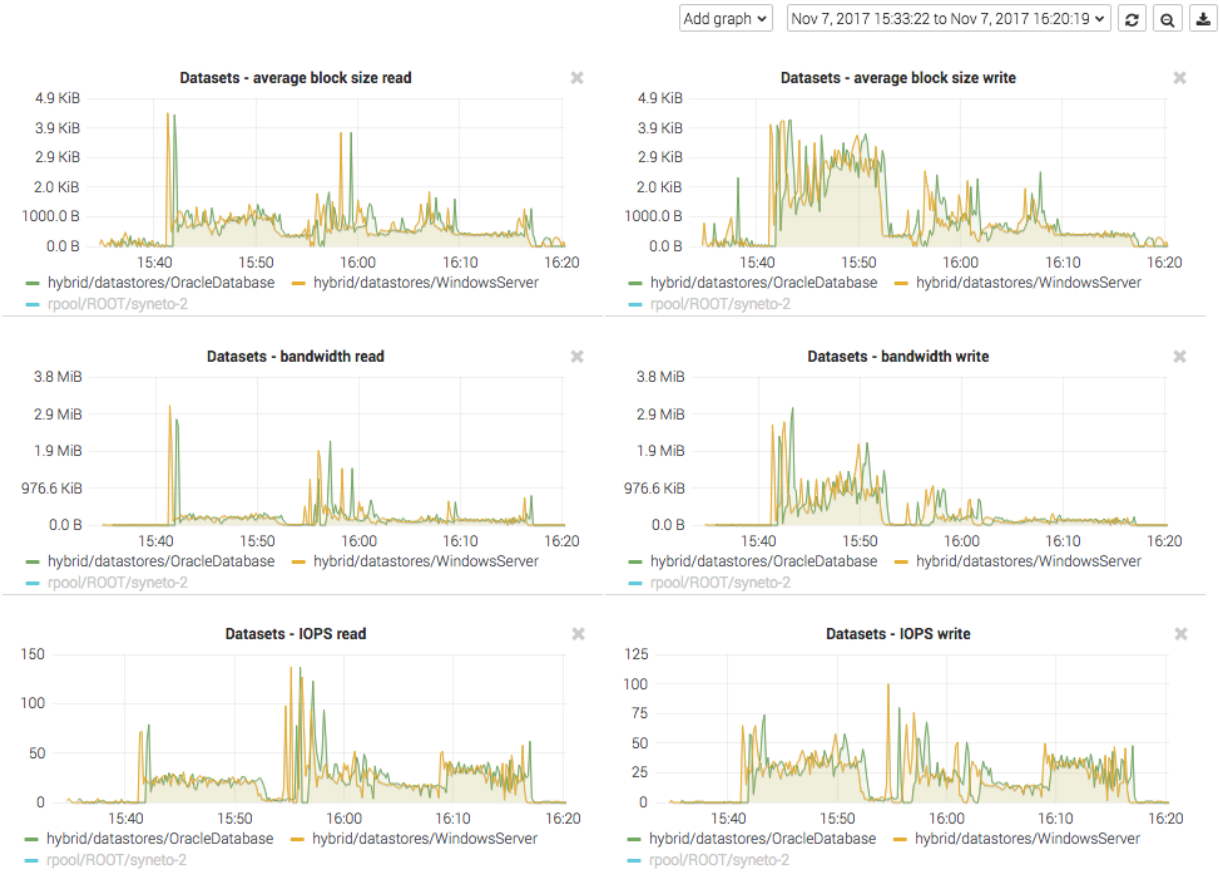

These graphics show the data transfer speed on the datastore from the Syneto HYPER's perspective.

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*.

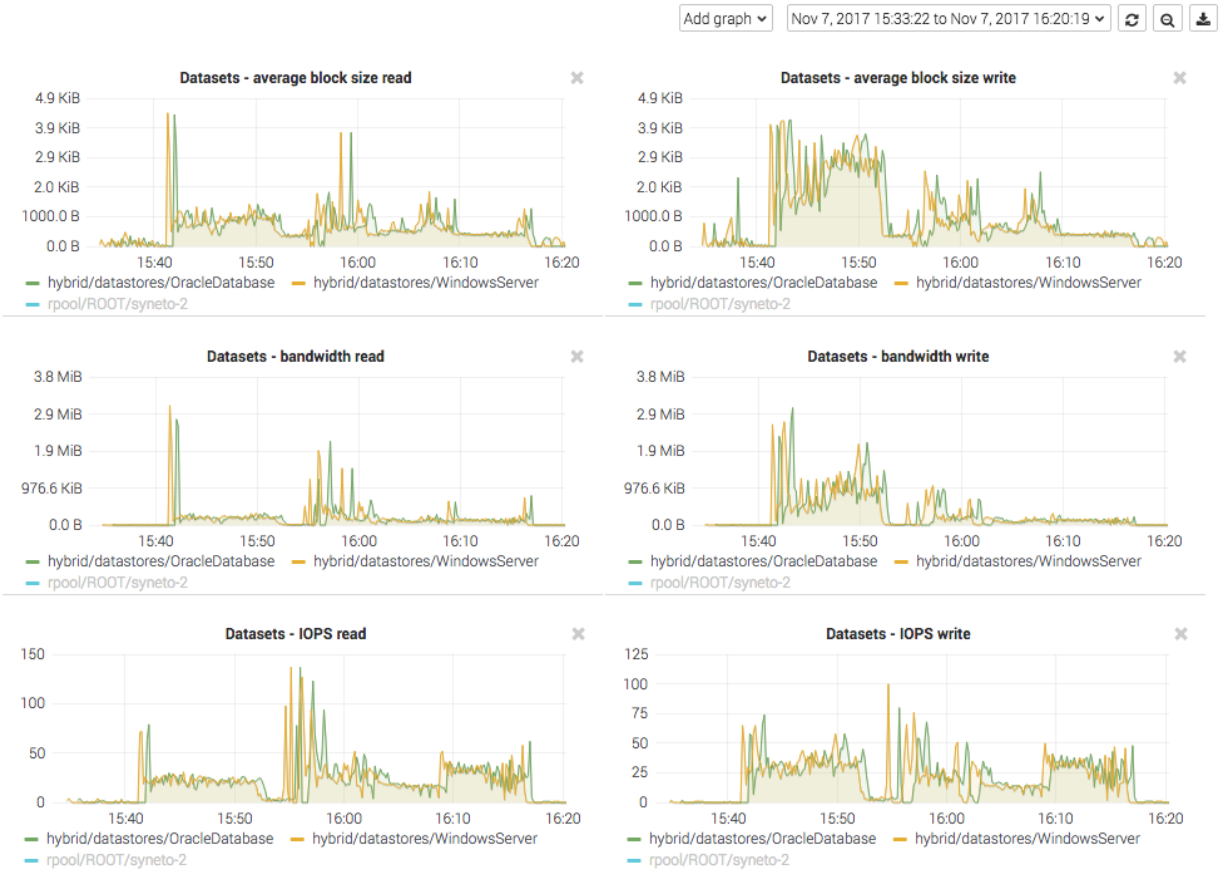

These graphics show the average block size of the data on the datastore from the Syneto HYPER's perspective.

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.

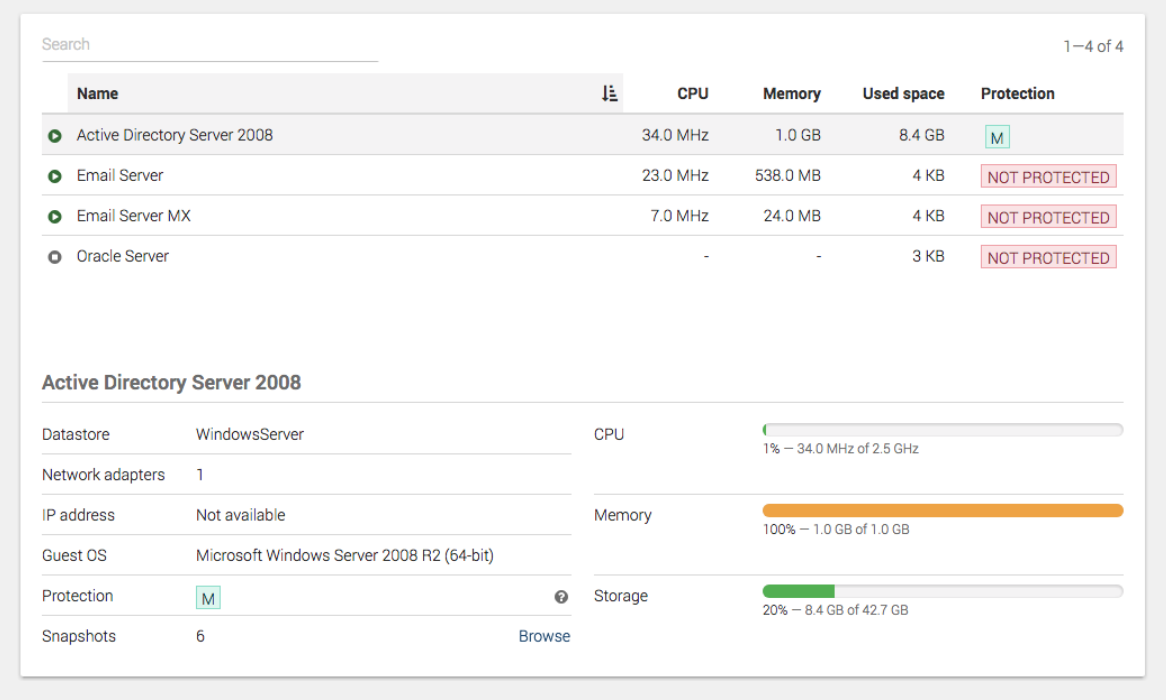

For all these, and many more cases, you can simply click *Browse* in the snapshots section of the detailed view of a virtual machine.

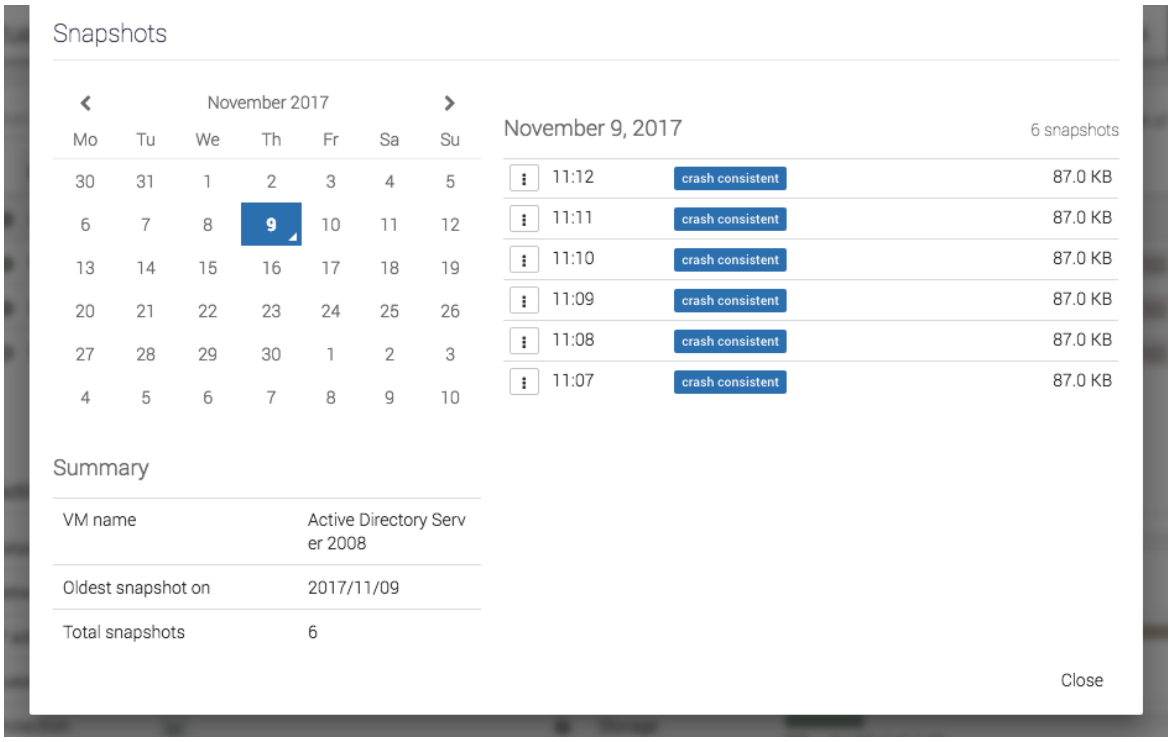

Pick a date from the calendar, and click on the contextual menu for a snapshot.

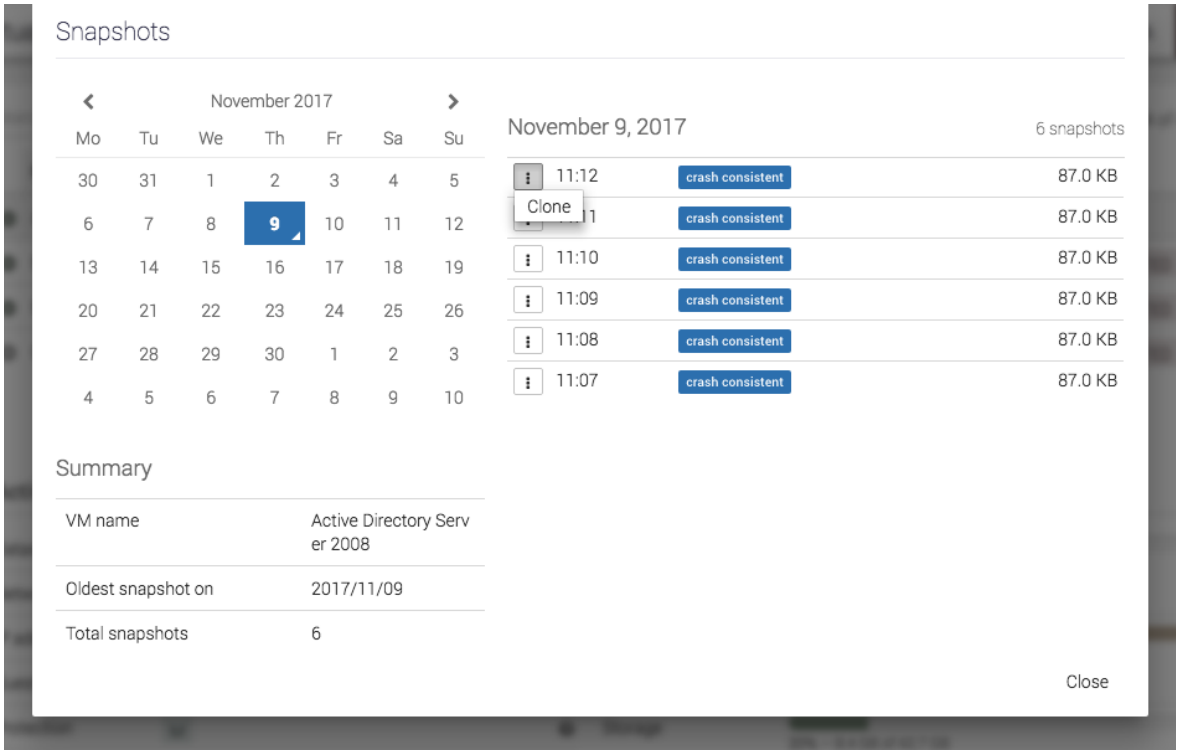

Click clone. The virtual machines will be preselected for you.

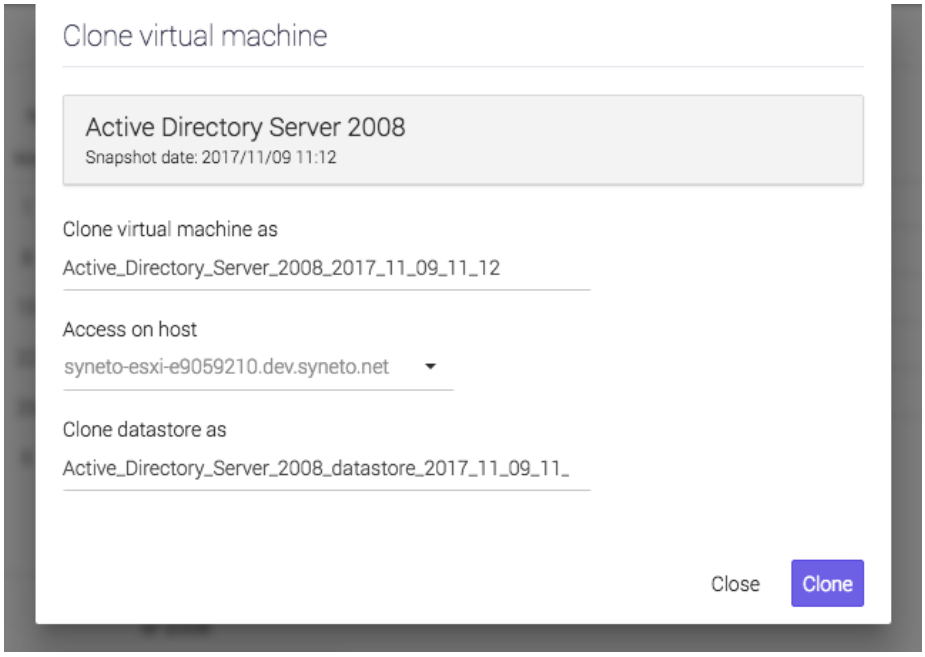

Click clone and wait for the process to finish.

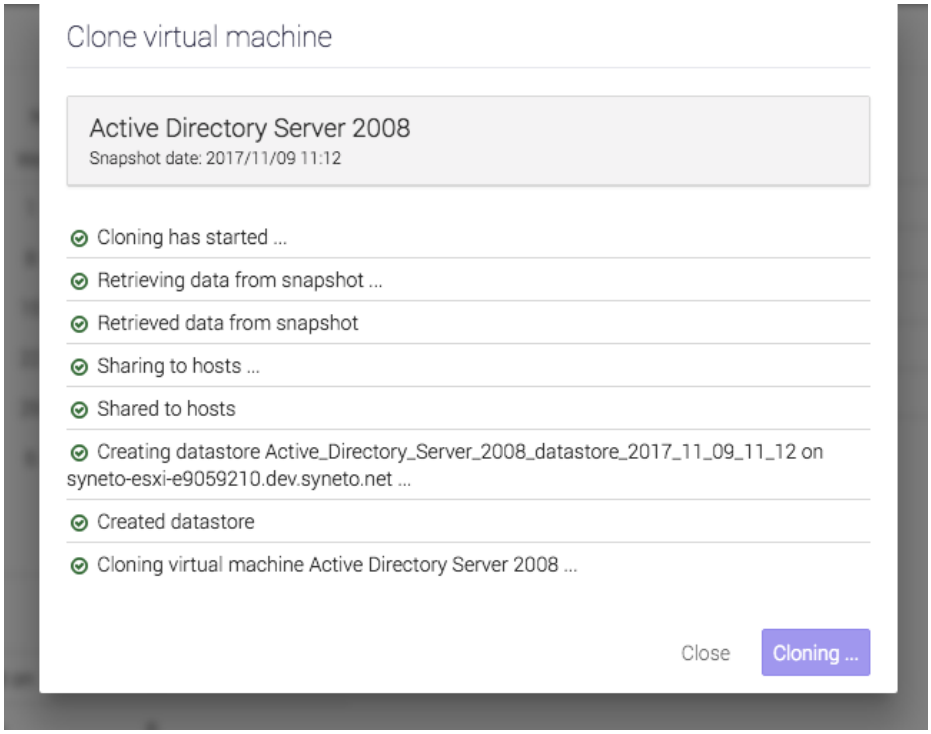

On the *Virtual machines* page we note that a new virtual machine was created. Using the ESXi web interface you can start the virtual machine and recover/copy the data you need.

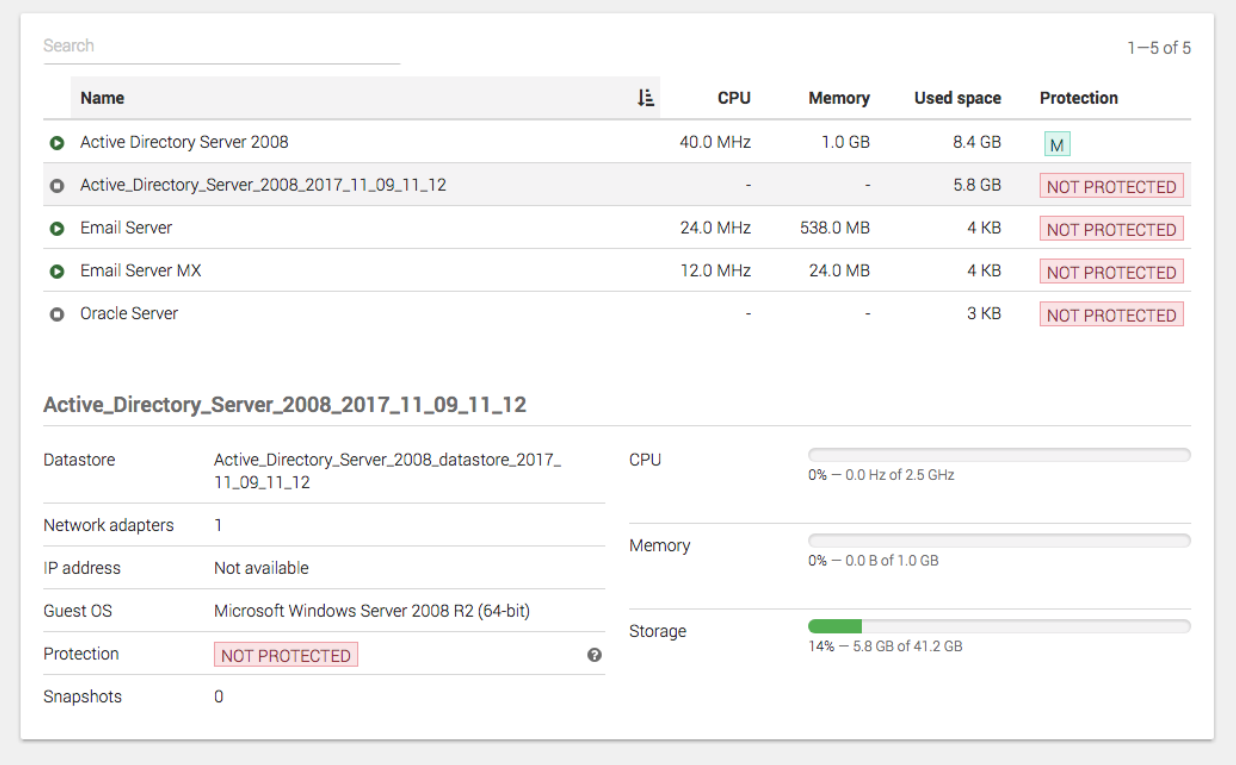

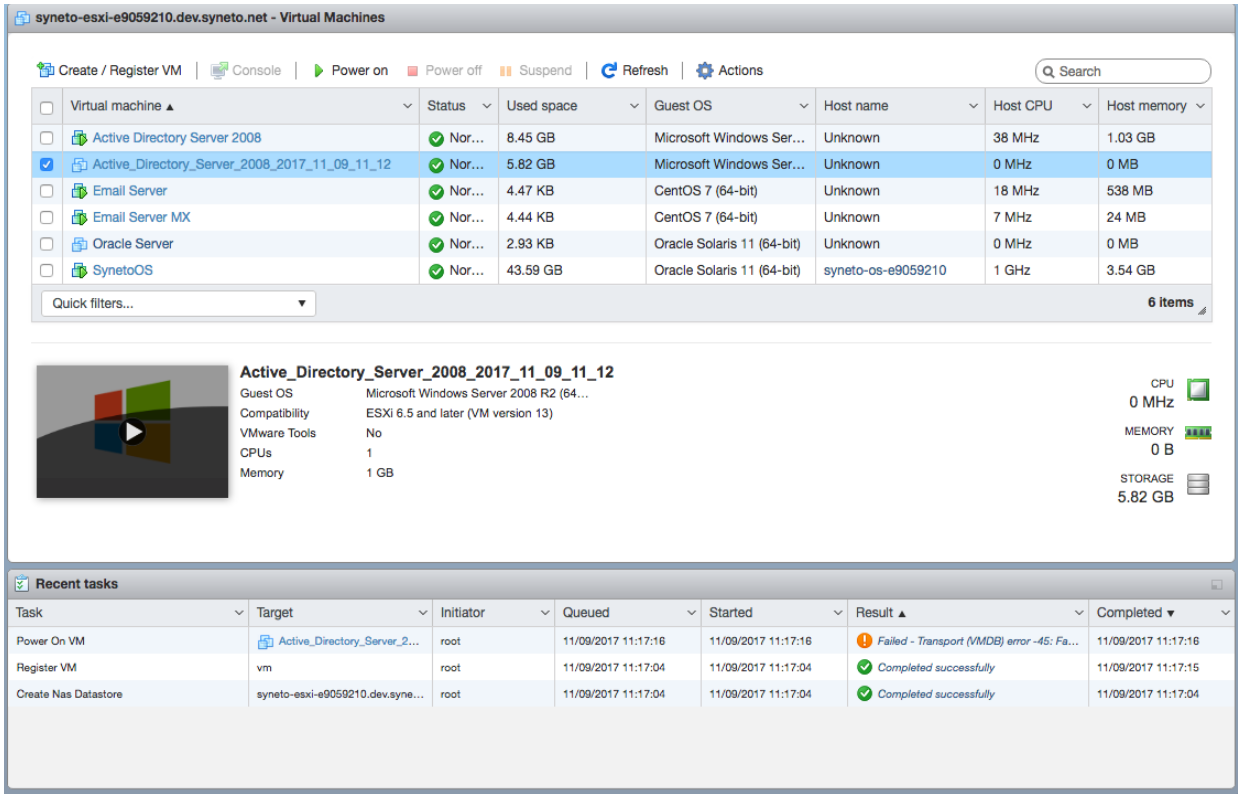

Clean up when you are done

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.

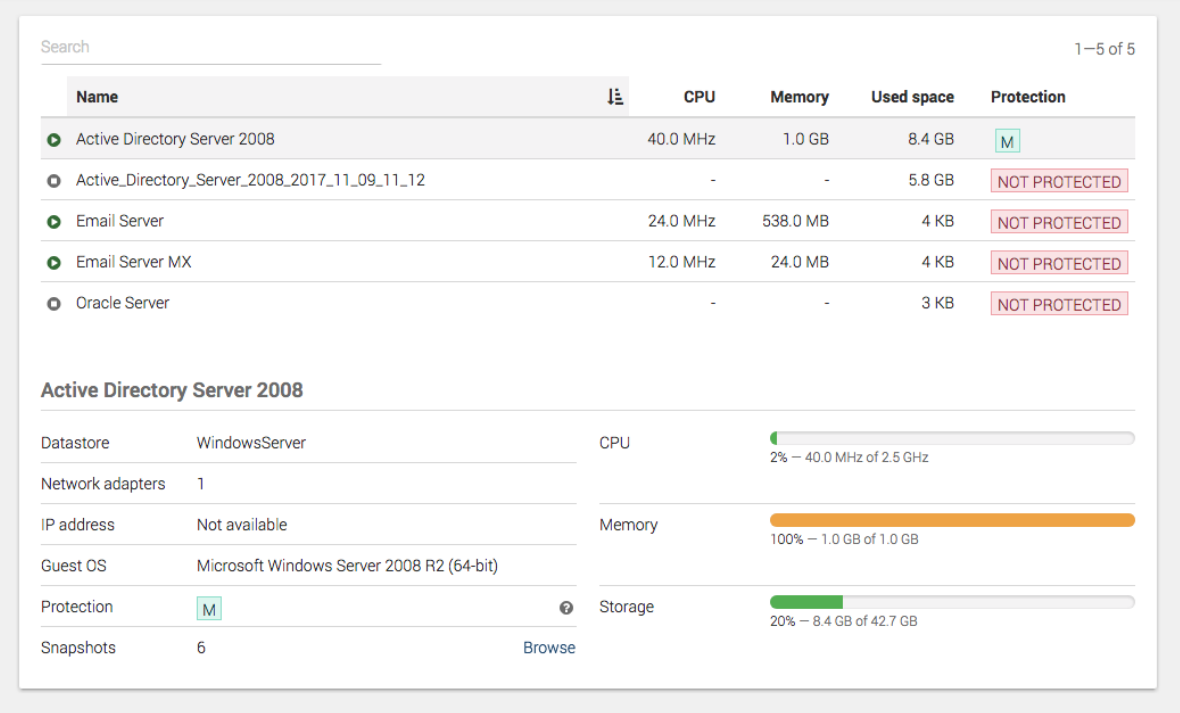

You will see that all cloned snapshots are marked with a label. Click the context menu for a cloned snapshot and select *Clean up clone*.

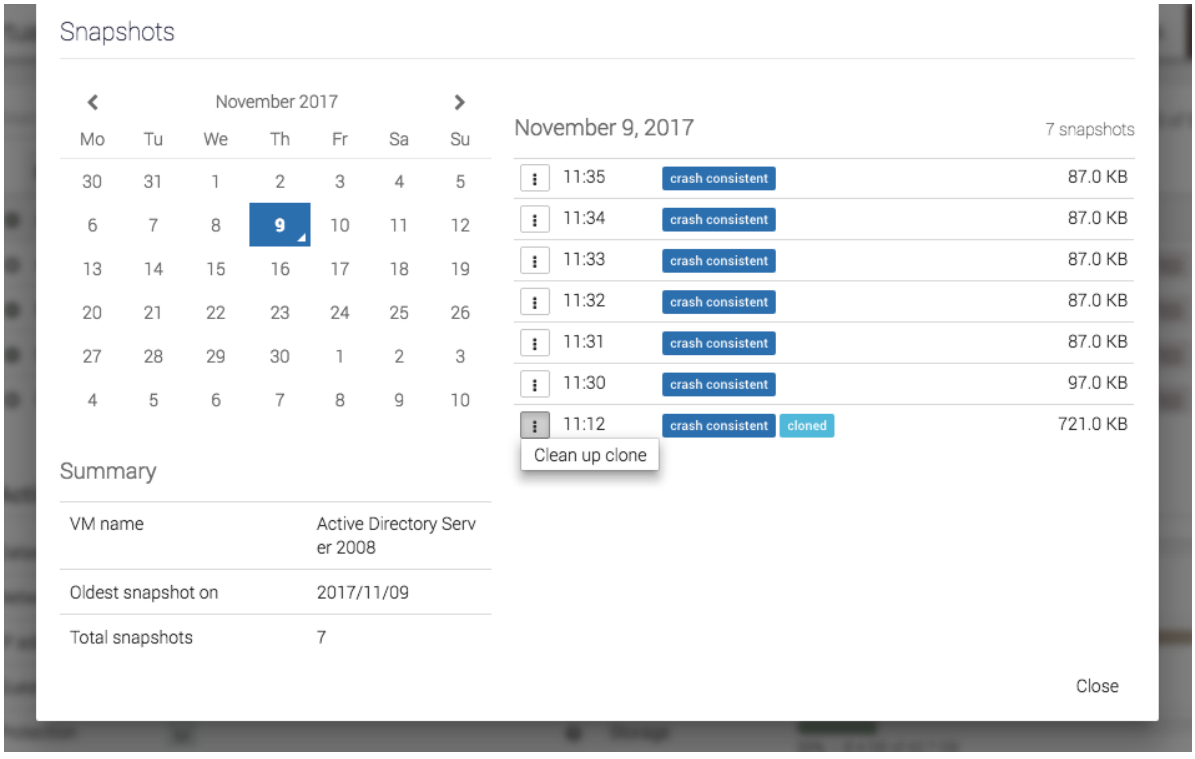

The system will ask you to confirm the process and it will explain all the things that will happen, such as virtual machines will be removed, datastore unmounted from VMware, etc.

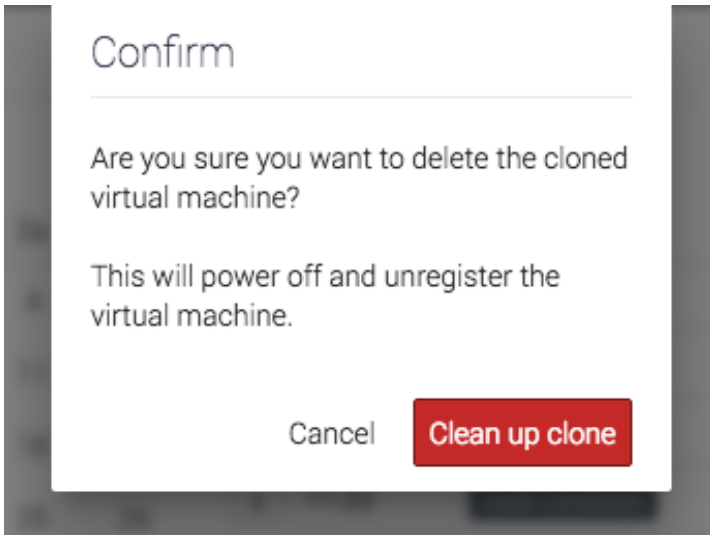

The process takes a little time, you will see a step by step progress.

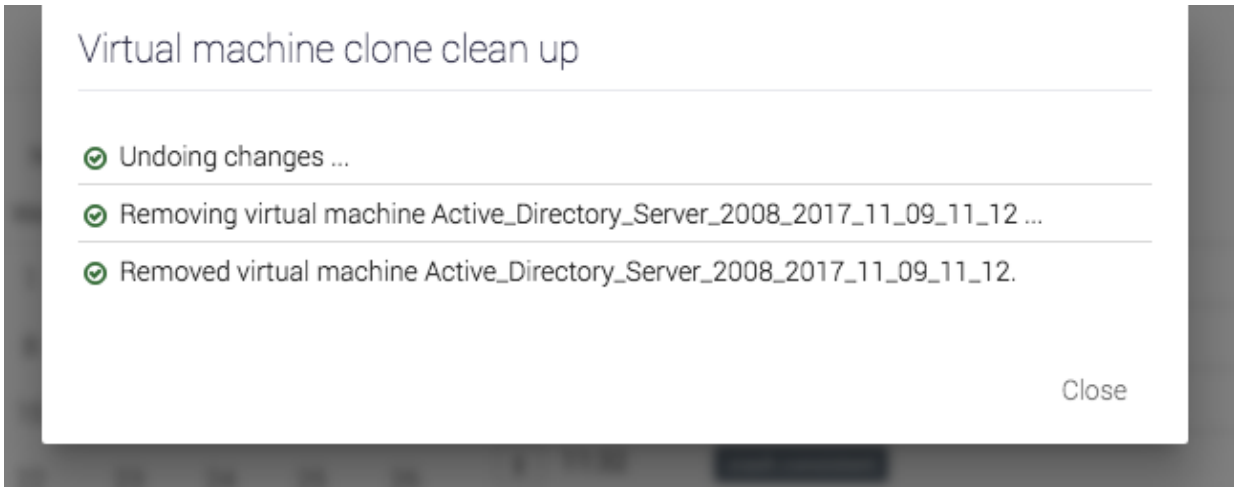

That's it, all that was created, shared, connected, imported for that snapshot clone was reverted. No more mess left behind by the recovery process.

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.

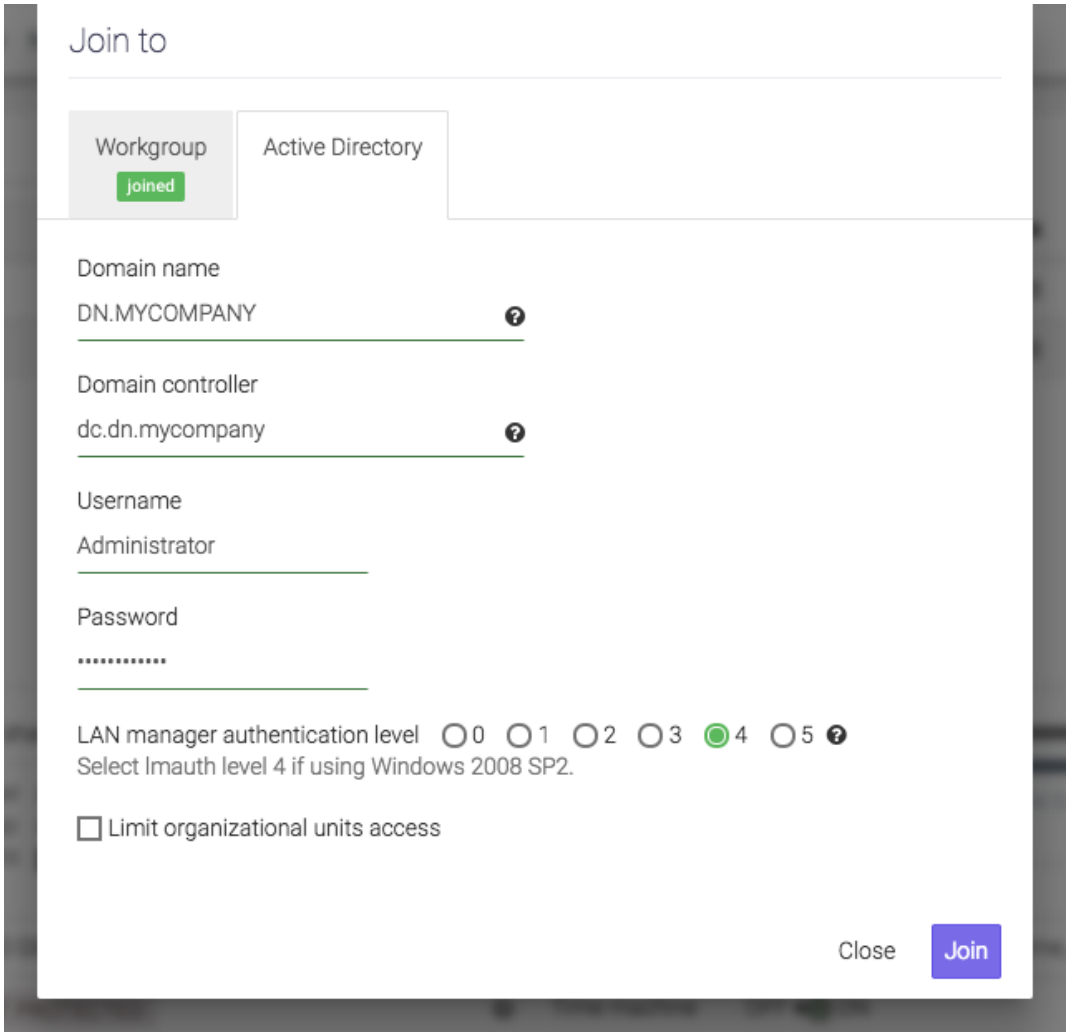

Click on the *Active Directory* tab. Fill in the form with the required data and click *Join*. Once joined you manage your users from the Active Directory server. File permissions are also set from Windows.

NOTE: *Username* must be an Active Directory user name 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*.

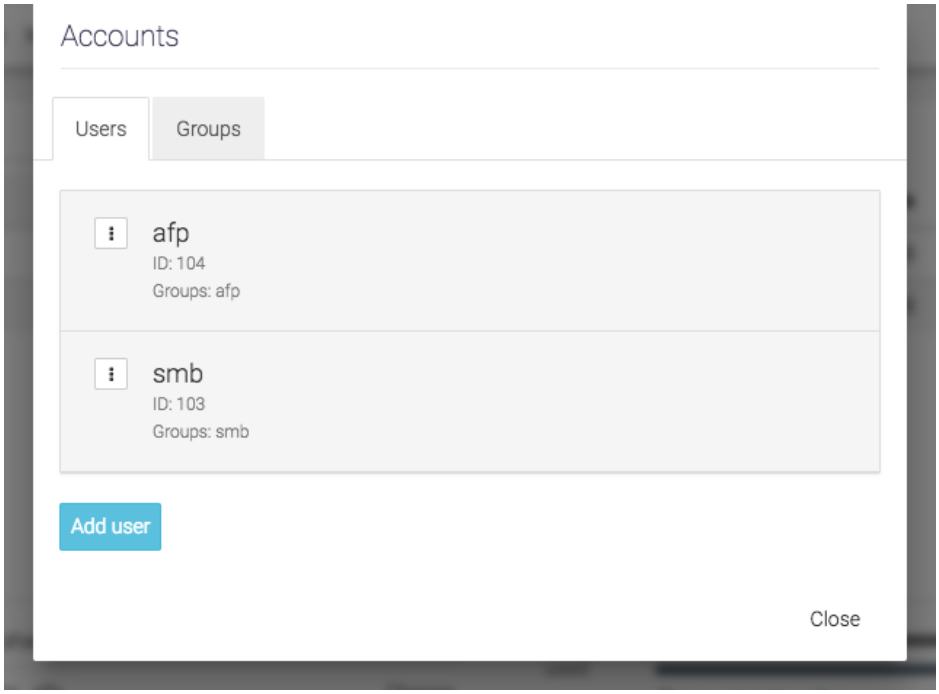

NOTE: By default there are two users created: **afp** with default password **afp** and user **smb** with default password **smb**.

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.

But, of course, this may not be enough. Just click *Add user* and add a new user.

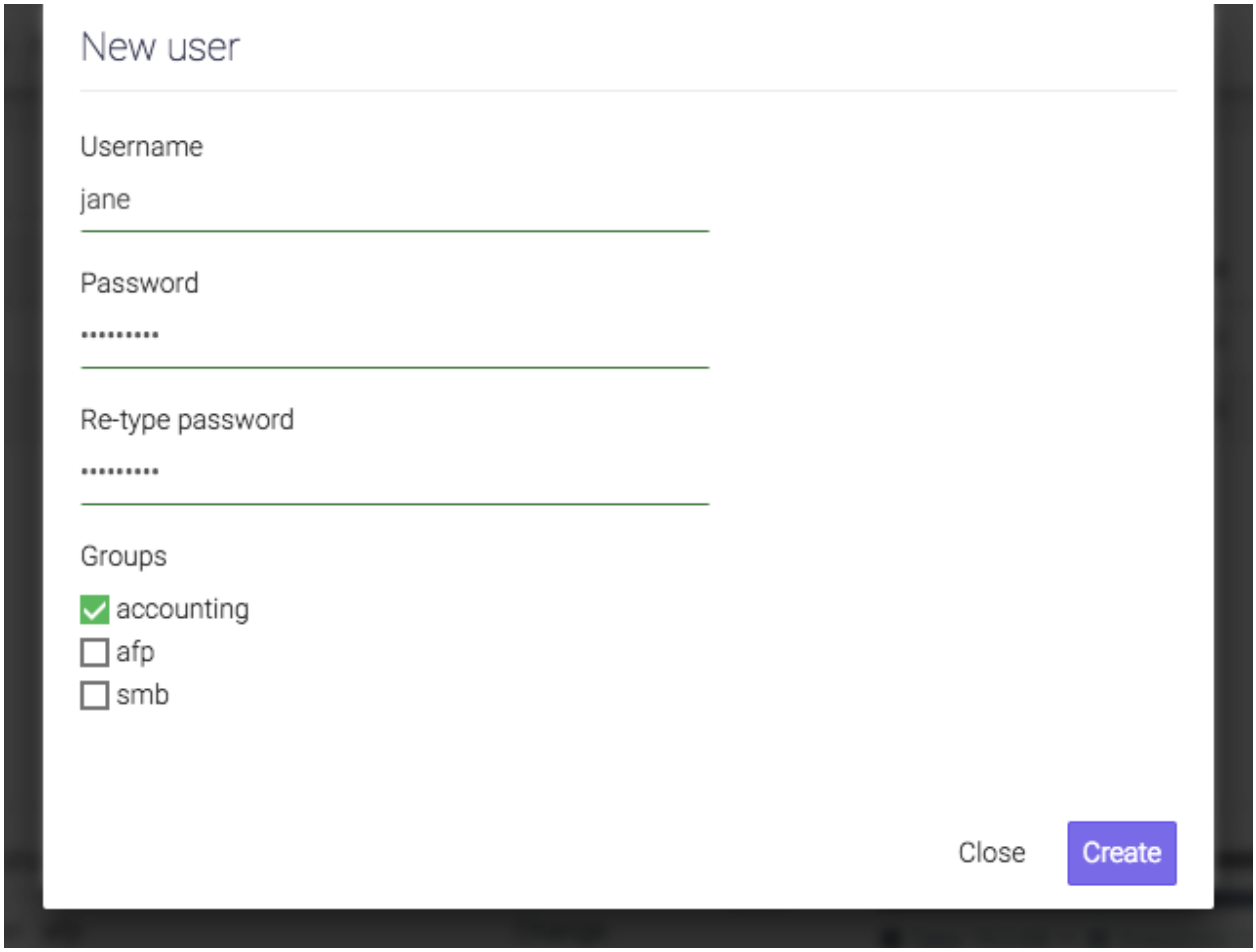

NOTE: each user will be created with a default group having the same name as the user. You can select secondary group for each user from the existing groups.

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

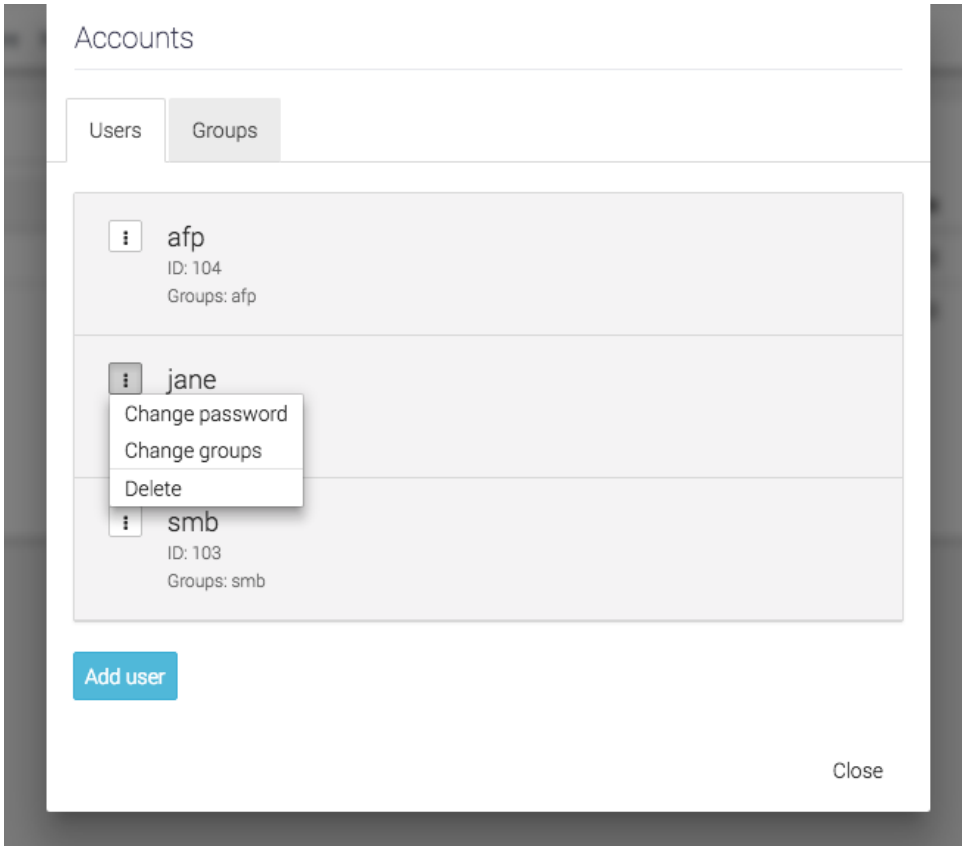

Adding and removing groups is similar to users.

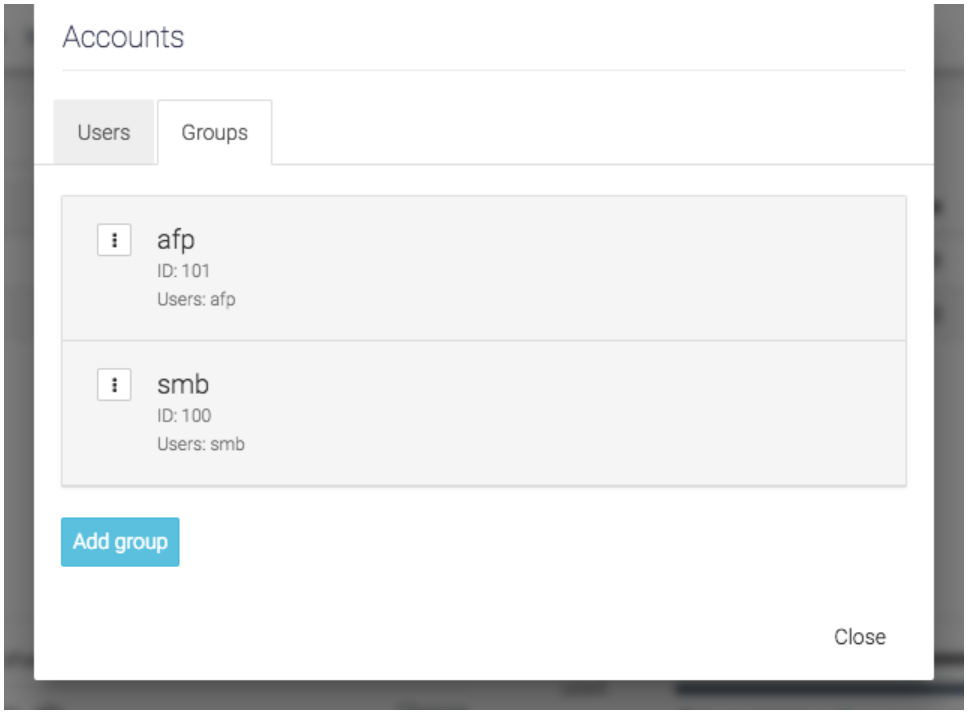

Just click *Add group* and fill in a name fill in a name for the group.

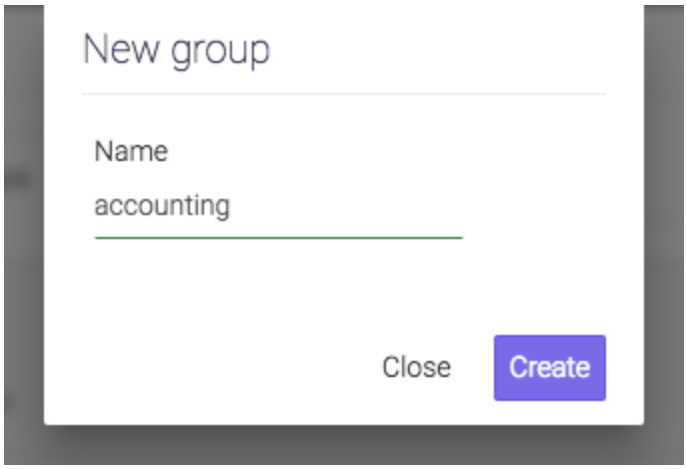

And then, click *Create* to create the group.

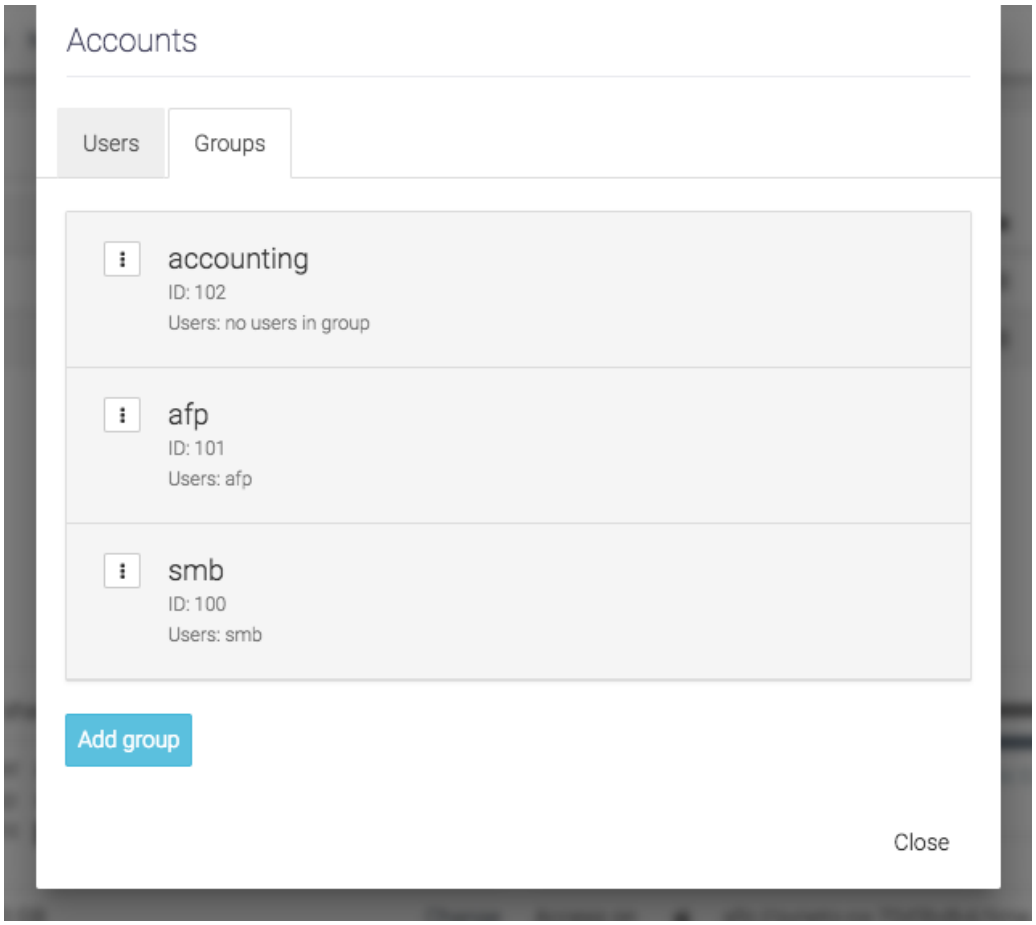

The new group will appear in the list.

Provisioning and managing

SMB - Create share

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

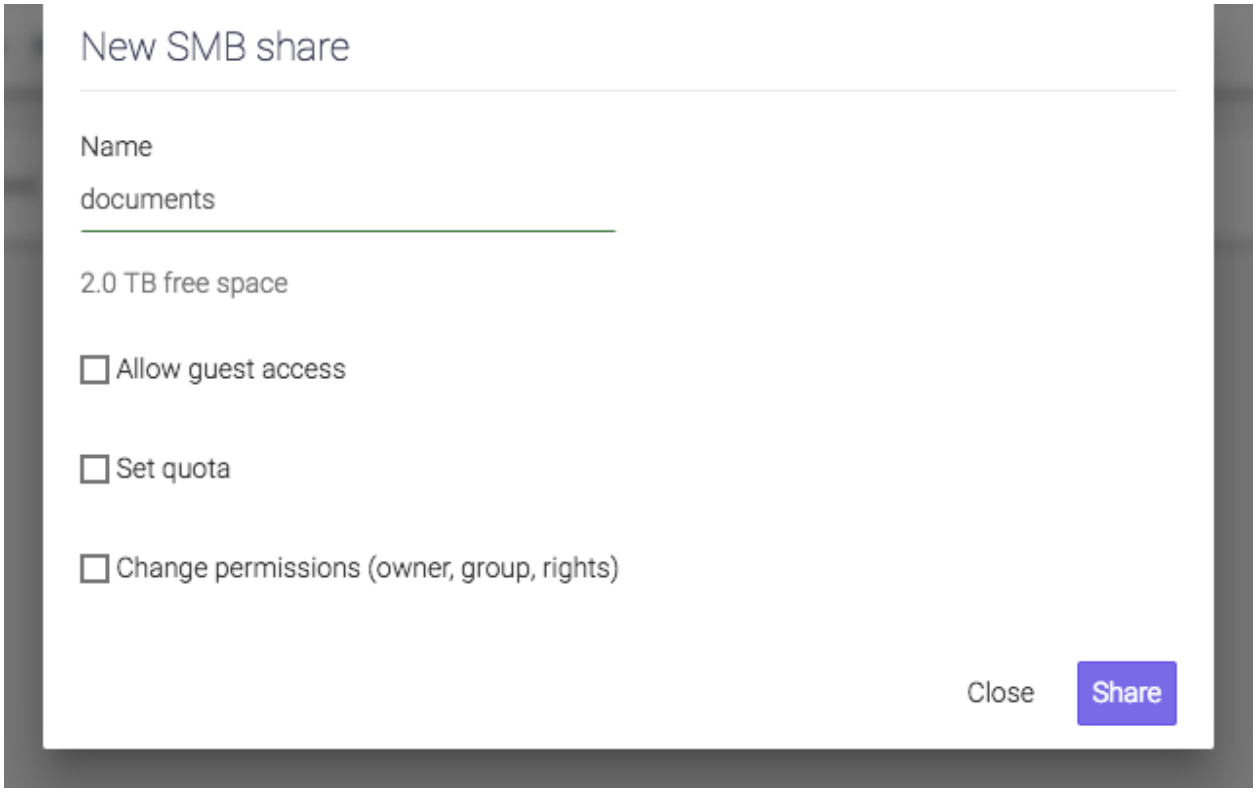

Provide a name for the share and optionally set one or more of the additional parameters:

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

  NOTE: Creating or setting a 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.

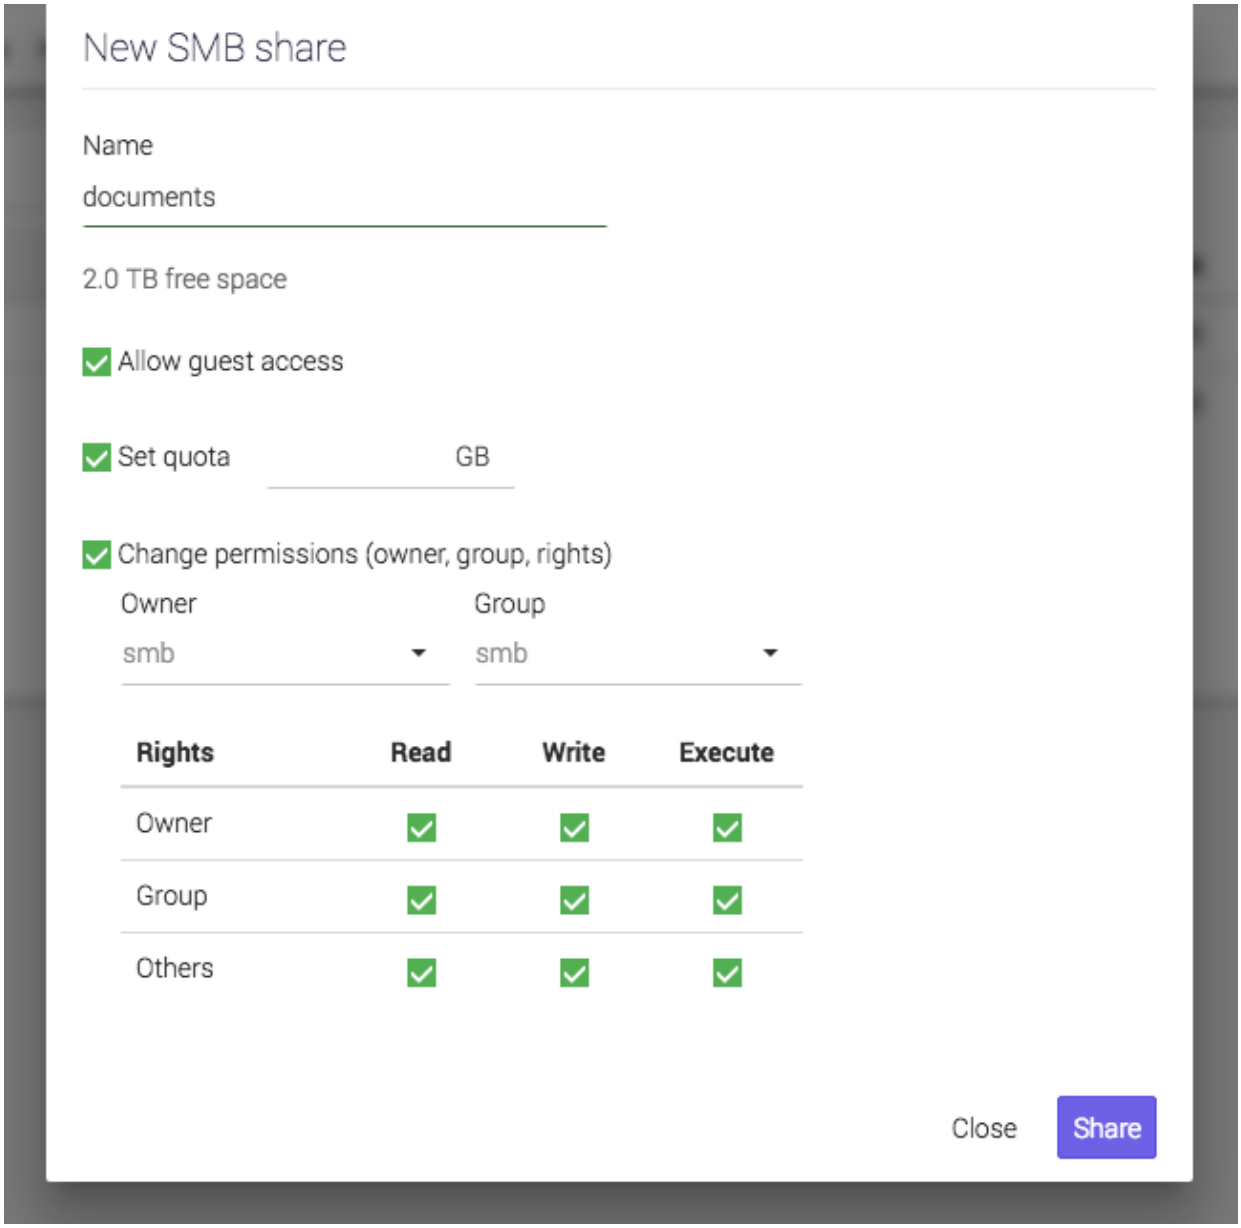

SMB - Manage share

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

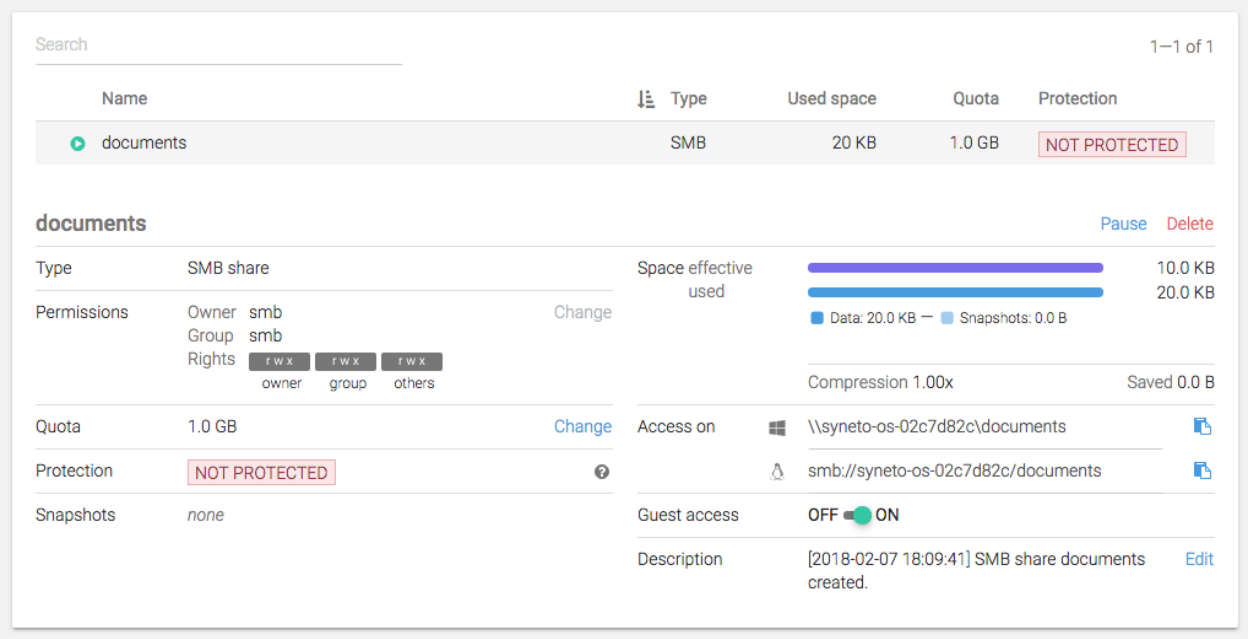

Click on it and you can see its details. From the details section you can perform a set of management actions:

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

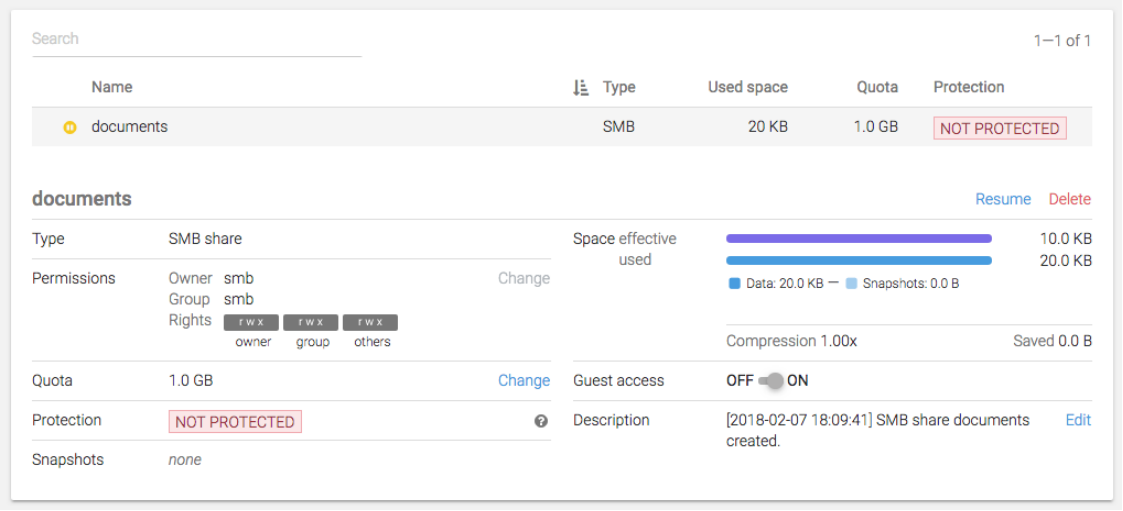

  When a share is paused its icon changes from the green triangle into an orange pause sign.

- *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 a 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*.

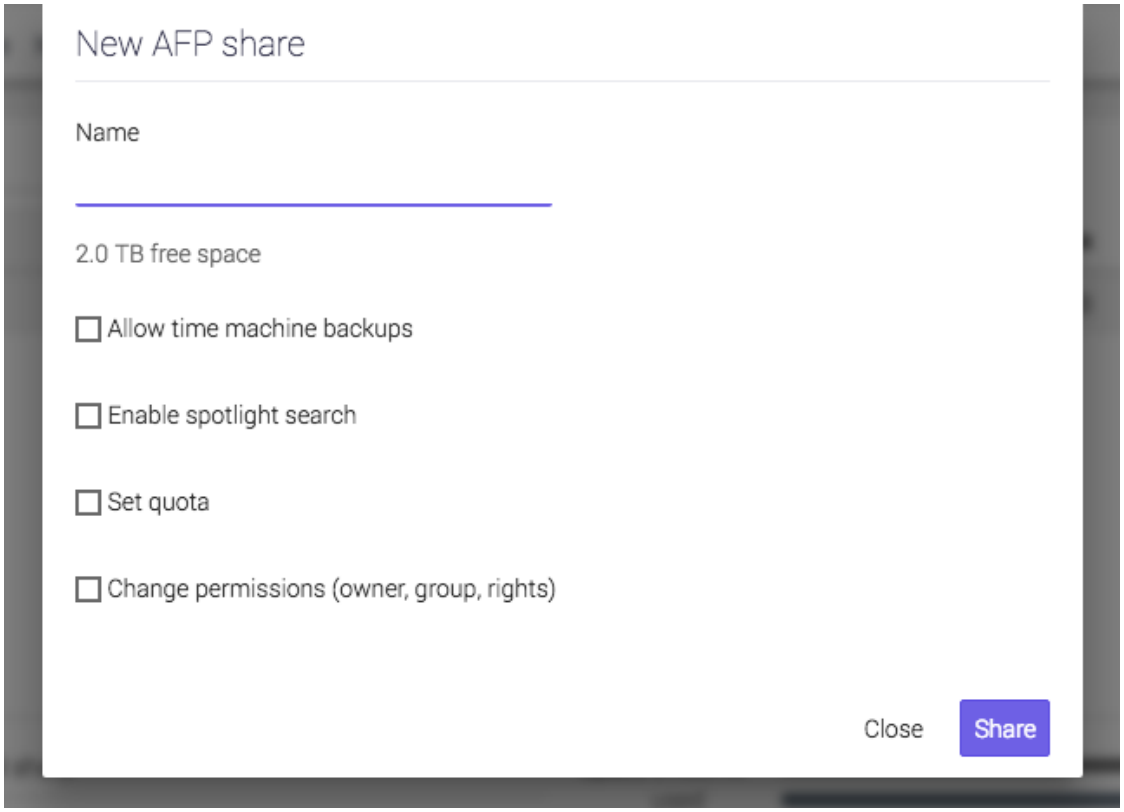

Provide a name for the share and optionally set one or more of the additional parameters:

- *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 shows by default the time when the share was created.

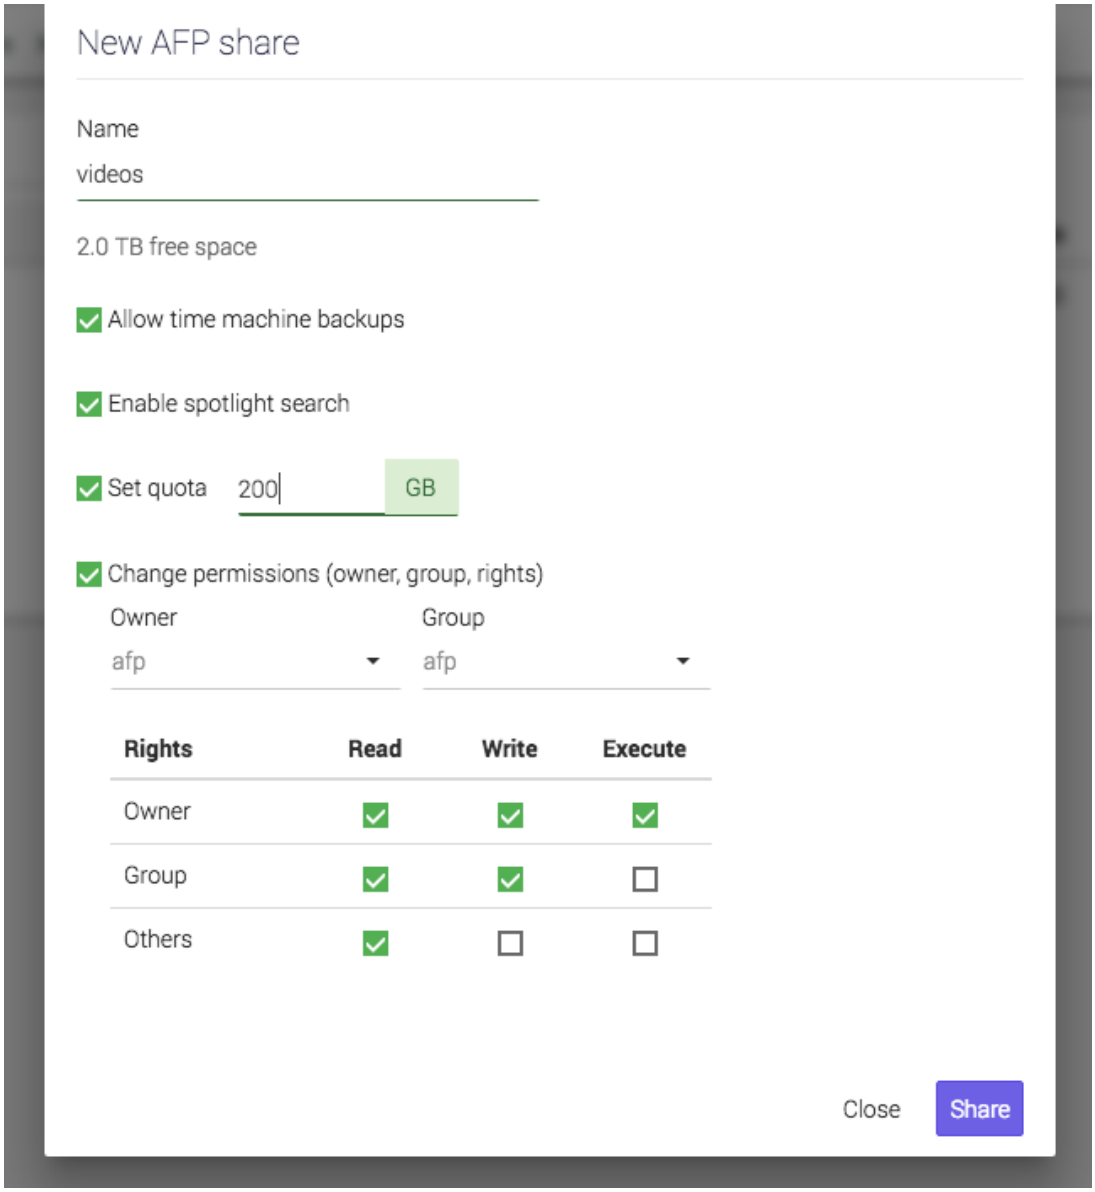

AFP - Manage share

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

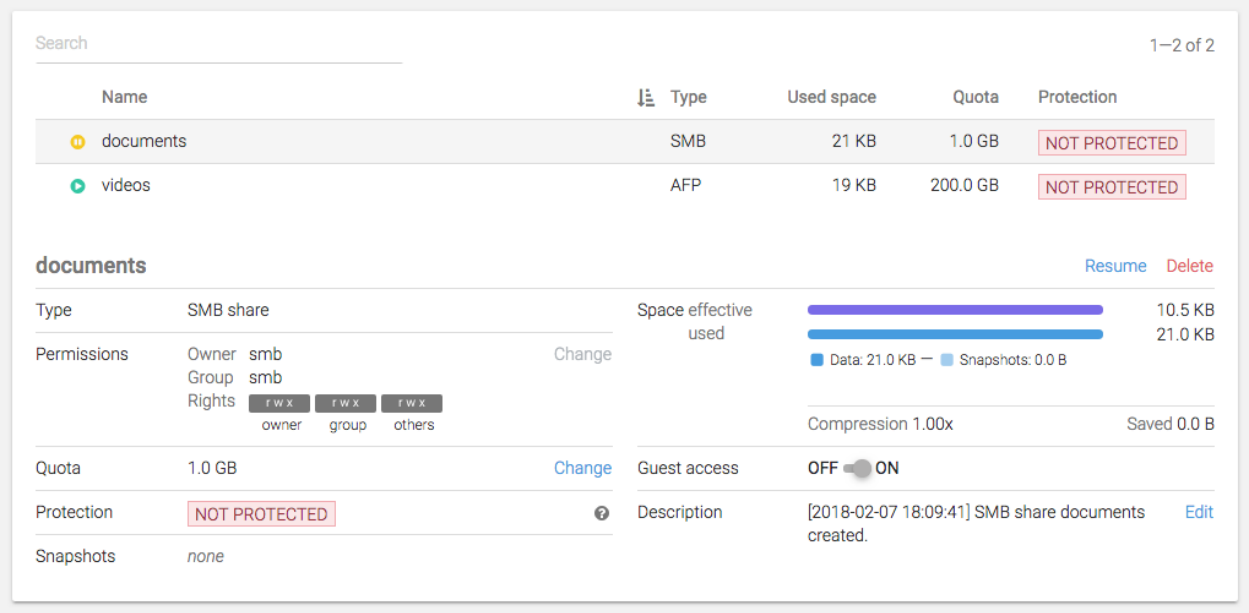

Click on it and your can see its details. From the details section you can perform a set of management actions:

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

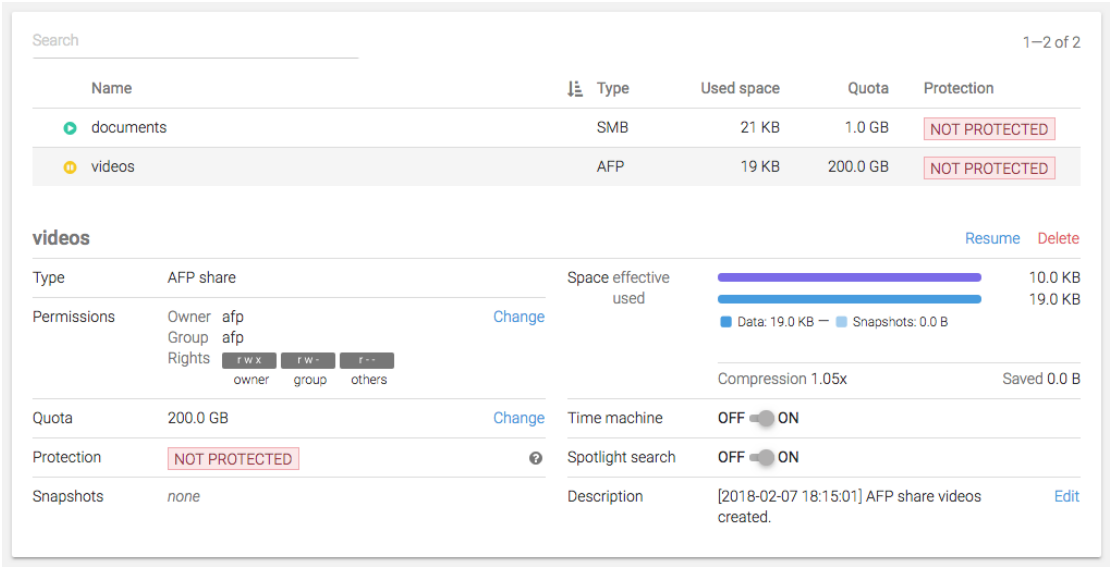

When a share is paused its icon changes from the green triangle into an orange pause sign.

- *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.

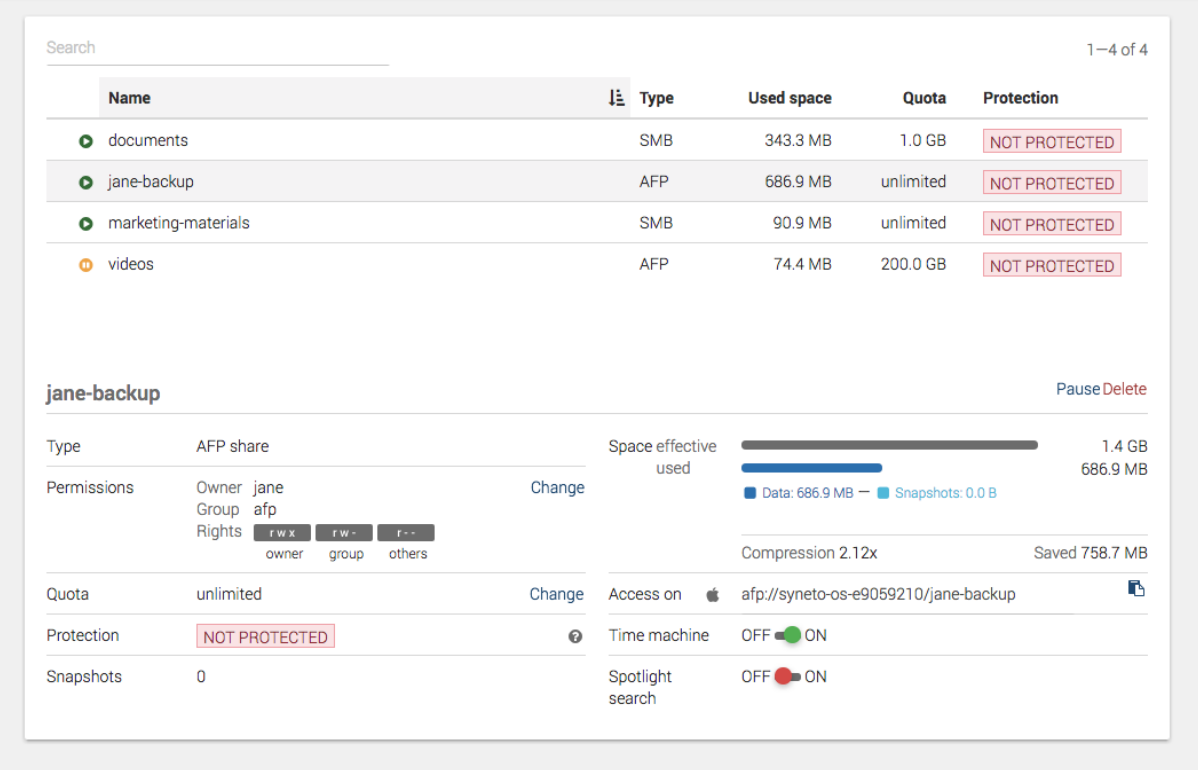

The list has several columns with information about the share: Type, Used space, Quota

Type

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

Used space

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

Quota

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

Analyzing performance

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

Dataset IOPS

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

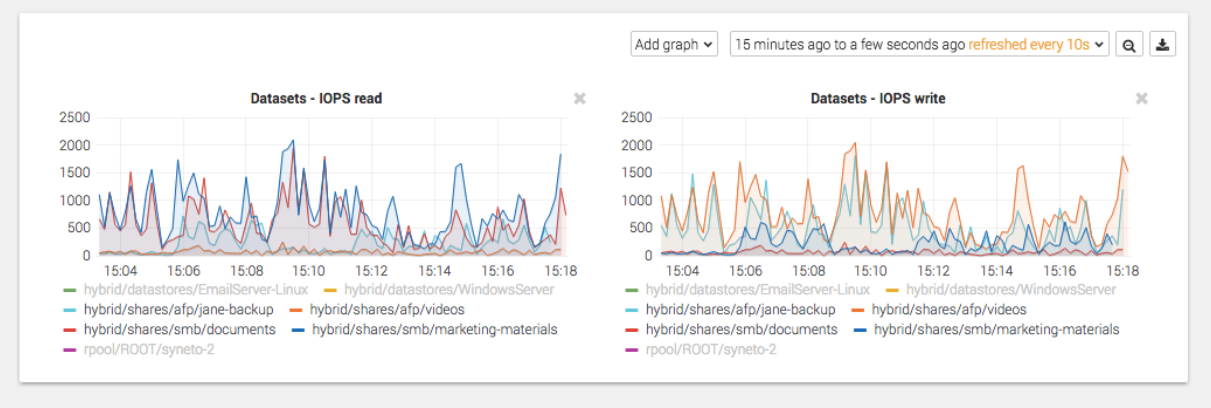

These graphics show the IOPS happening on the share from the Syneto HYPER 's perspective.

Dataset bandwidth

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

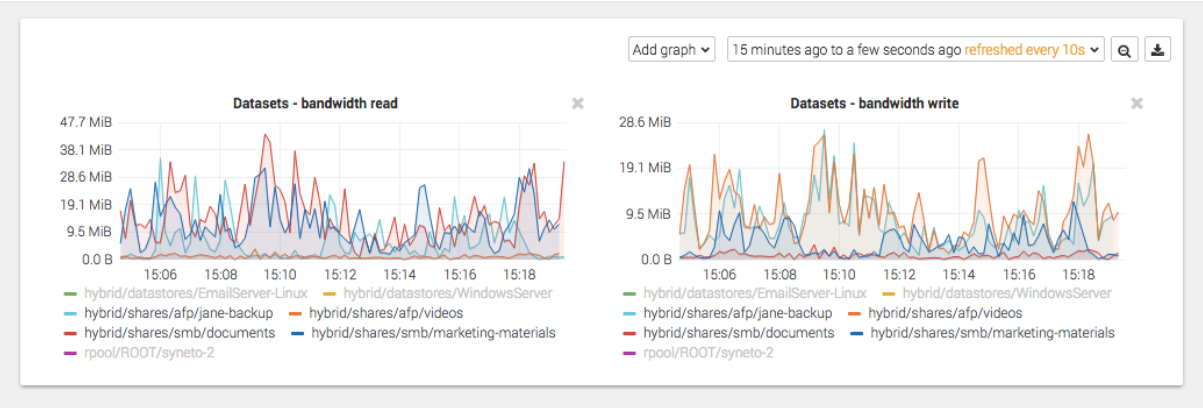

These graphics show the data transfer speed on the share from the Syneto HYPER's perspective.

Dataset average block size

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

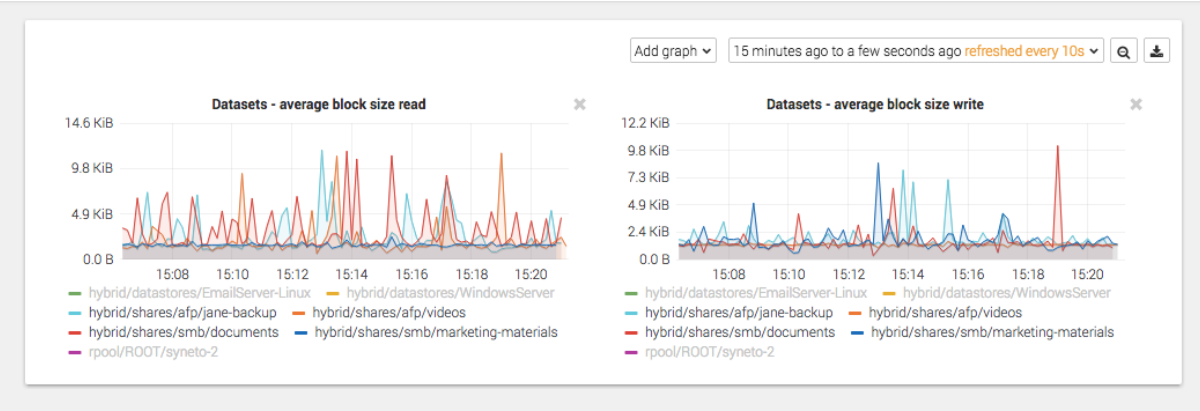

These graphics show the average block size of the data on the share from the Syneto HYPER's perspective.

Recovering data from a snapshot

Access a previous version of your share

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

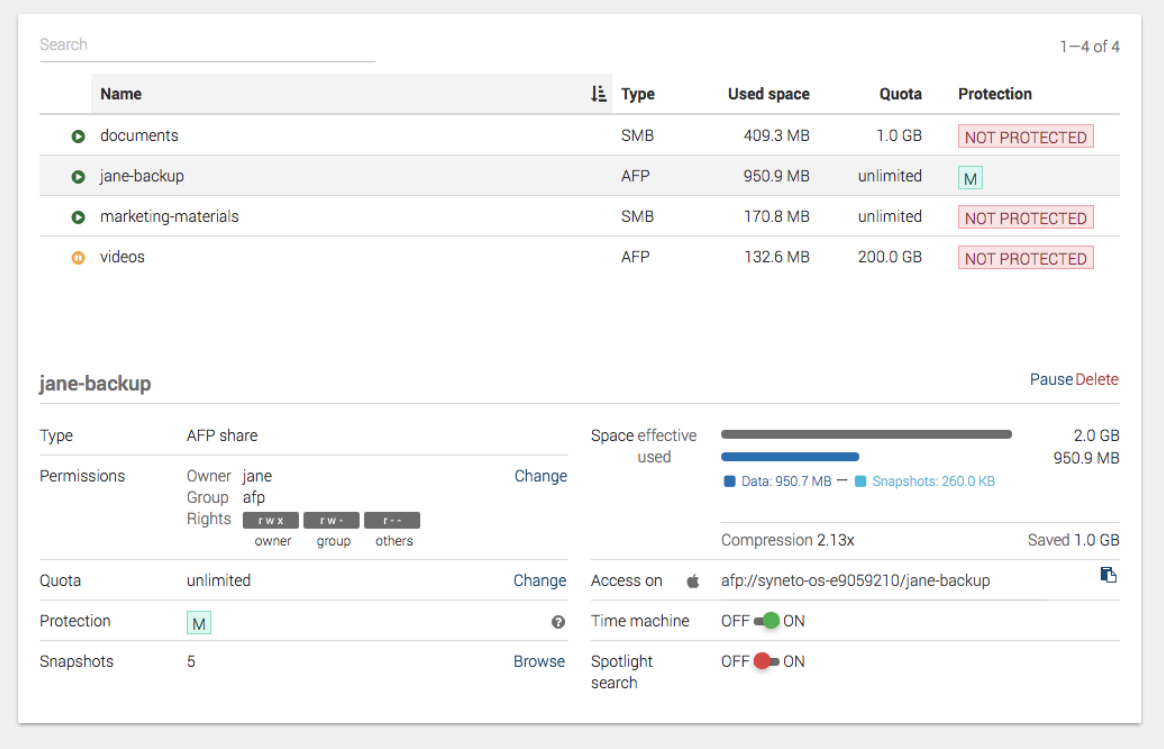

For all these, and many more cases, you can simply click *Browse* in the snapshots section of the detailed view of a share.

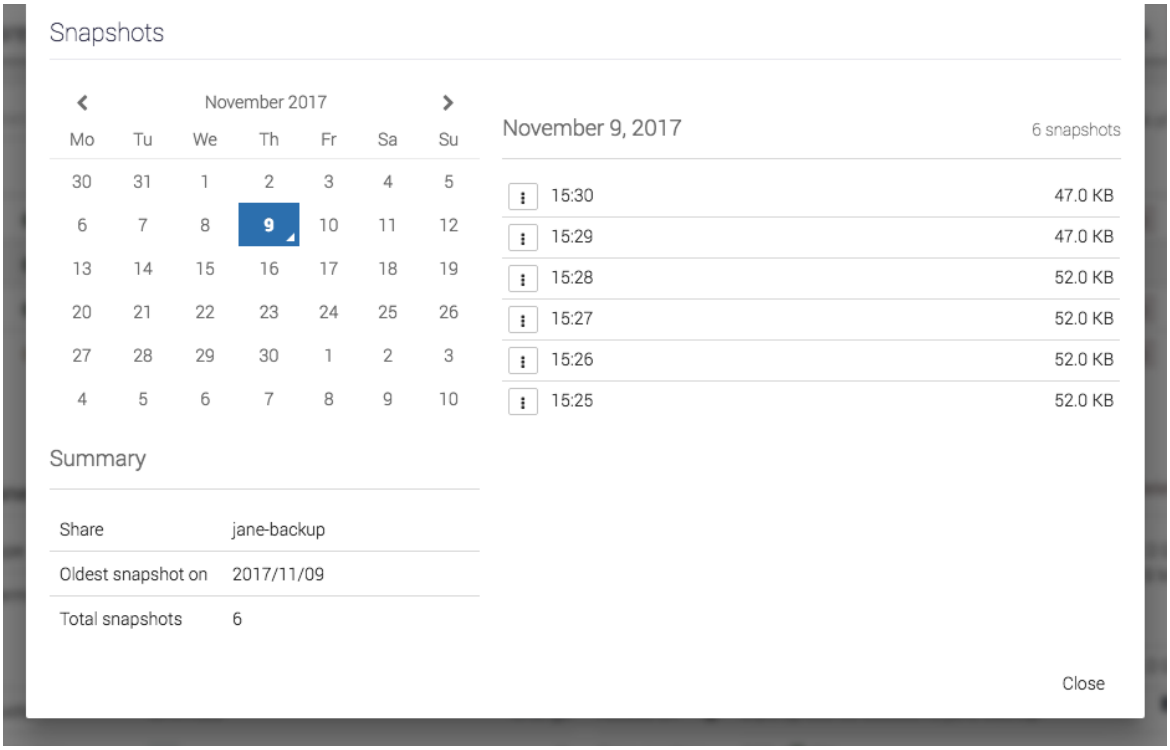

Pick a date from the calendar, and click on the contextual menu for a snapshot.

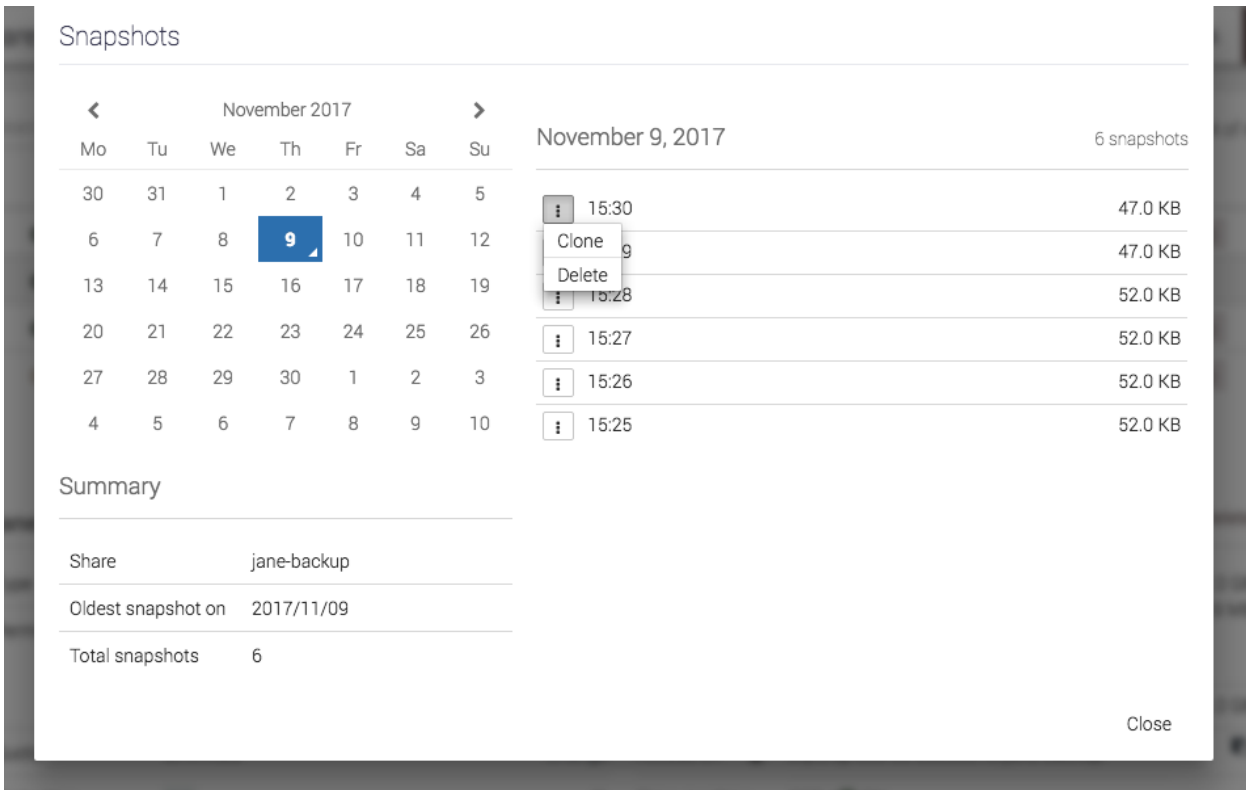

Select clone.

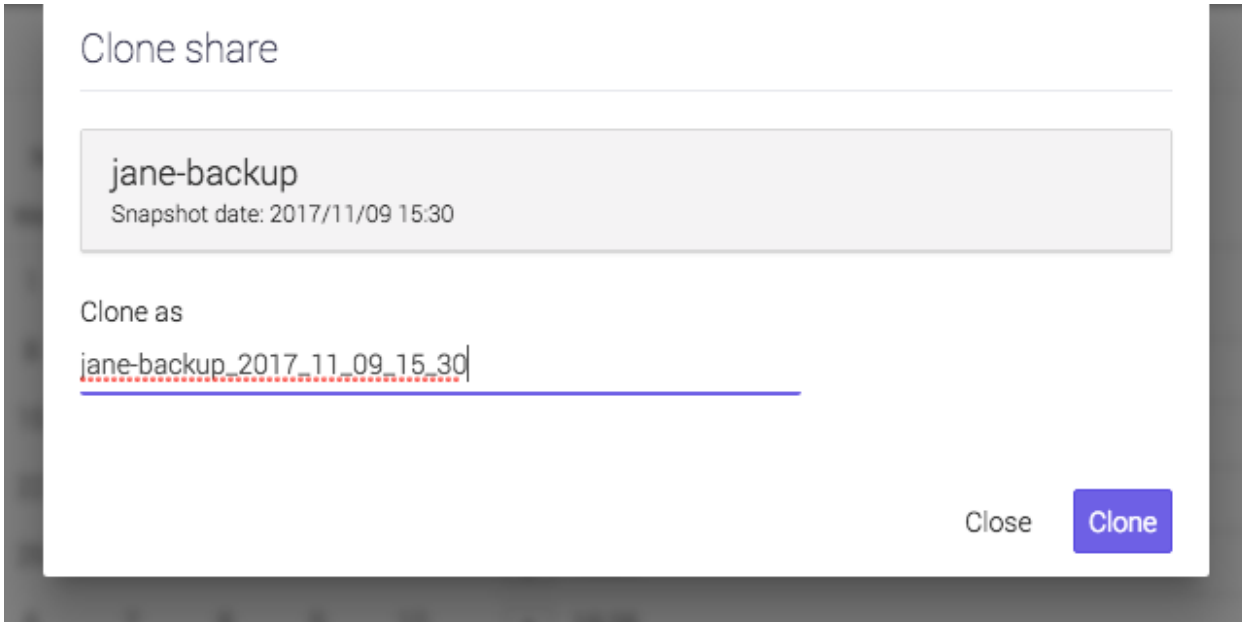

Click clone and wait for the process to finish.

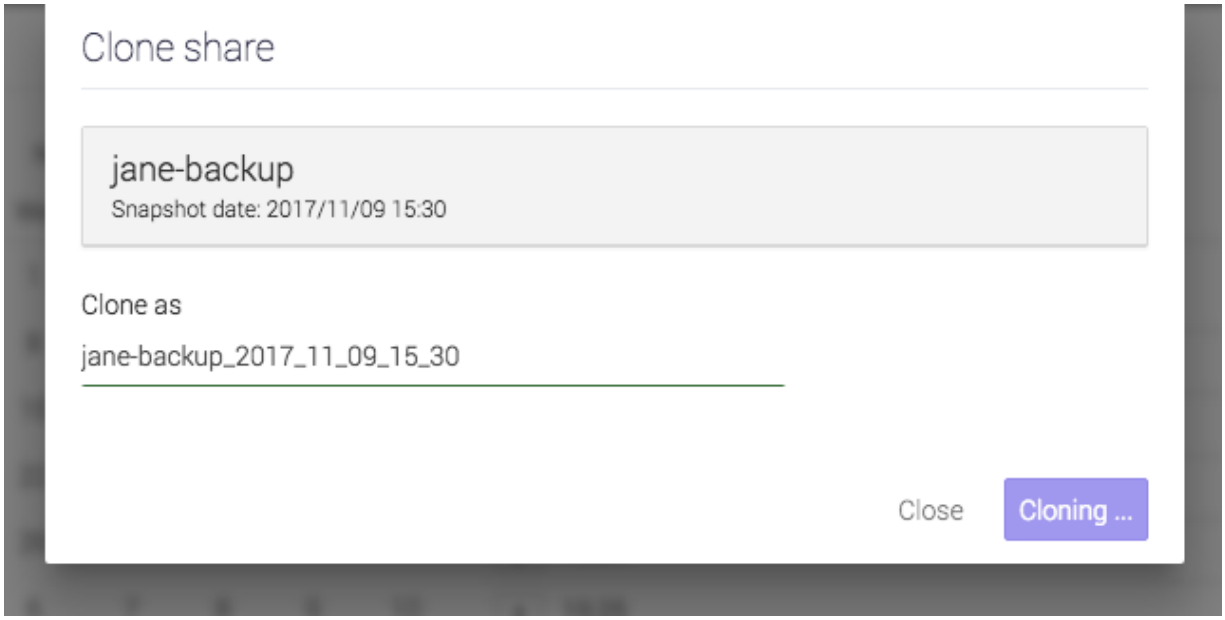

Now you can access a share with the name specified at the clone step. It will contain the data as it was present at the moment of taking the snapshot.

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

Windows previous versions

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

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

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

Clean up when you are done

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

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

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

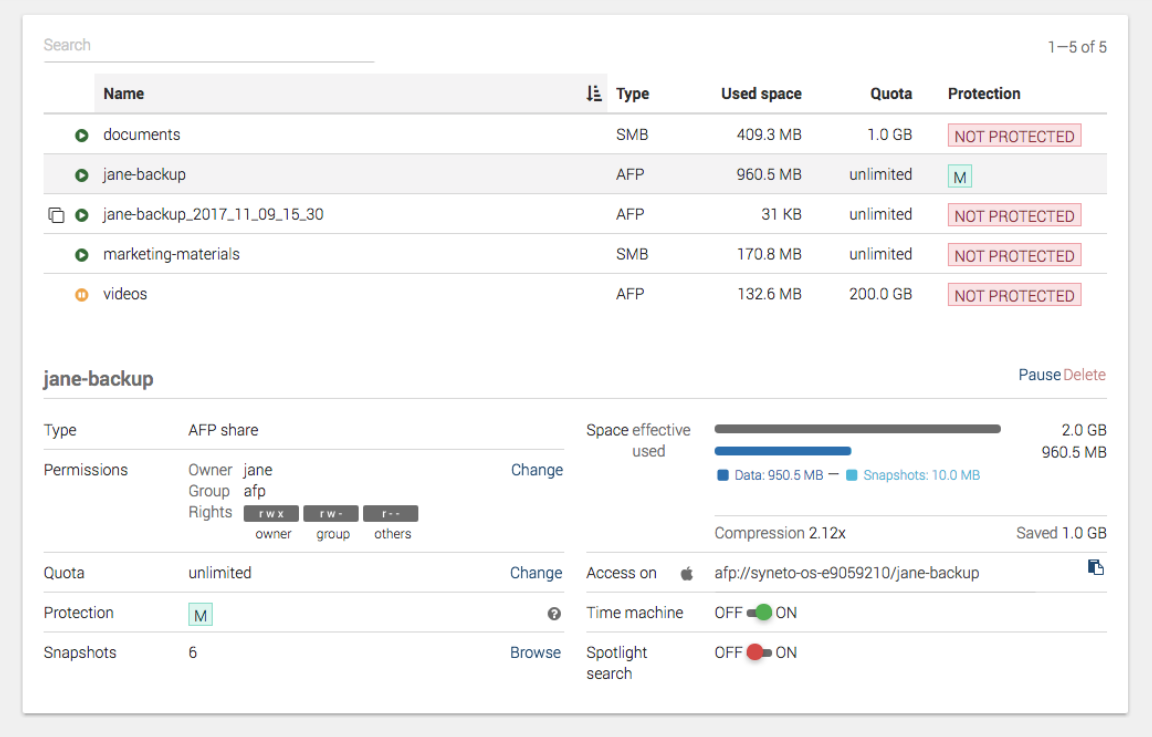

You will see that all cloned snapshots are marked with a label. Click the context menu for a cloned snapshot and select *Clean up clone*.

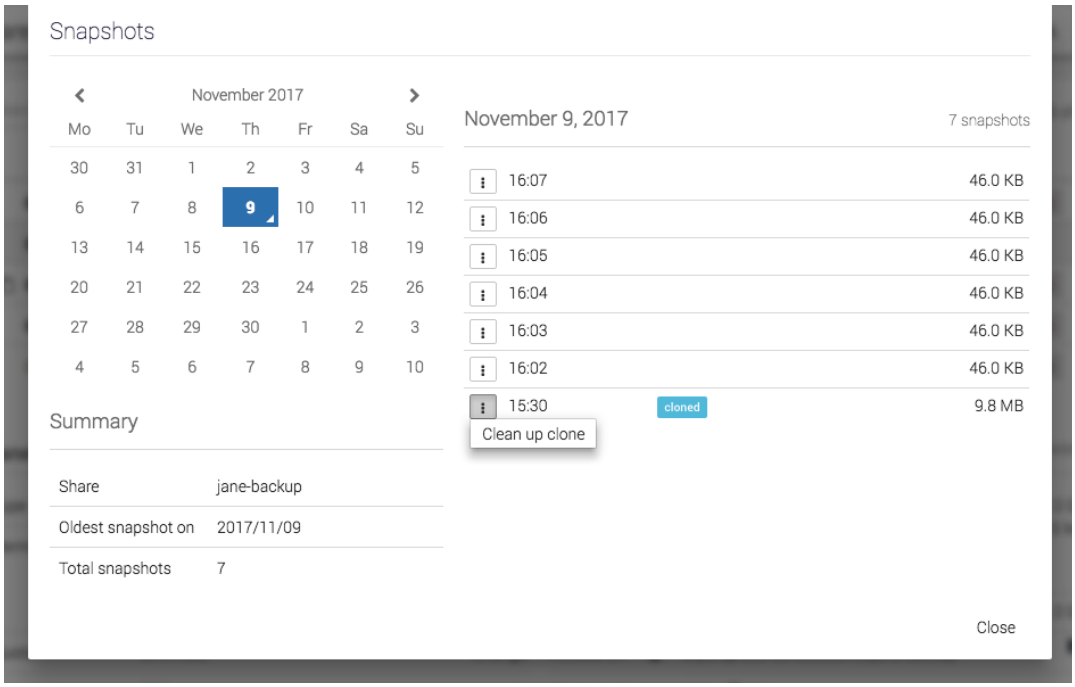

The system will ask you to confirm the process and it will explain all the things that will happen.

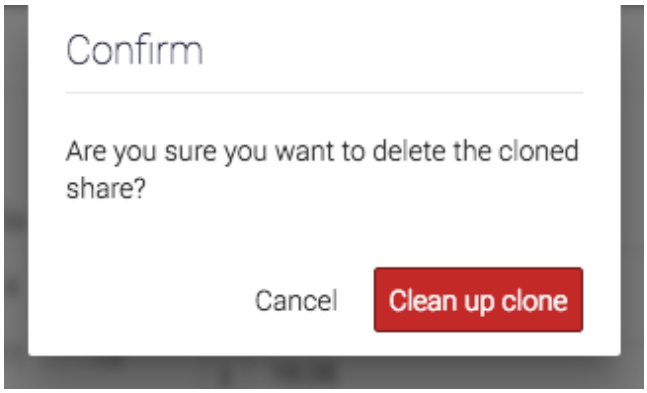

The process takes only a second, you will see a little spinner in the list.

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

Reverting a clone

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

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

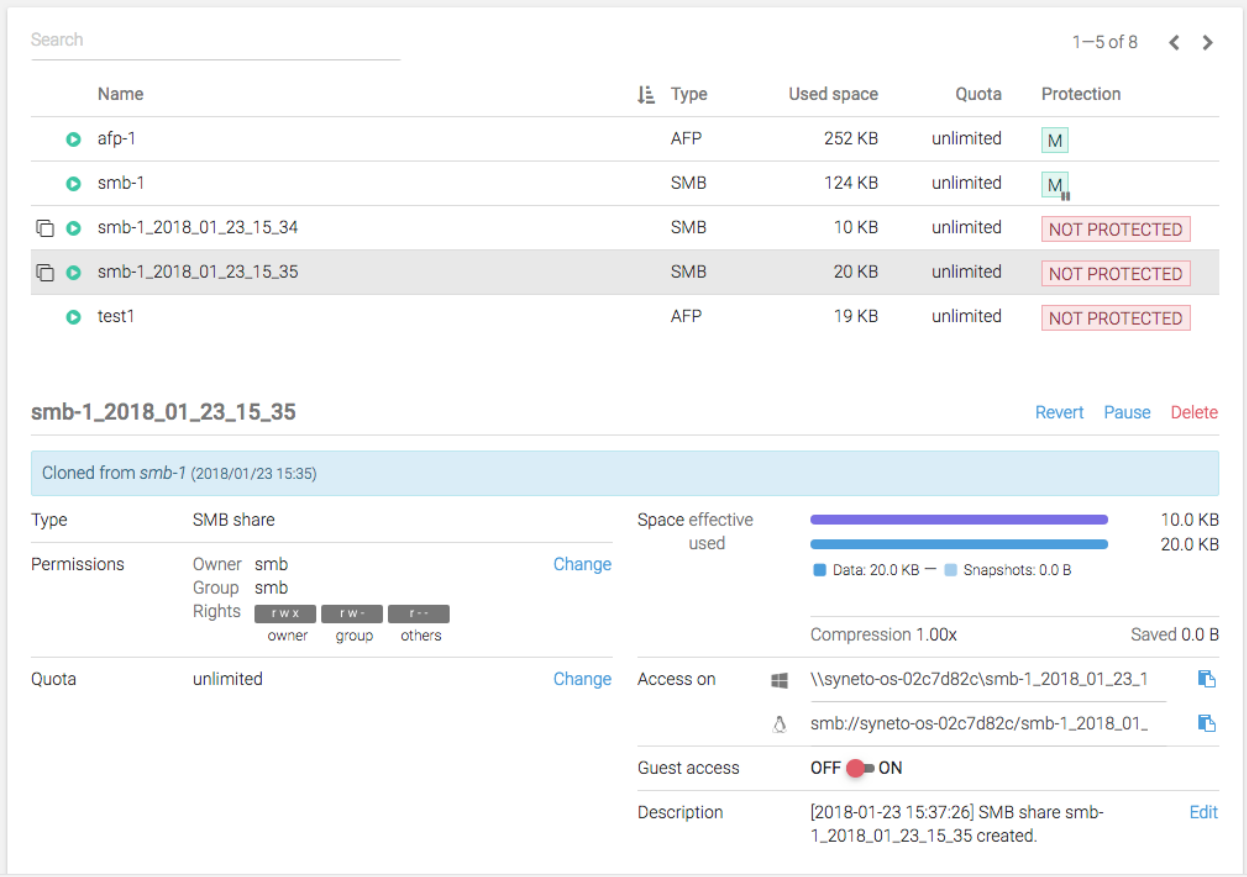

A clone revert process will do the following steps:

- 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.

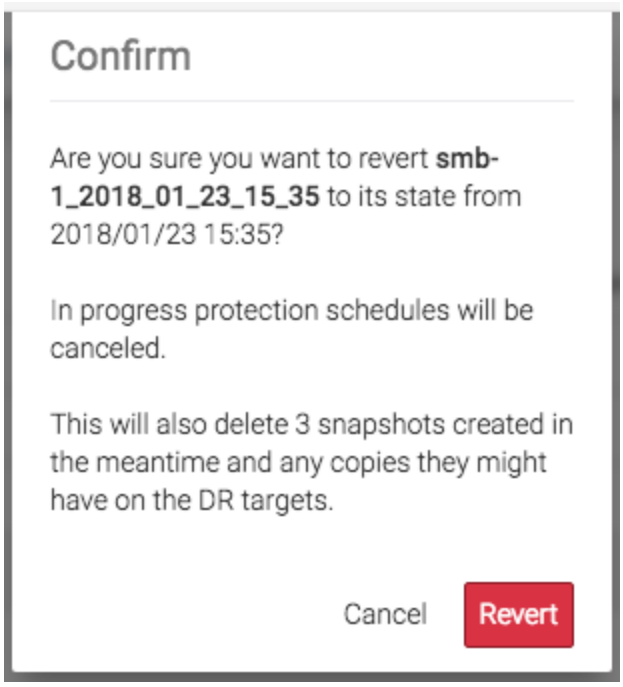

During the revert process you will see a step-by-step progress. If something goes wrong, an appropriate message will be shown.

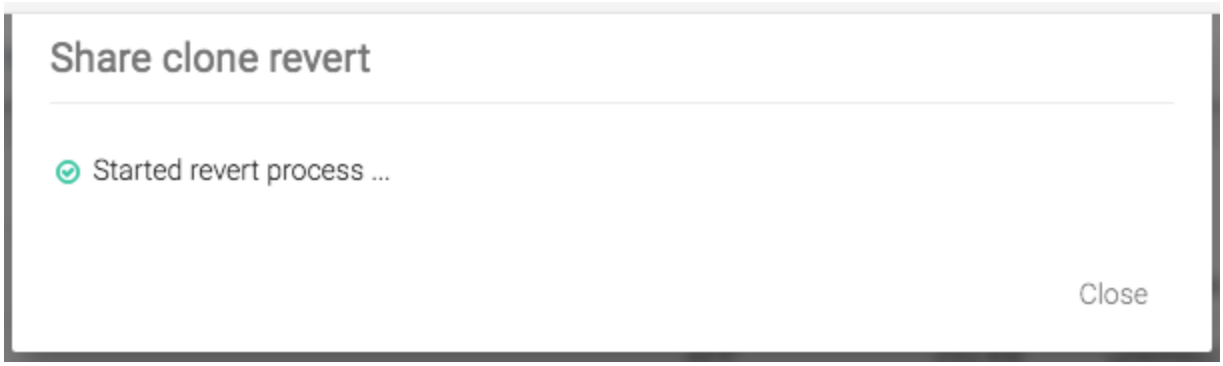

WARNING: On DR unit, on a received share, you can create clones but you cannot revert the clone.

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

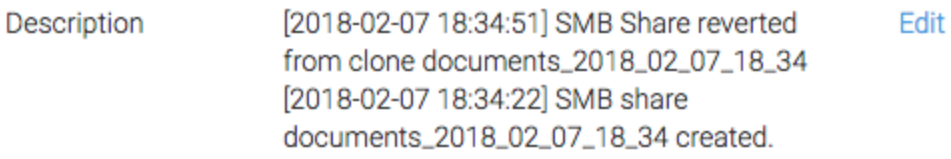

VOLUMES

Introduction

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

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

Access management

iSCSI remote initiators

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

Go to *Main menu → Shares → Initiators*.

This will show a dialog with the currently defined initiators.

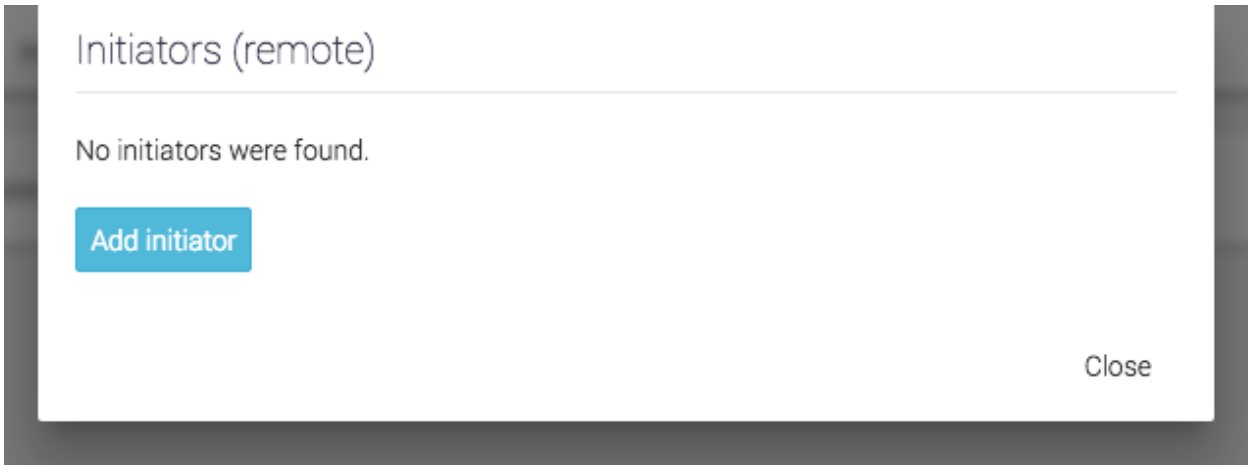

By default, no initiators are defined. To add a new initiator click *Add initiator* and specify the group name and identifier.

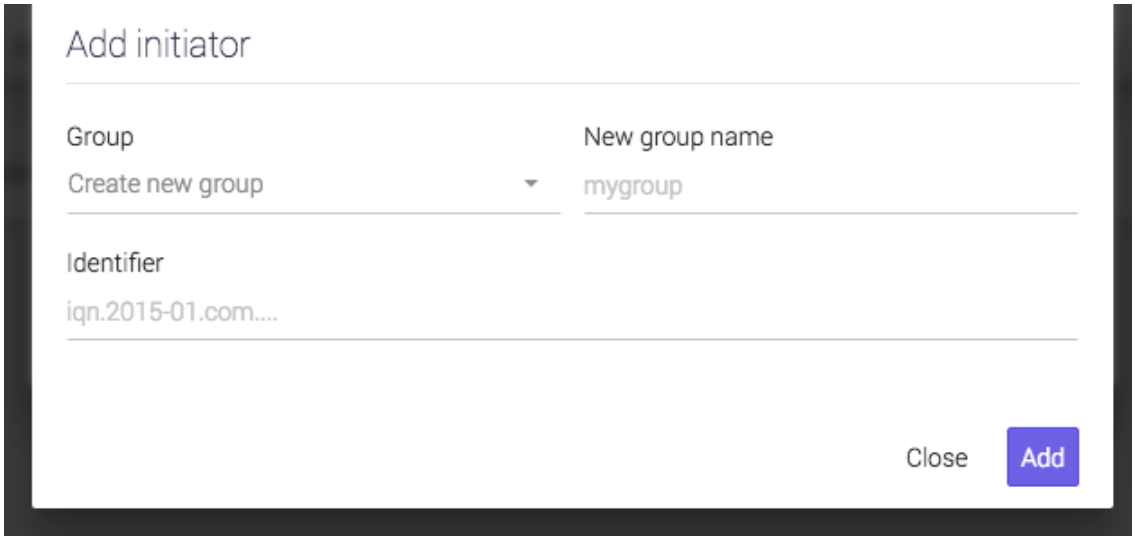

NOTE: each identifier must be in a group named by you. This is necessary even if there is only one initiator in the group.

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

iSCSI local targets

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

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

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

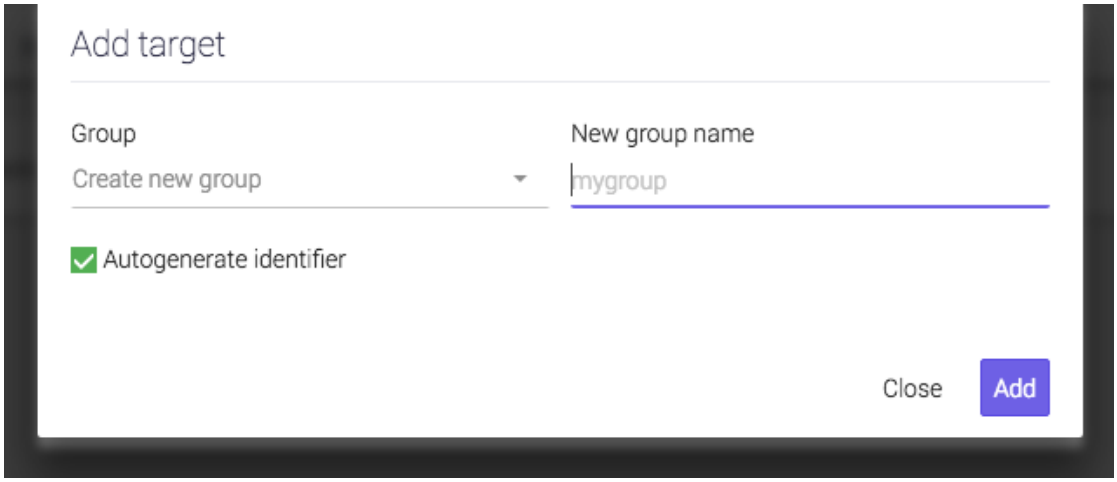

To create a new target click *Add target*. The form will ask you the group. Options are:

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

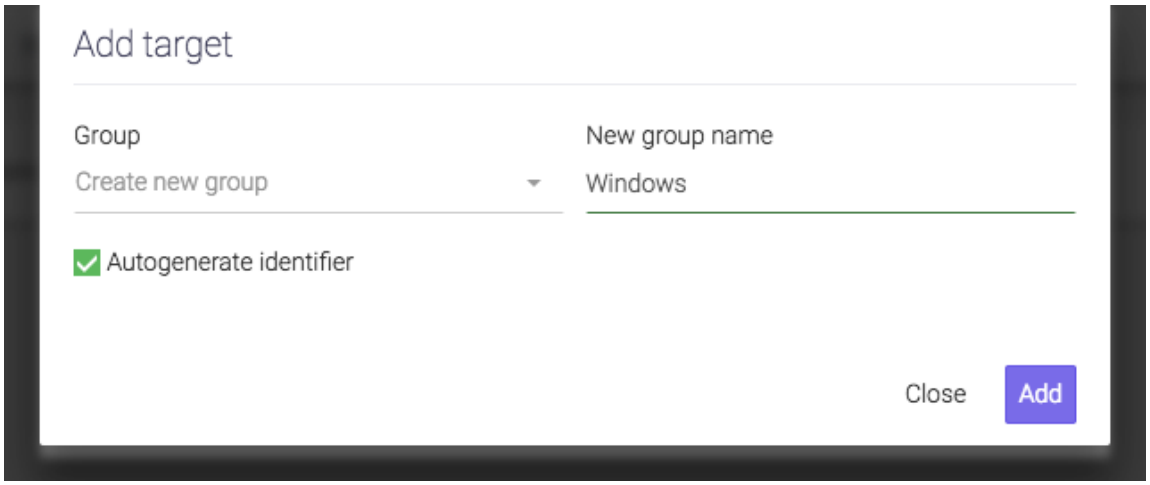

By default the *Autogenerate identifier* option will be selected. Optionally, you can uncheck it and specify an IQN yourself.

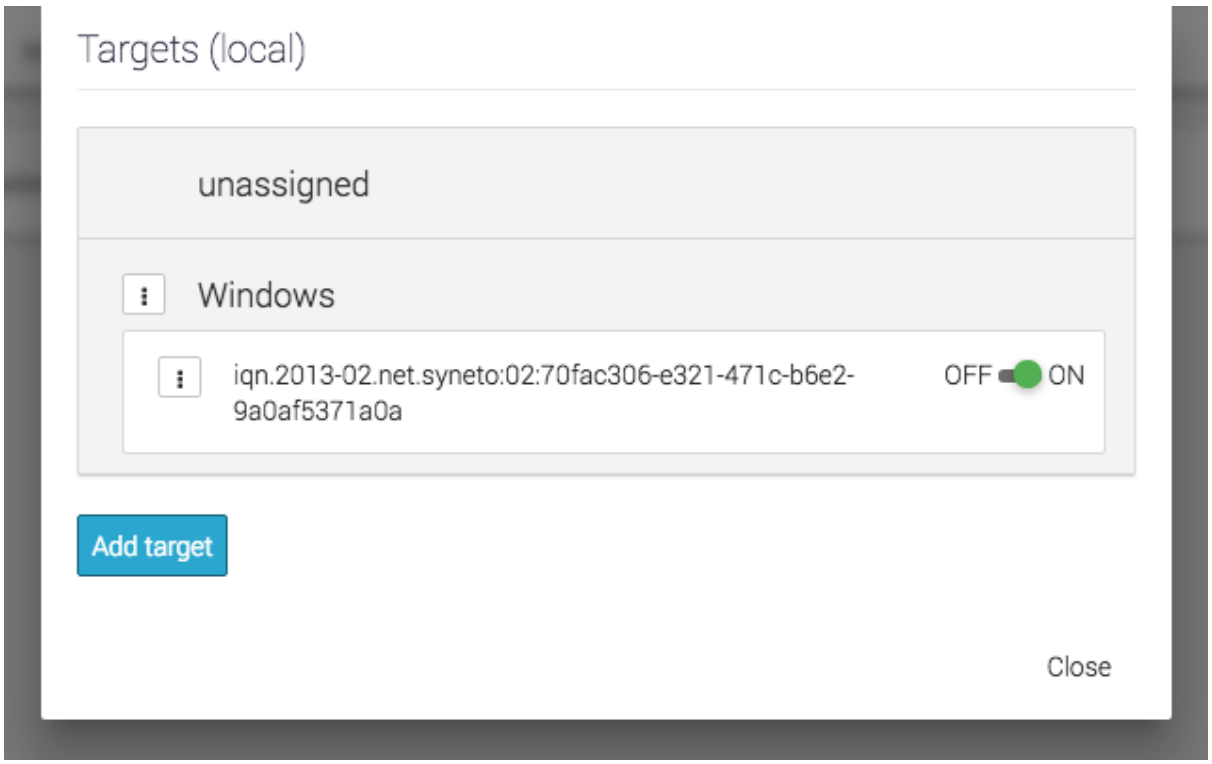

Once you are done, click *Add*. The new target will be show in the list. From this list you can manage the targets, move them from one group the another, turn them on or off, or delete them.

Provisioning and managing

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

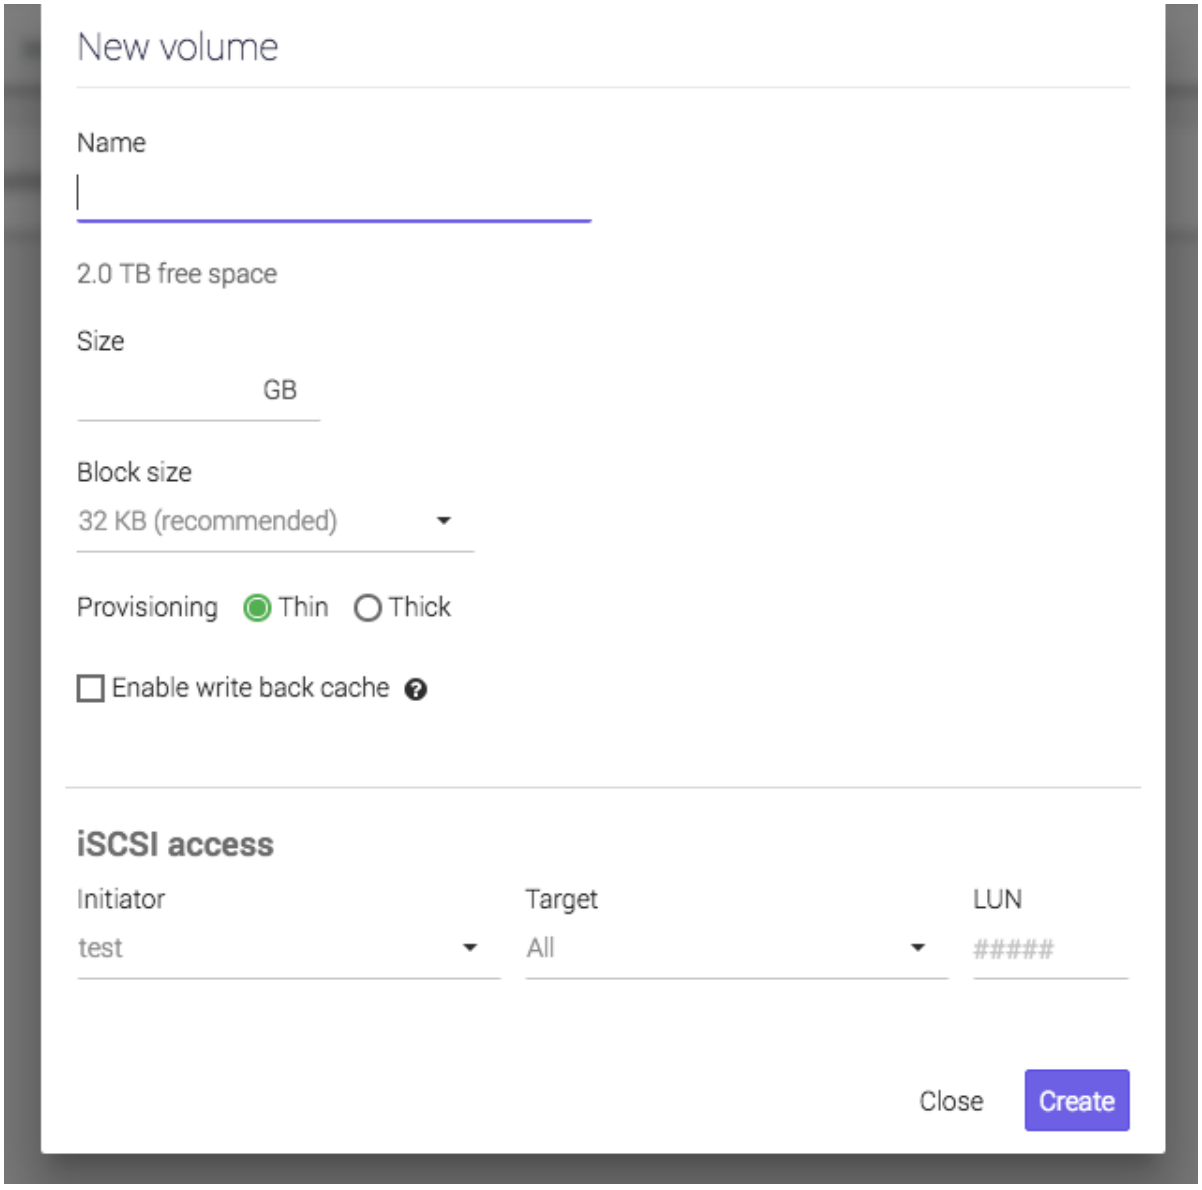

NOTE: An initiator must be defined before provisioning a new volume.

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:

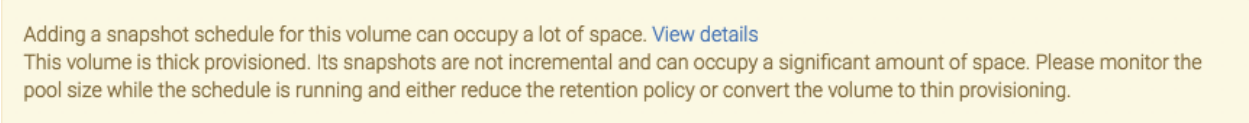

- *Enable write back cache* - enable an iSCSI protocol integration that will receive data to the memory first, confirm to the sender that the data is written, and flush the data to the disk at specific time intervals. This speeds up write operations a lot, but it may lead to data loss at sudden power 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).

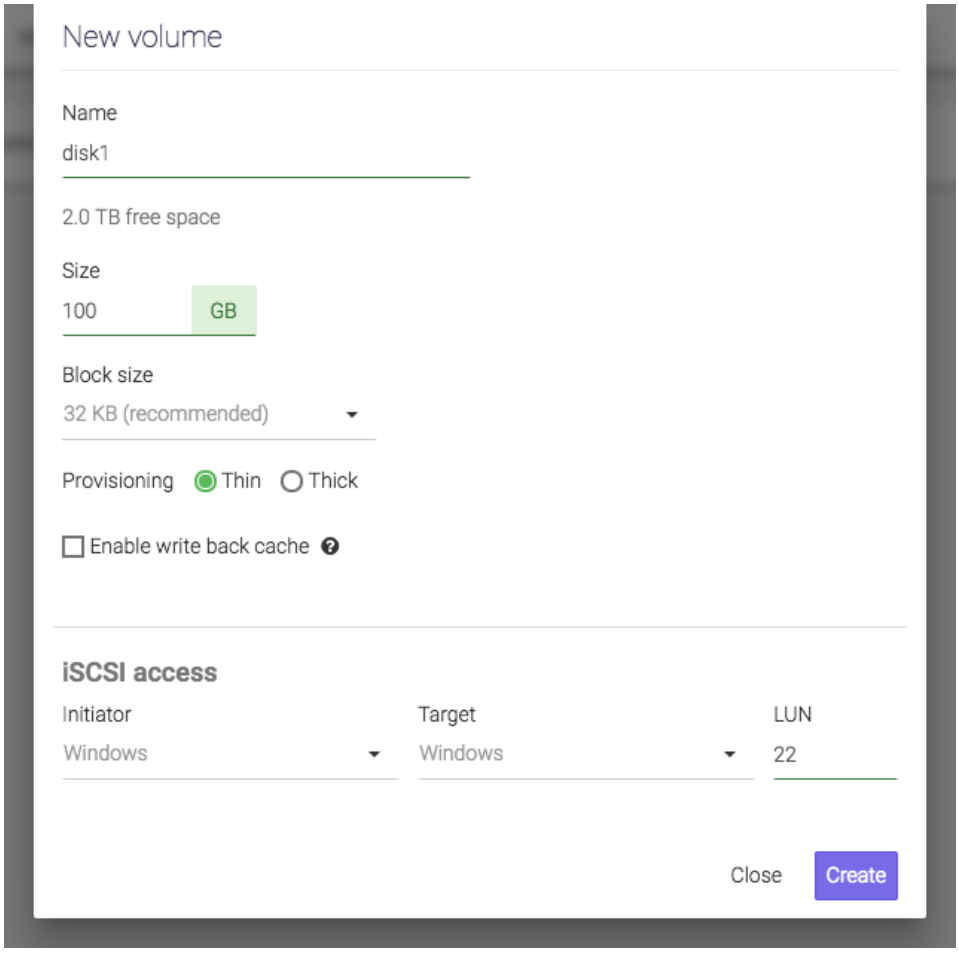

After you are done, click *Create*. The Syneto HYPER will create the local volume, it will reserve the space if you requested thick provisioning, it will set up all the iSCSI configuration for it so that it will be immediately accessible from your network.

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

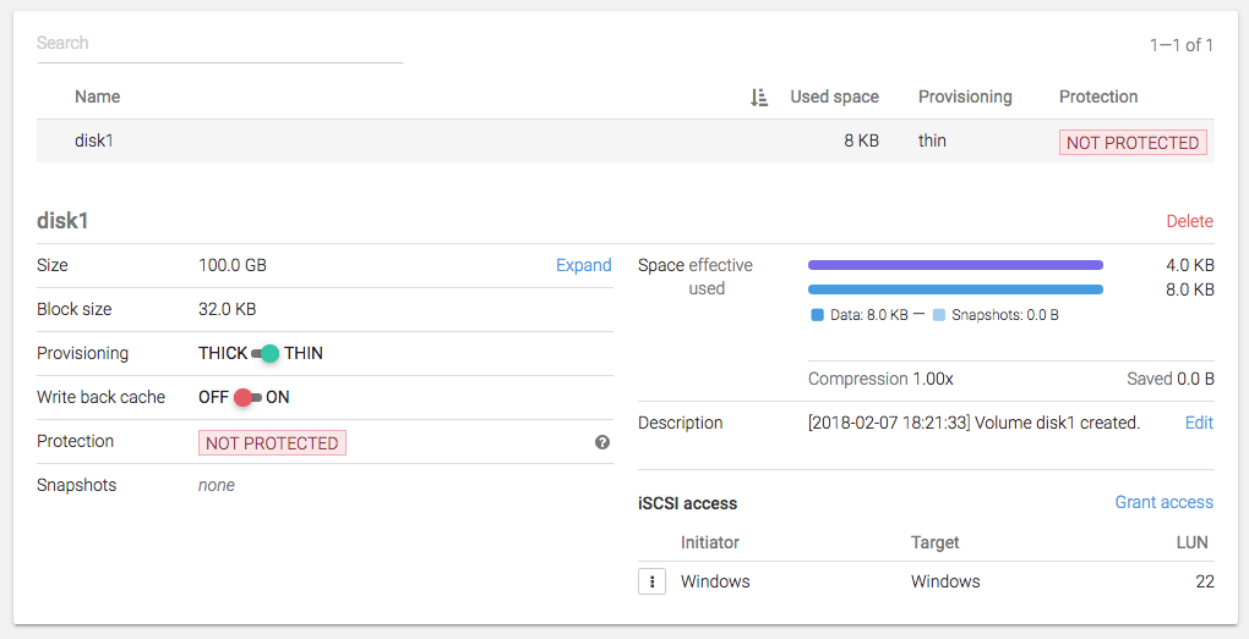

Volumes can be clicked in the list, and a detailed box in the bottom will appear. Here you can take a few volume management steps:

- *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.

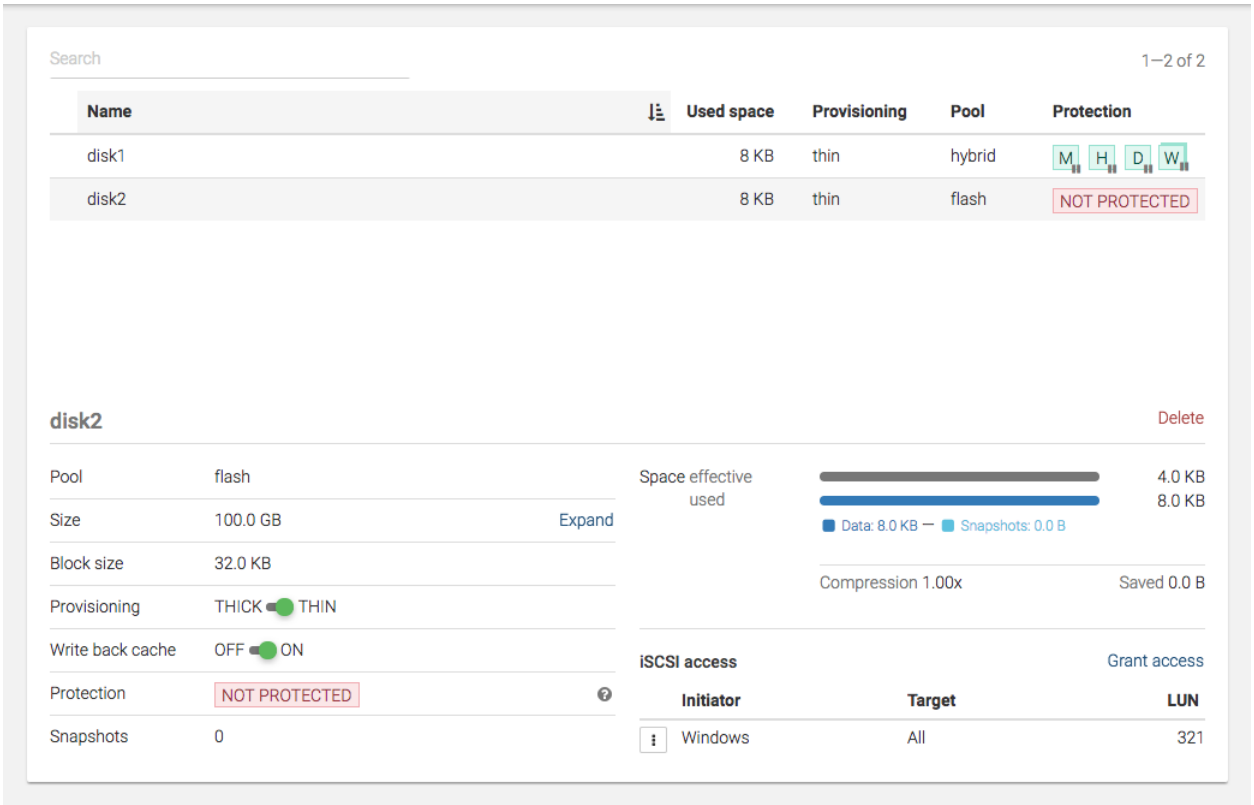

The list has two columns that can help you monitor what is happening: Used space and Provisioning

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*.

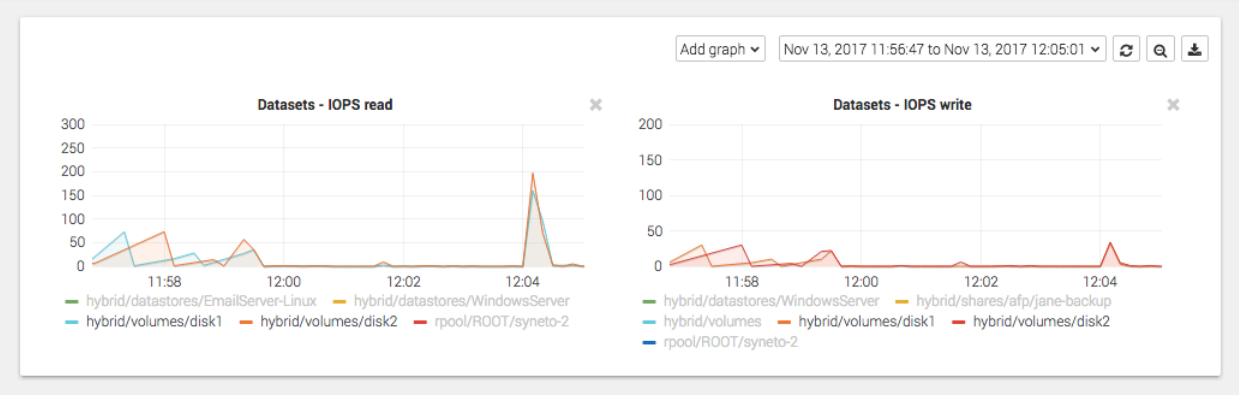

These graphics show the IOPS happening on the volume from the Syneto HYPER perspective.

Dataset bandwidth

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

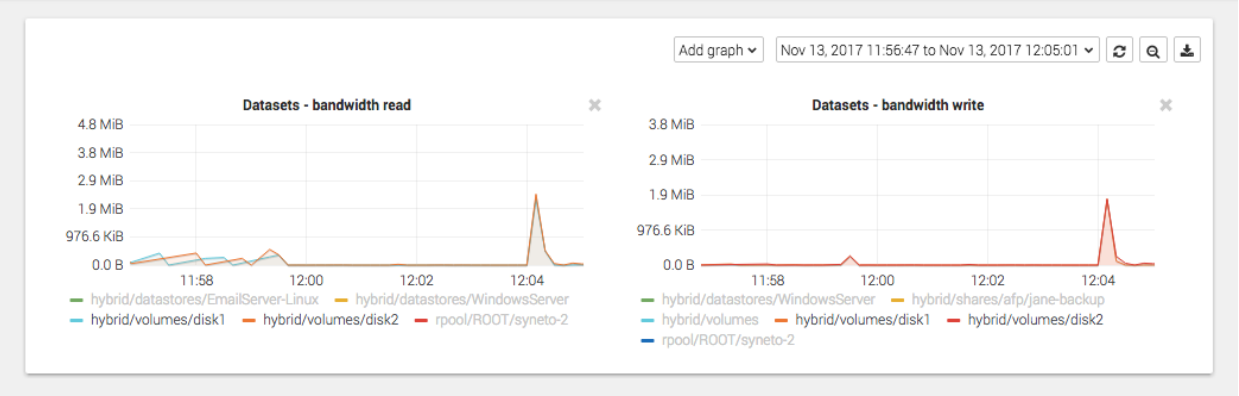

These graphics show the data transfer speed on the volume from the Syneto HYPER's perspective.

Dataset average block size

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

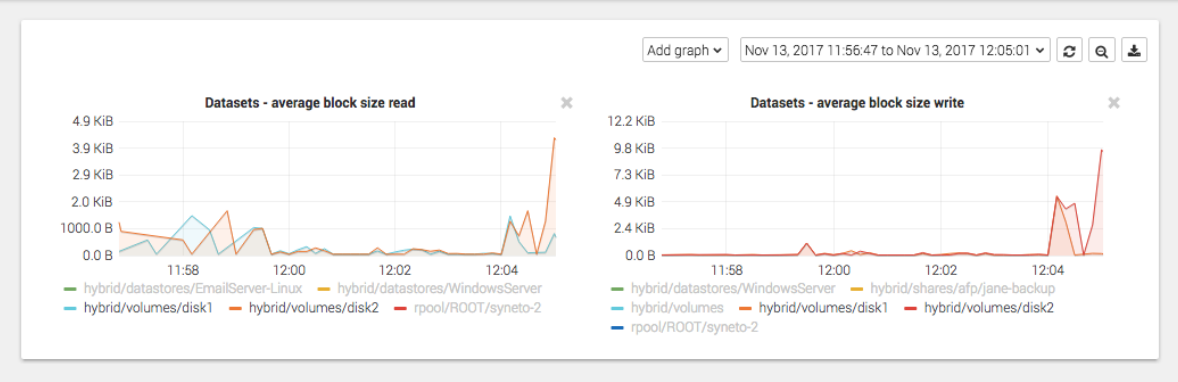

These graphics show the average block size of the data on the volume from the Syneto HYPER's perspective.

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.

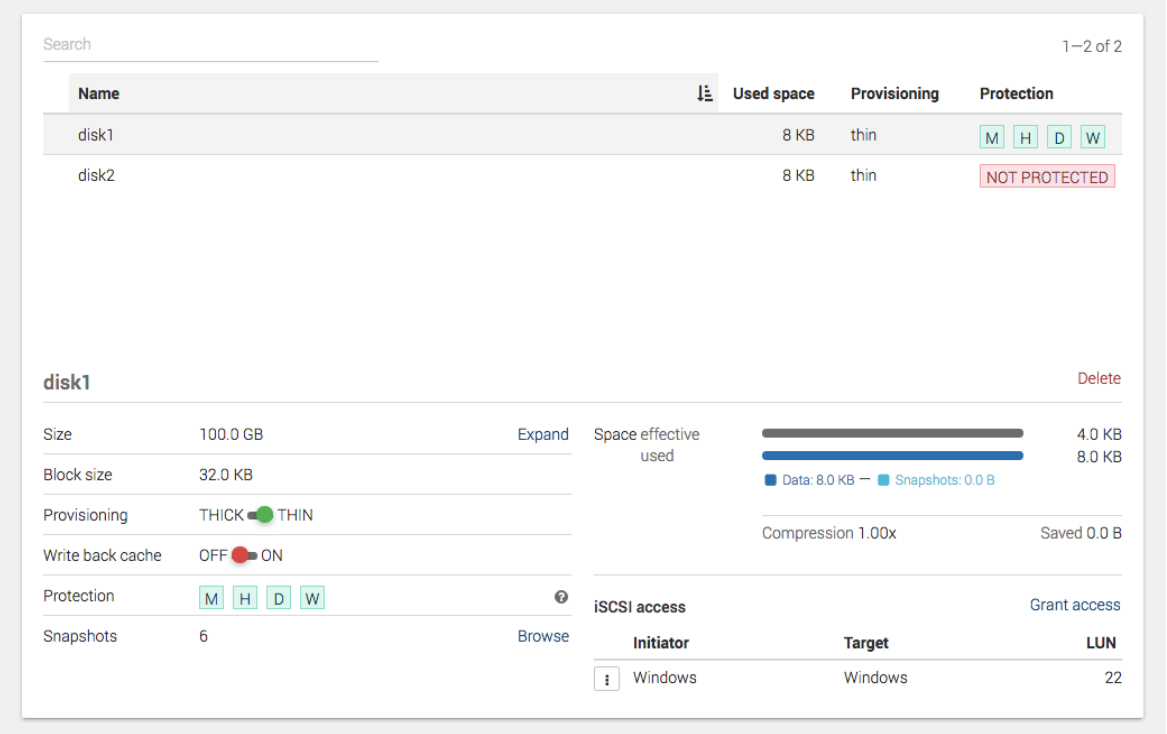

For all these, and many more cases, you can simply click *Browse* in the snapshots section of the detailed view of a volume.

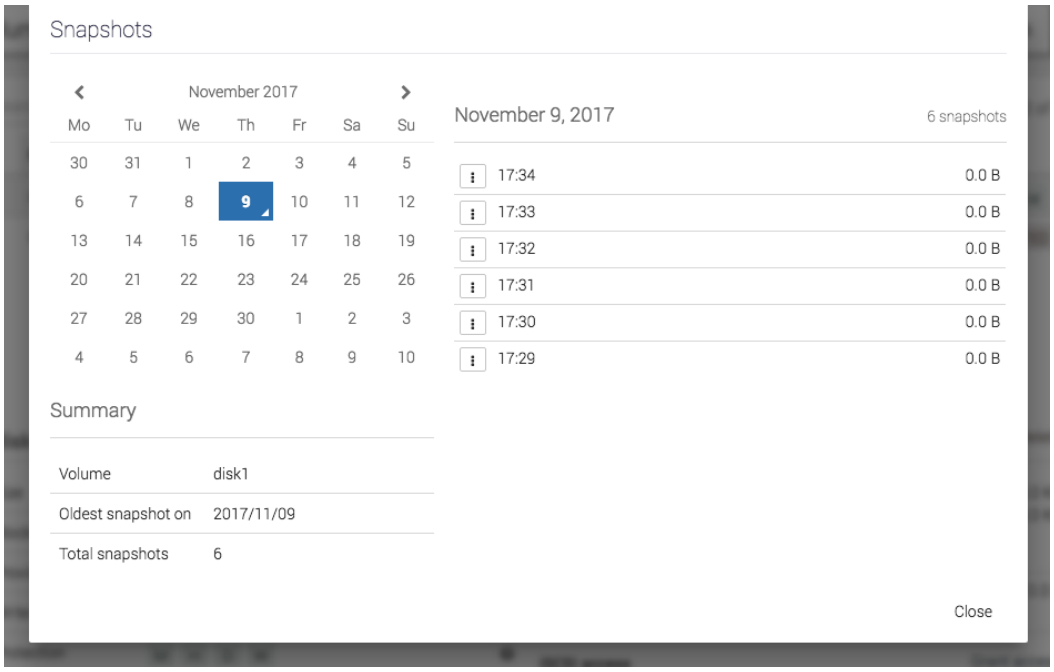

Pick a date from the calendar, and click on the contextual menu for a snapshot.

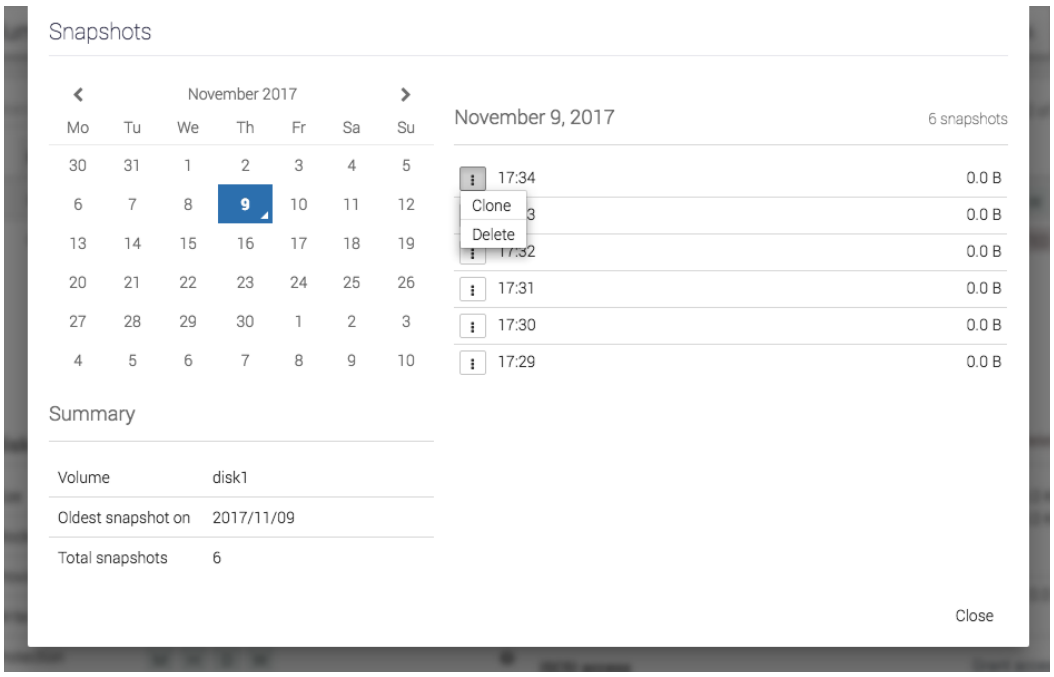

Select *Clone*.

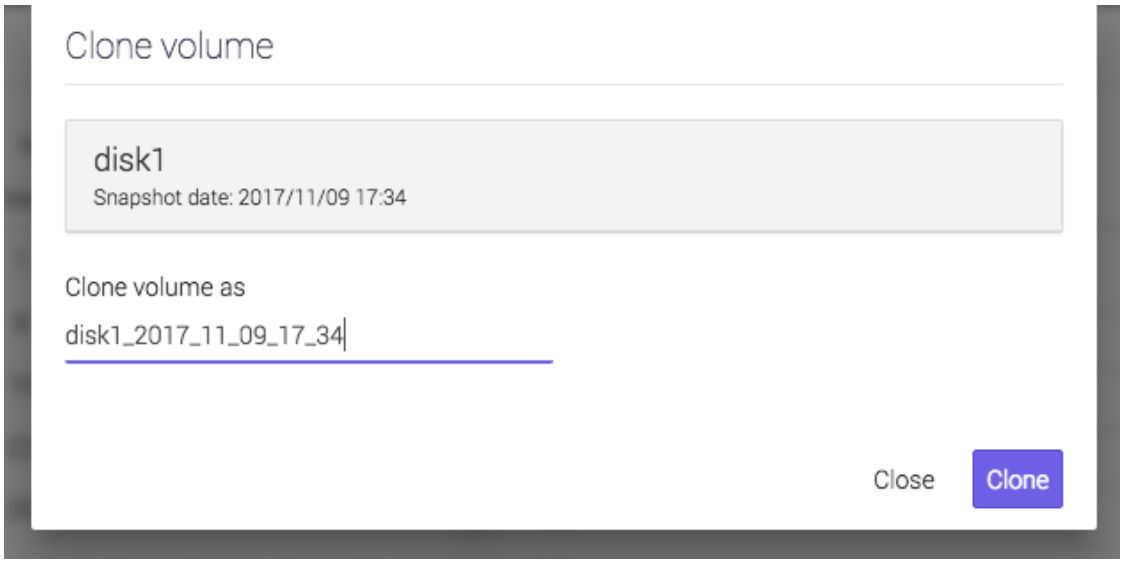

Click clone and wait for the process to finish.

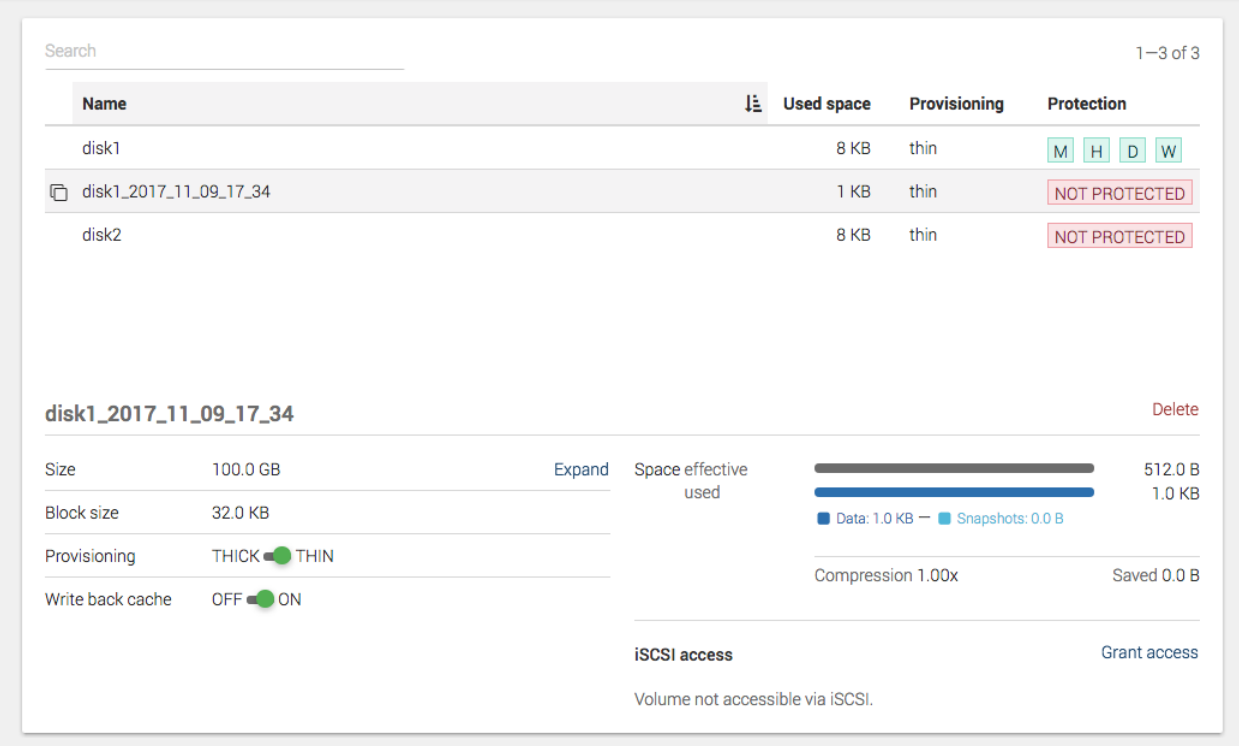

Now you can access a volume with the name specified at the clone step. It will contain the data as it was present at the moment of taking the snapshot.

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.

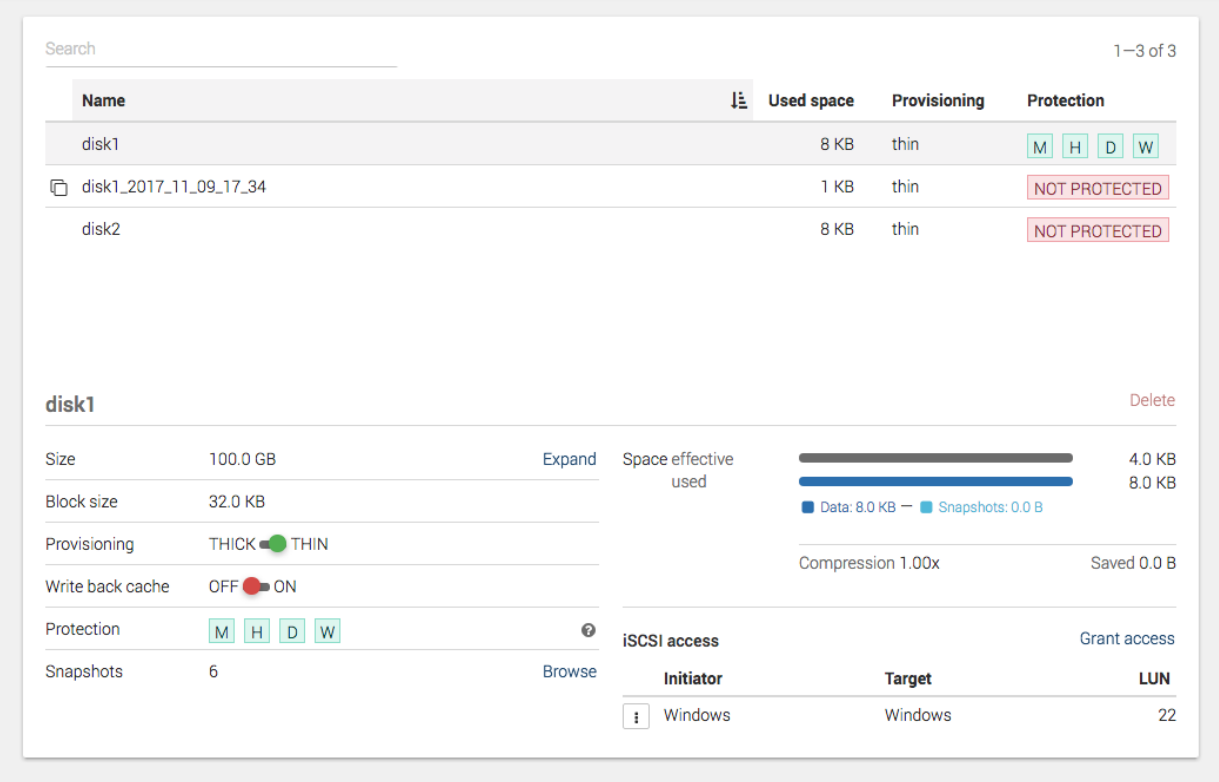

You will see that all cloned snapshots are marked with a label. Click the context menu for a cloned snapshot and select *Clean up clone*.

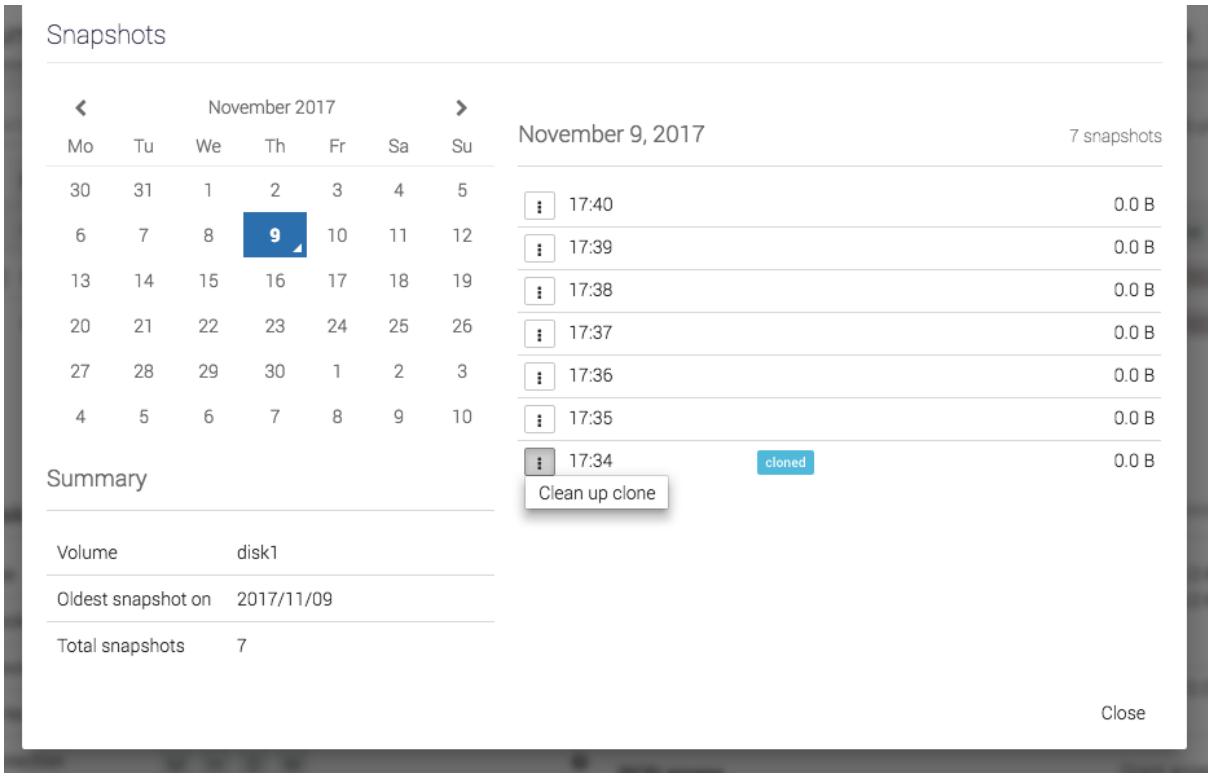

The system will ask you to confirm the process.

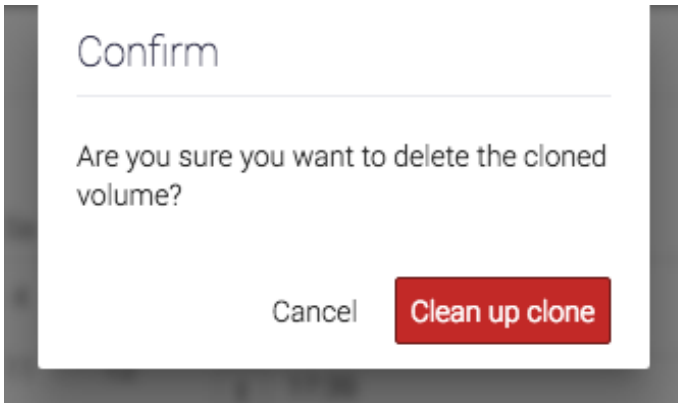

Wait a few seconds for cleanup to finish. A green notification in the bottom right will tell you when it's done.

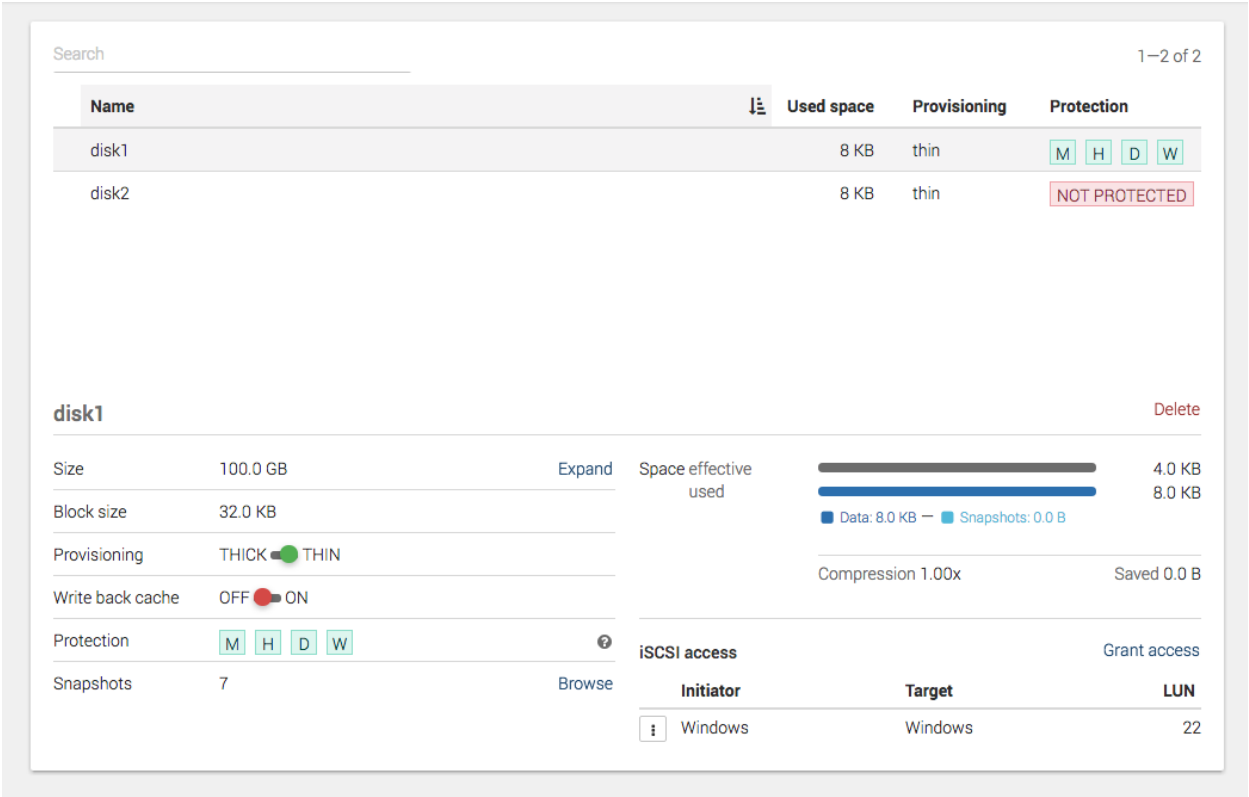

That's it, all that was created and shared for that snapshot clone was reverted. No more mess left behind by the recovery process.

DISKS AND POOLS

Introduction

In the secondary menu, we have the *Hardware* section. This contains *Disks* and *Pools*.

On the *Disks* page you can monitor the health of the disks and replace a faulted disk. By clicking on a disk, more details will be shown.

Pools are used to group disks 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.

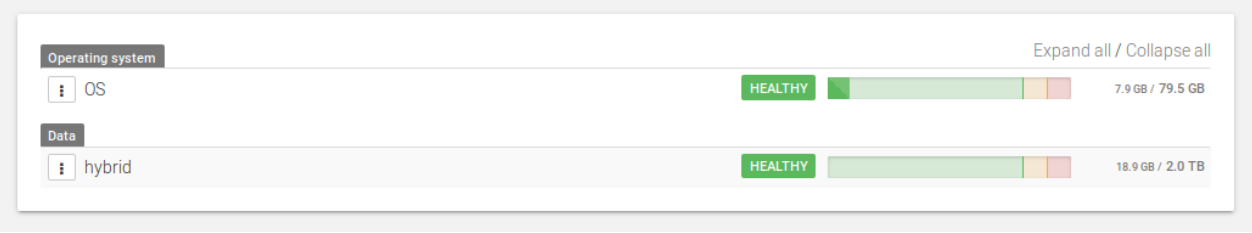

Replacing a Faulted disk

Go to *Secondary Menu* → *Disks*.

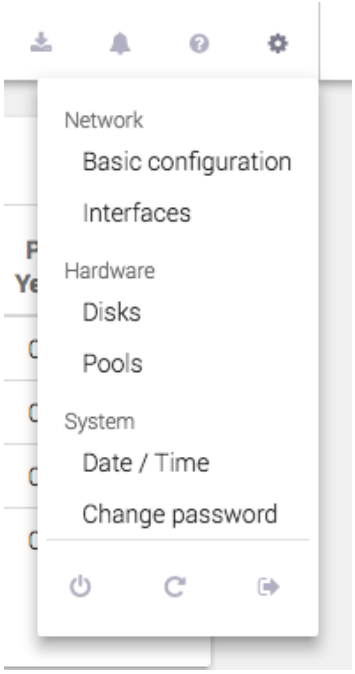

The *Disks* page will open. The faulted disk will be highlighted

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

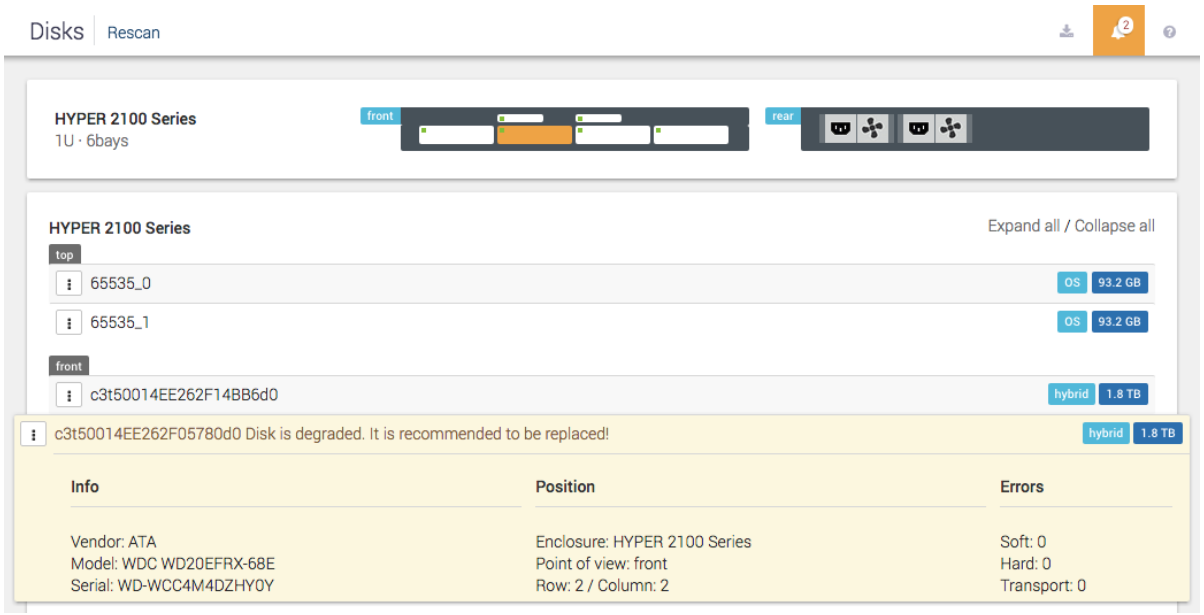

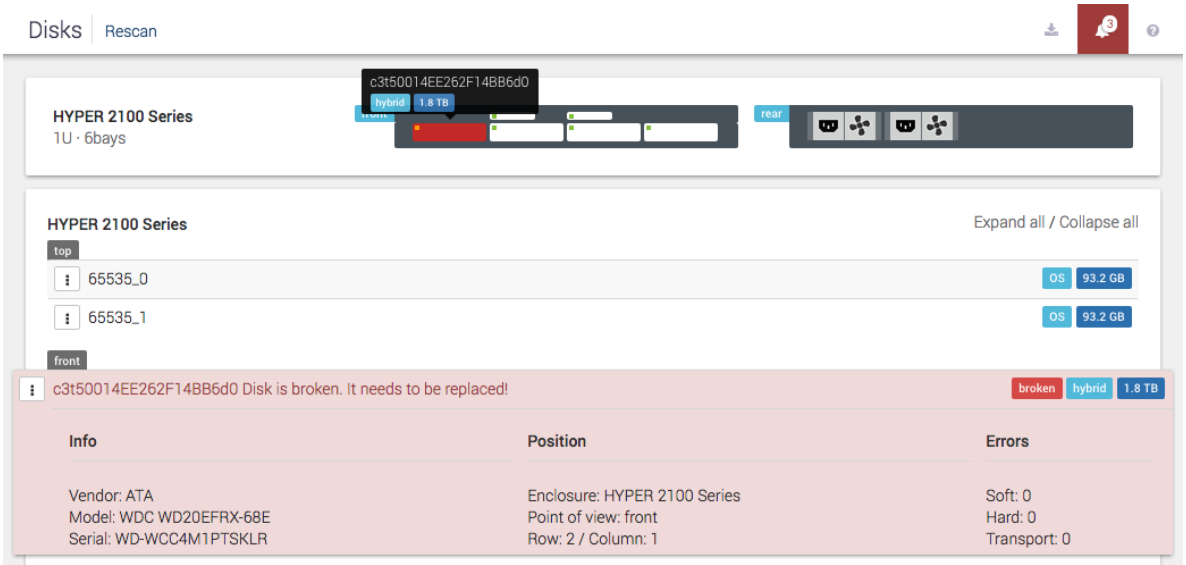

Click on the disk for more details. See in which bay it is placed. Our machines have support for blink, so you can detect the disk easily. To activate the blink for a disk, click on *Context Menu → Blink*.

NOTE: OS disks do not have support for blinking.

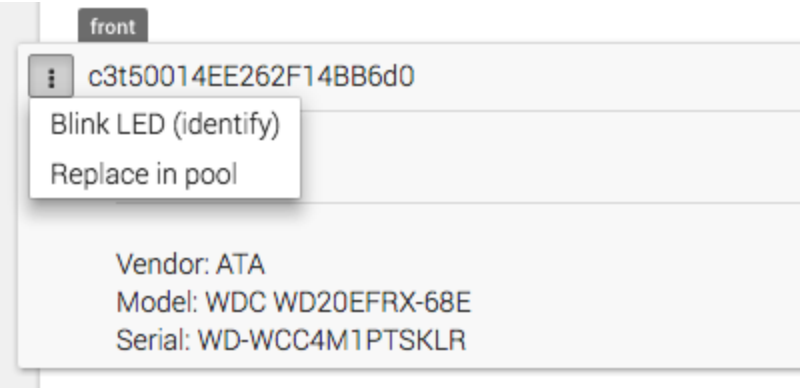

Go to the physical machine. Identify the faulted disk, take the bay out. Replace the disk in the bay with a new disk.

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

Place the bay in the machine.

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

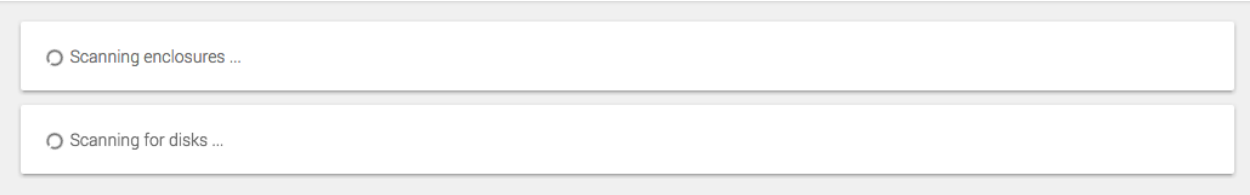

Click on *Context Menu* → *Replace in pool*.

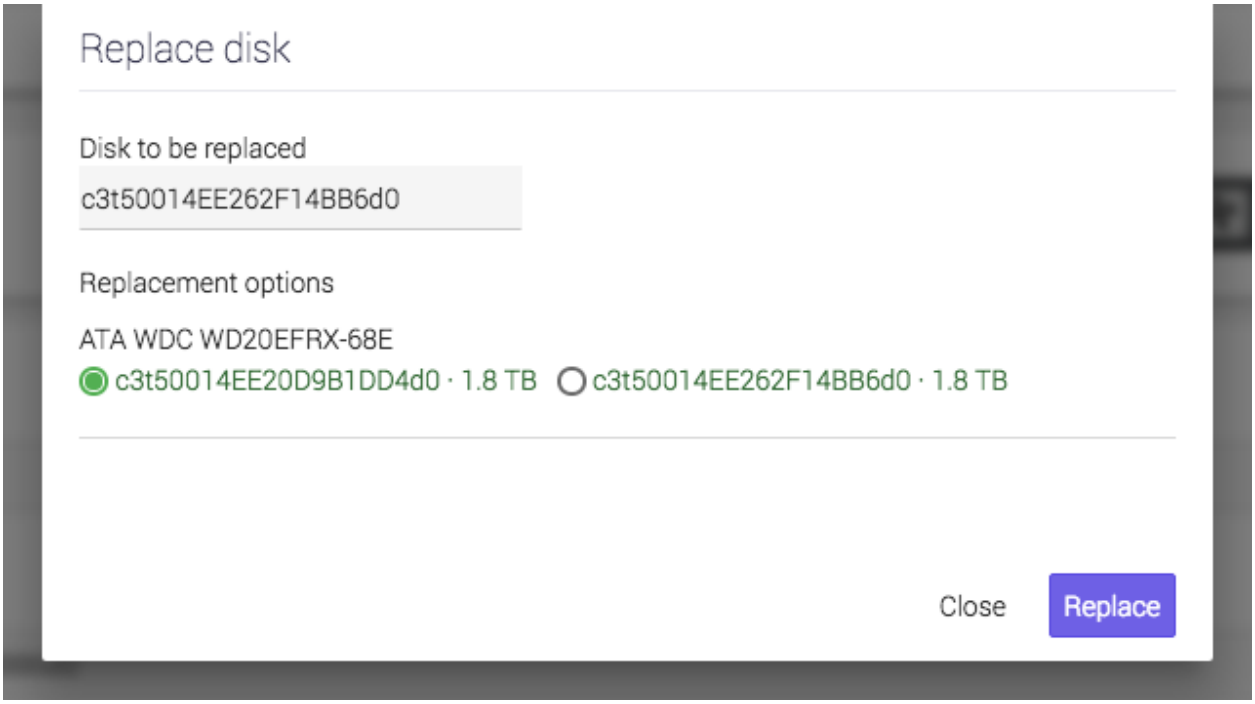

There will be a list with the available disks for replacement. Select the newly introduced one and click *Replace*. Now, the faulted disk is being replaced with the new disk. The system will automatically start restoring the data on the new disk.

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

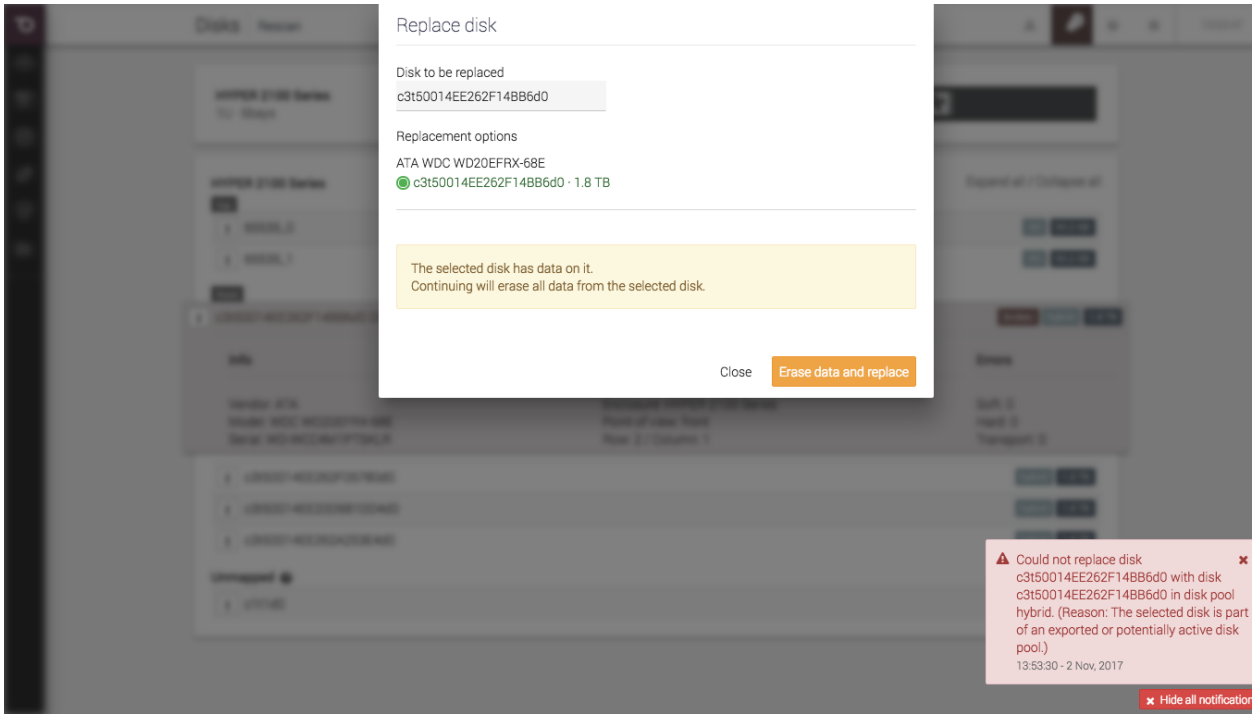

WARNING: If you are sure there is no important data on the disk, click on Erase data and replace.

You can see the building redundancy progress on the *Pools* page.

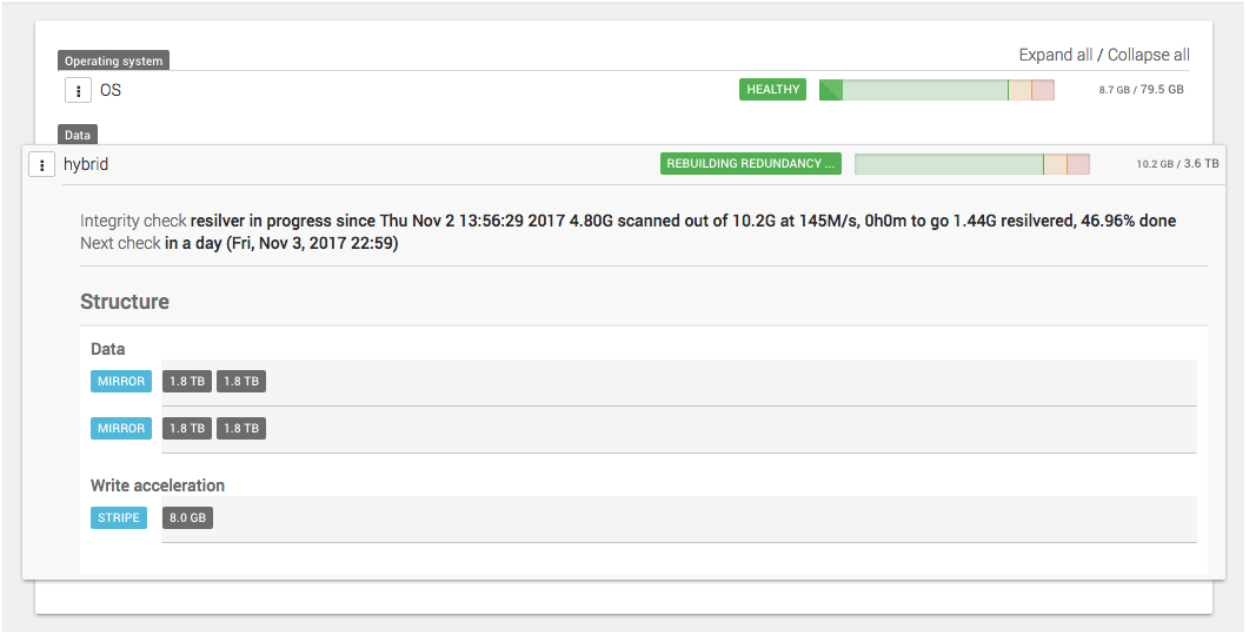

Replacing a Faulted OS Disk

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

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

Replacing a Faulted cache disk

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

Attaching an expansion shelf

If the space on your machine is not sufficient for your needs, you can always expand it using a JBOD.

NOTE: Hyper2000 series and DR-Play do not support expansion.

Connect the JBOD to the Syneto machine. Go on *Disks* page and click *Rescan.* The new disks will appear on the page with the tag *available*.

You can either create a new pool or expand an existing one with the newly available disks.

WARNING: Disks in JBOD need to be mapped manually.

Creating a pool

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

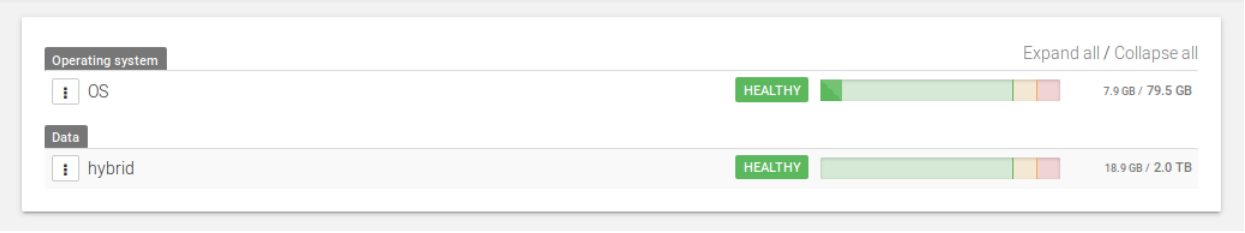

Click on the *New Pool* button in the action bar, top of the page. A dialog for creating a new pool will open.

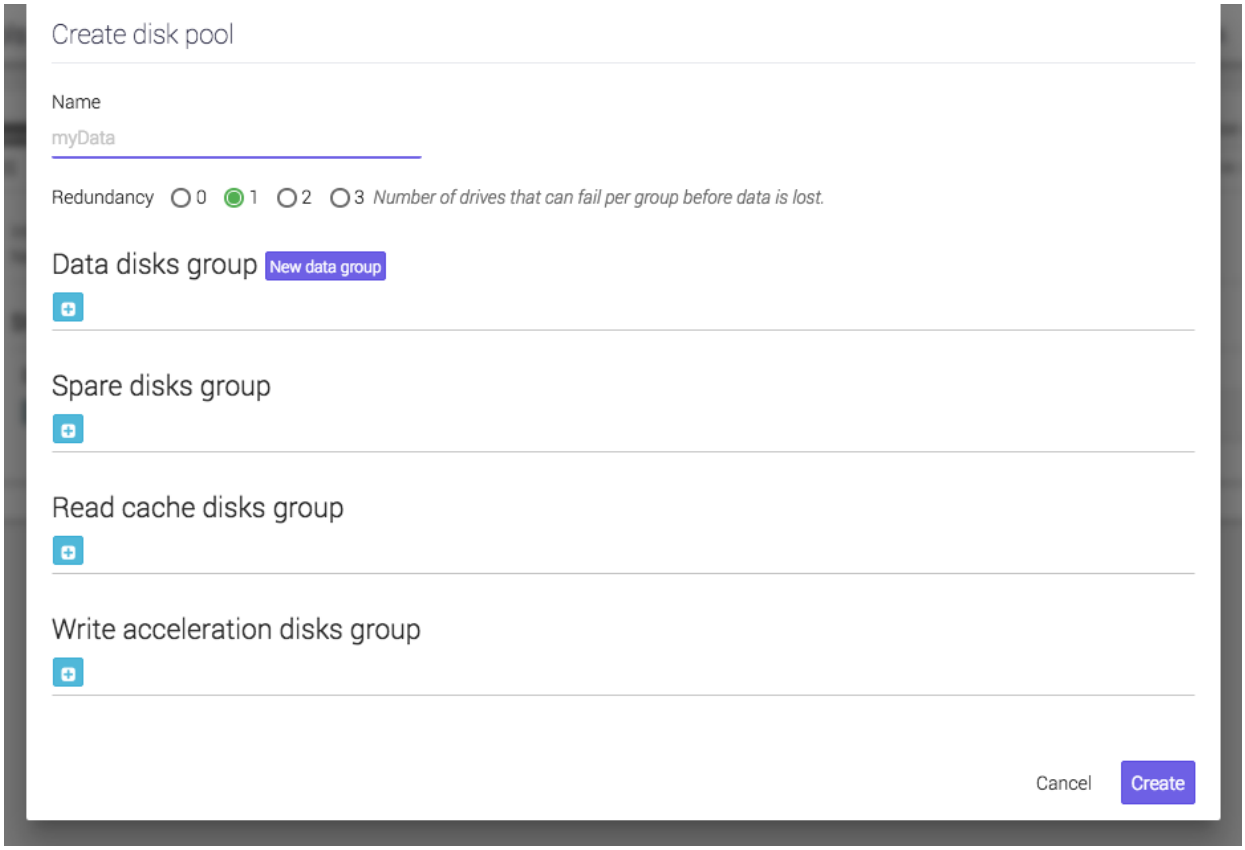

You will have to give a name to the pool.

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

- *Data disk groups* – disks that contain the actual data. Each data disk group has a number of disks that may fail before losing data. Choosing data disk groups that will support disk failures will result in lower available space on the group – depending on the number of disks that may fail (e.g. if a data disk group of 5 disks will allow 2 disks to fail before losing data then the total available space on the disks group will be reduced by 40%).

  Recommended: Select *redundancy* value 1 and two *disks* per *data group*.
- *Spare disk group* – the disks in this group are not actively used but will automatically replace a data disk that has failed. Adding spare disks is optional.
- *Read cache disk group* – this group contains disks that are actively used to speed up the data reading process. It is recommended to use high speed disks (SSD) in this group. Adding disks to this group is optional.
- *Write acceleration disk group* – this group contains disks that are actively used to speed up the data writing process. It is recommended to use only high speed disks (SSD) in this group. Adding disks to this group is optional.

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

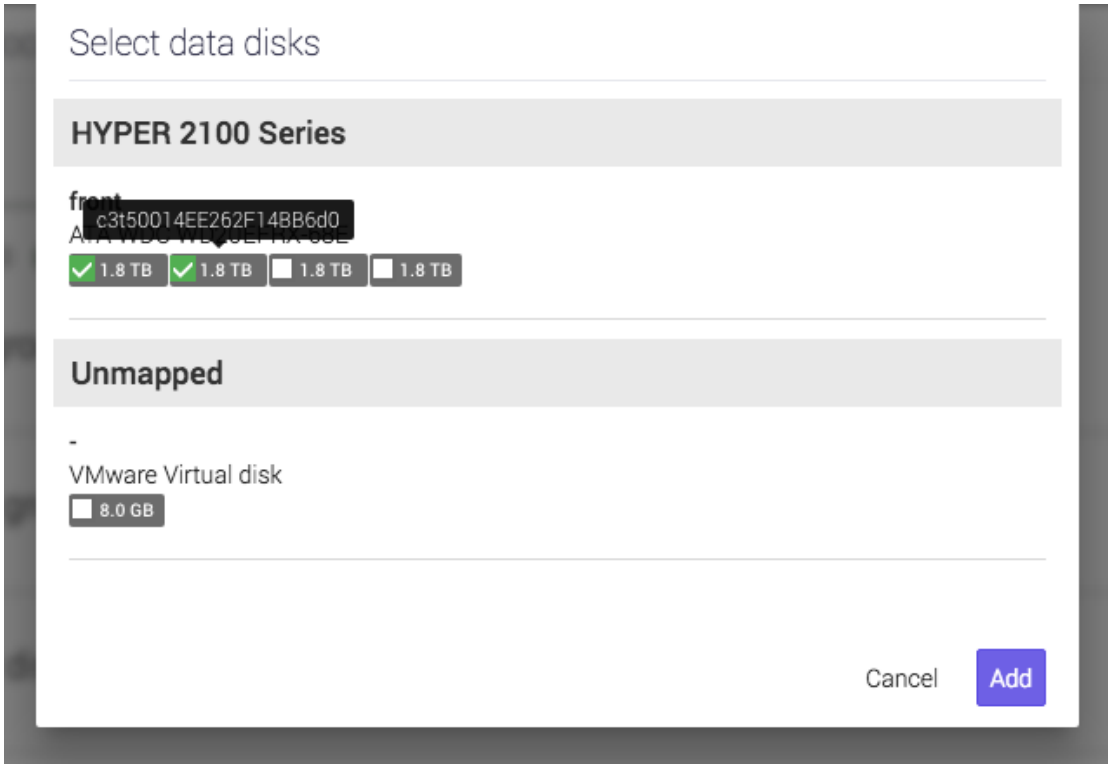

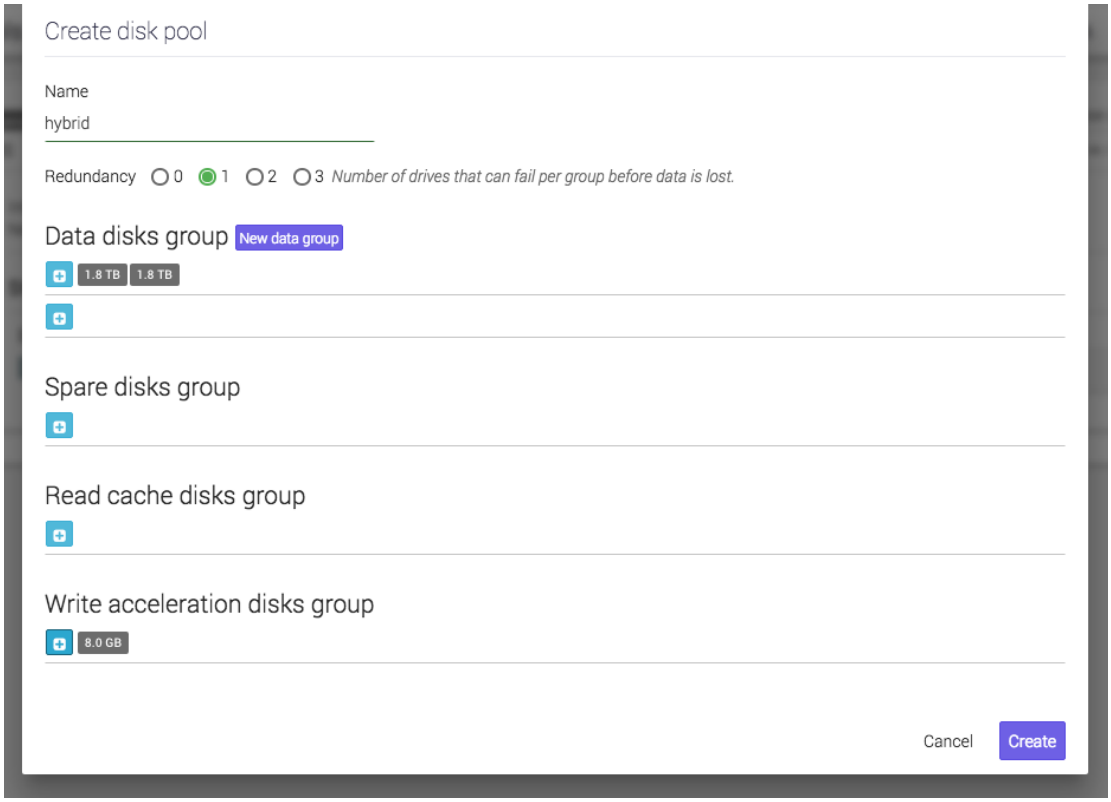

Expanding a pool

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

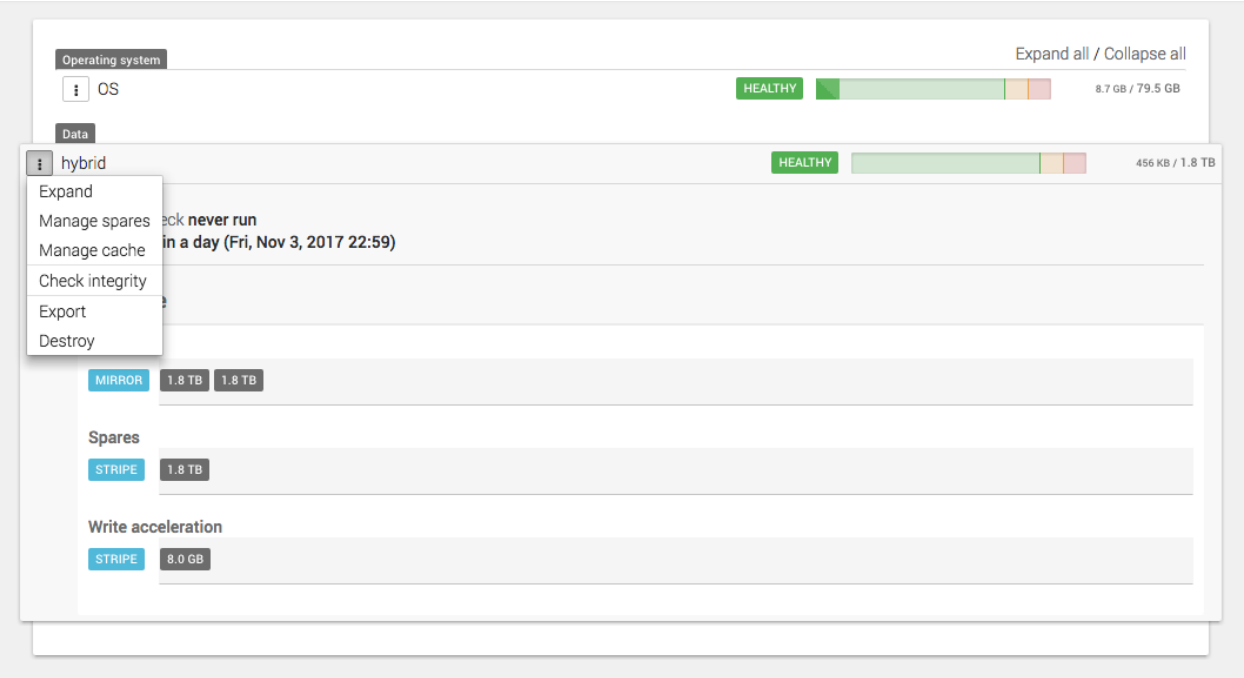

This will only expand the data disks.

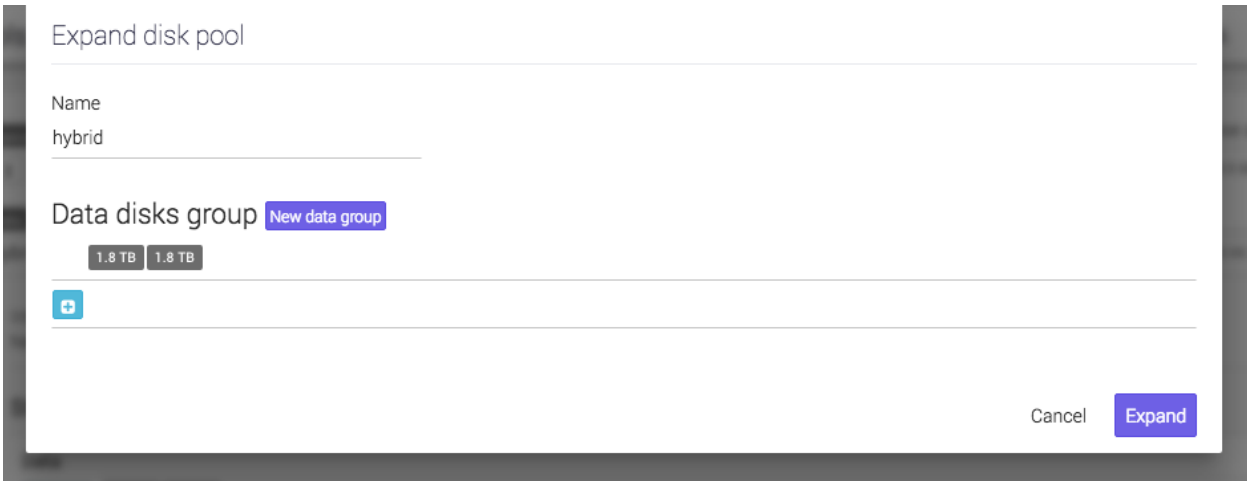

For adding disks, click the *Plus* button.

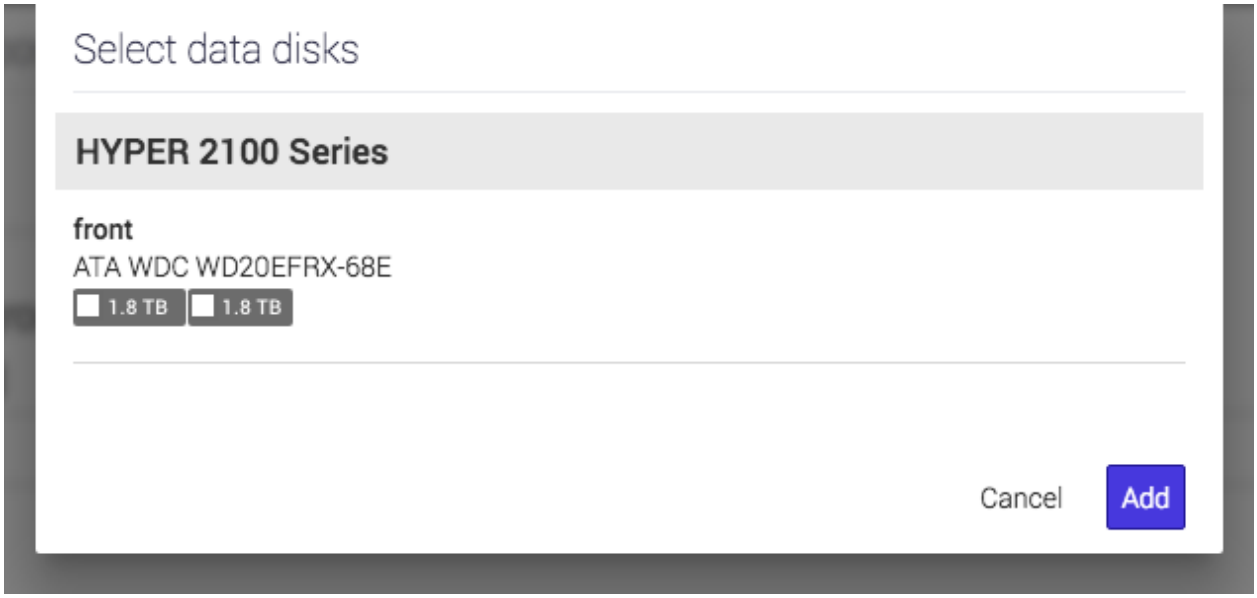

Select the ones you want by clicking on them.

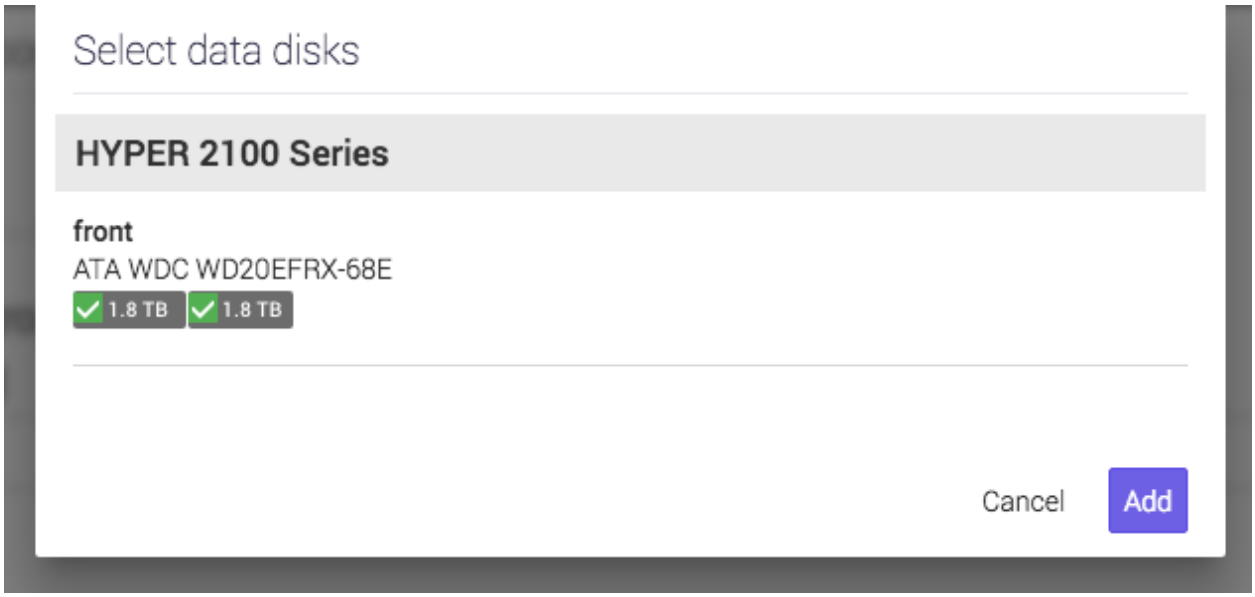

After you finished selecting, click *Add*.

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

To expand a pool with *spares disks*, click on *Context Menu → Manage spares*.

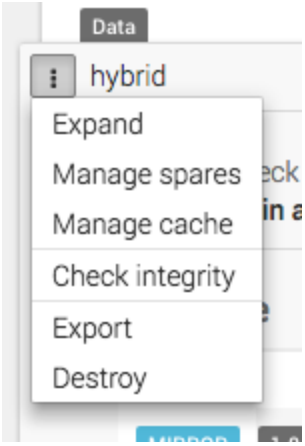

A dialog for adding *spare disks* will open.

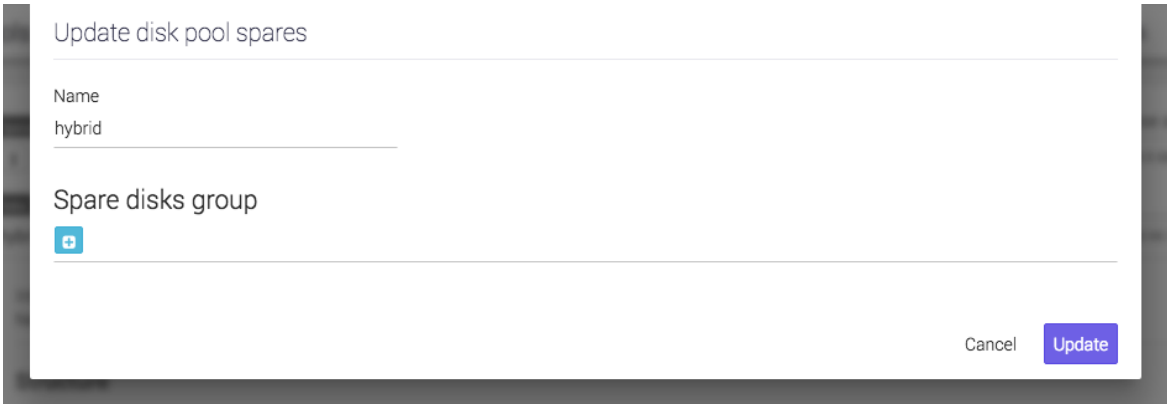

Click on the *Plus* button. A dialog will open where you can select which disks to be used.

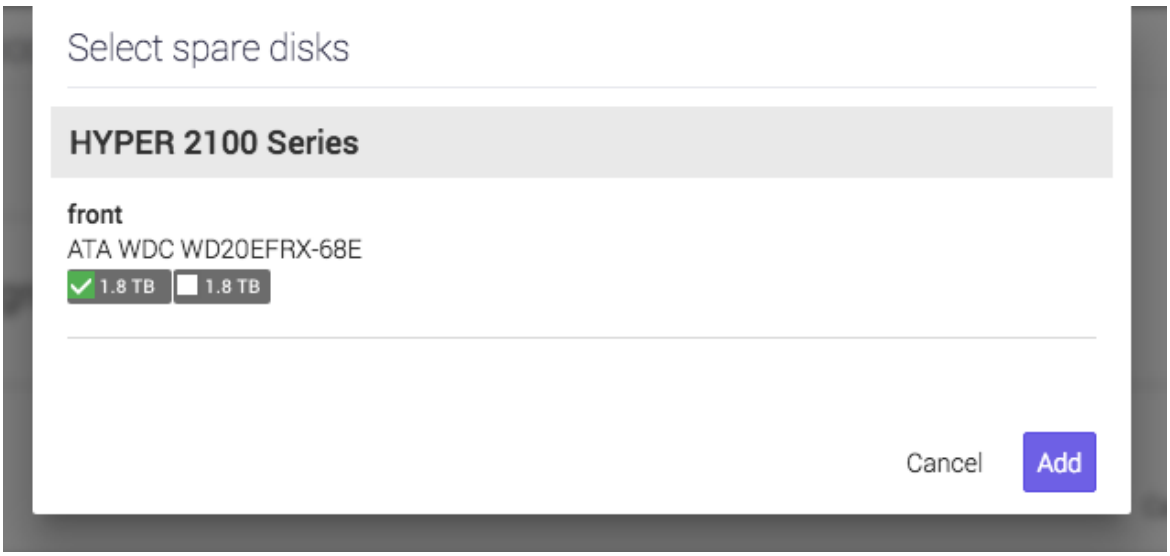

After selecting the disks you want click *Add*.

To expand a pool with *cache disks*, click on *Context Menu → Manage cache*.

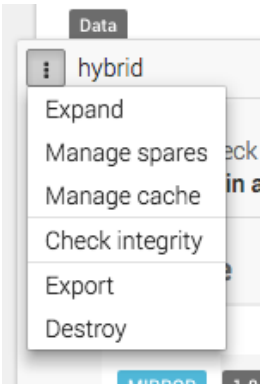

A dialog will open where you can add *read cache disks* and *write acceleration disks*.

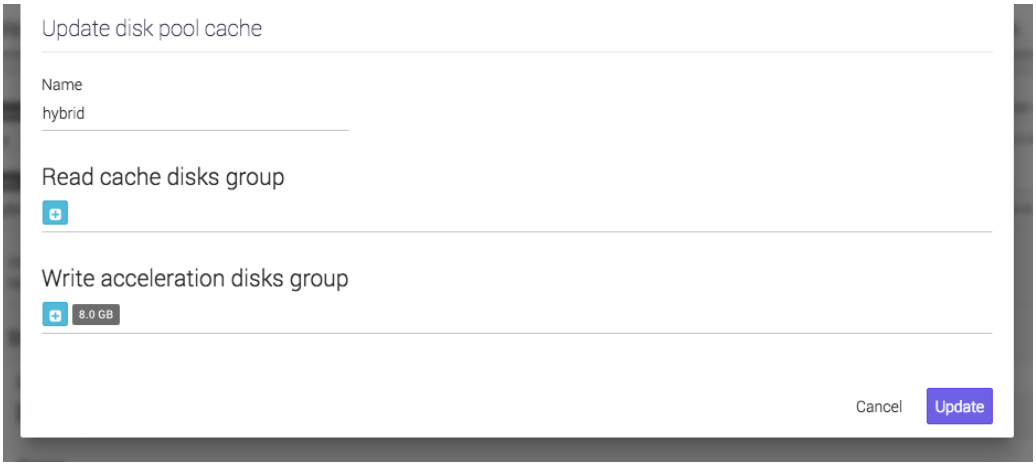

Click on the *Plus* button. A dialog will open where you can select which disks to be used.

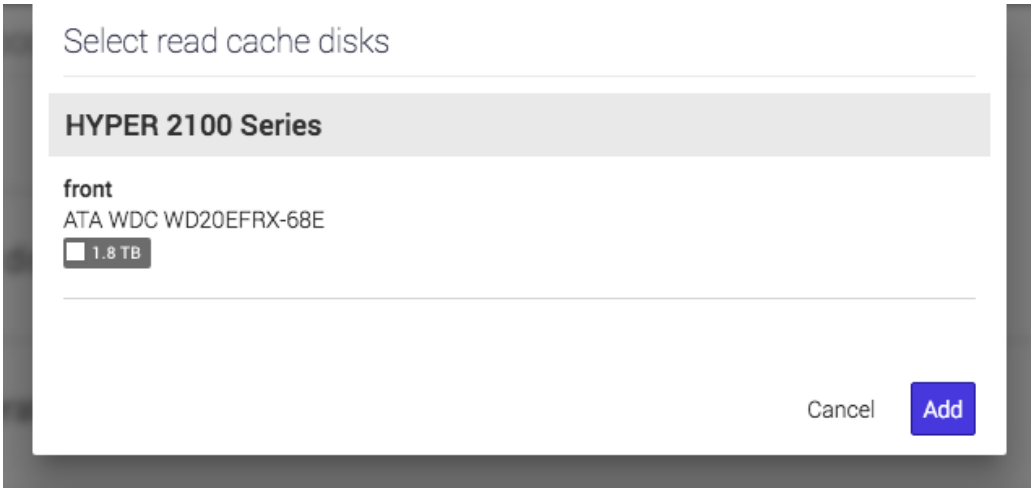

After you finish your selection, click *Add*.

Importing a disk pool

If you reinstall SynetoOS or you have introduced disks 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.

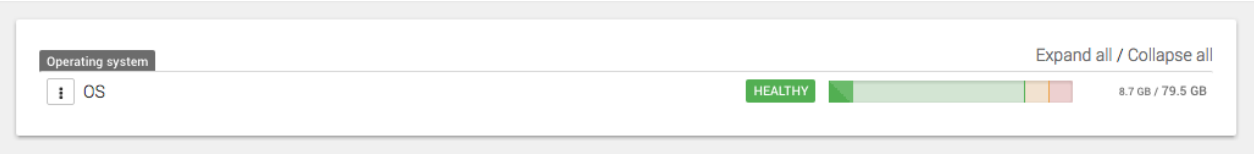

Click *Import existing* in the action bar. If there are disk pools that can be imported a list will be shown.

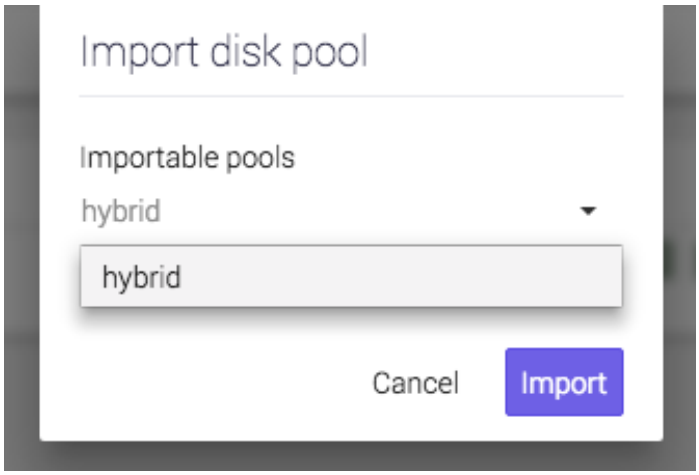

Select the disk pool to be imported and click *Import*.

If a disk pool was previously used in another machine, a warning will appear. You will need to force import it. This option will be offered to you when clicking Import.

Exporting a disk pool

If you reinstall SynetoOS or you want to move disks 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.

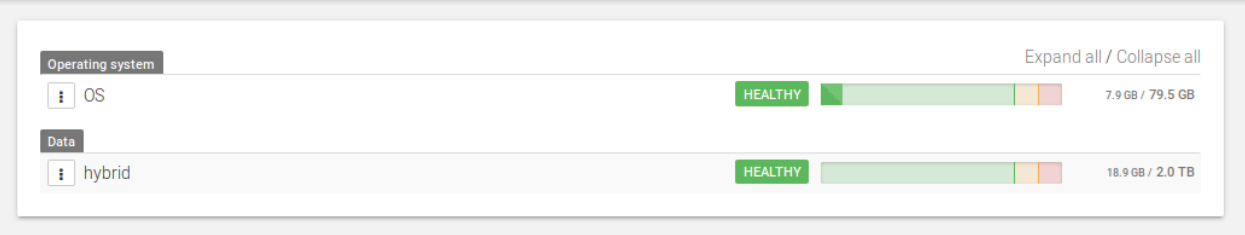

To export a pool, click on *Context Menu* → *Export*.

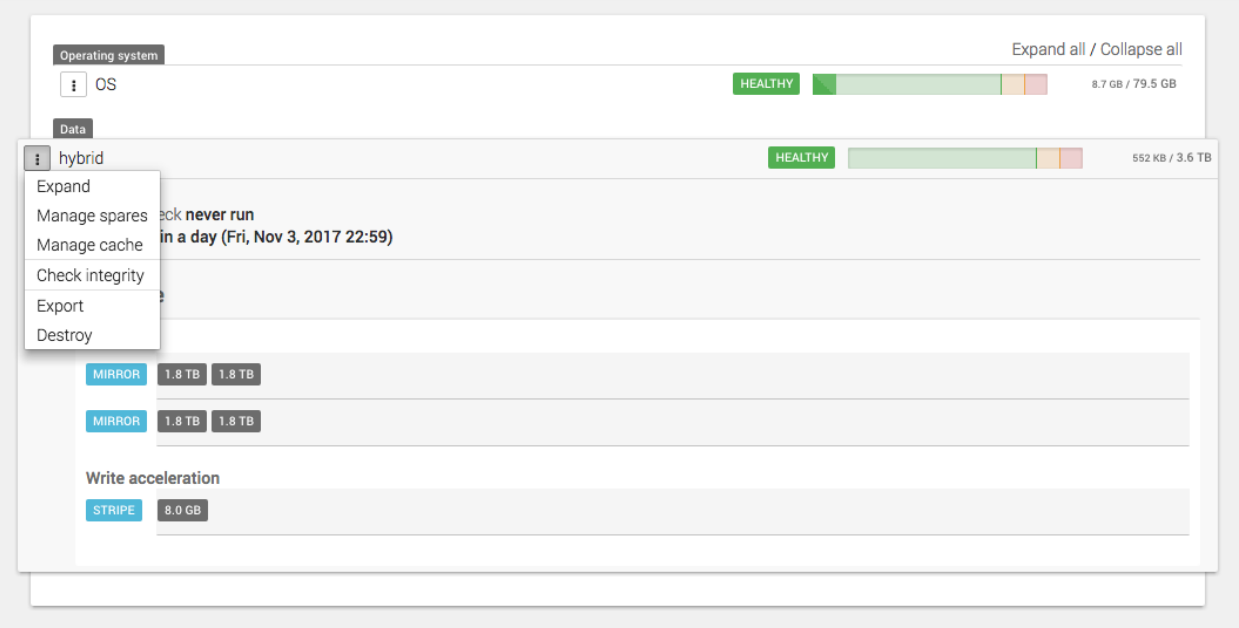

NOTE: The OS pool can not be exported.

WARNING: All virtual machines and datastores related to this pool from all connected ESX 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.

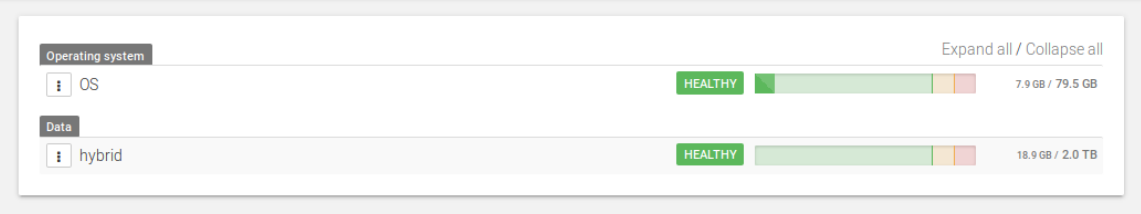

To set *integrity check*, click on *Context Menu → Check Integrity*.

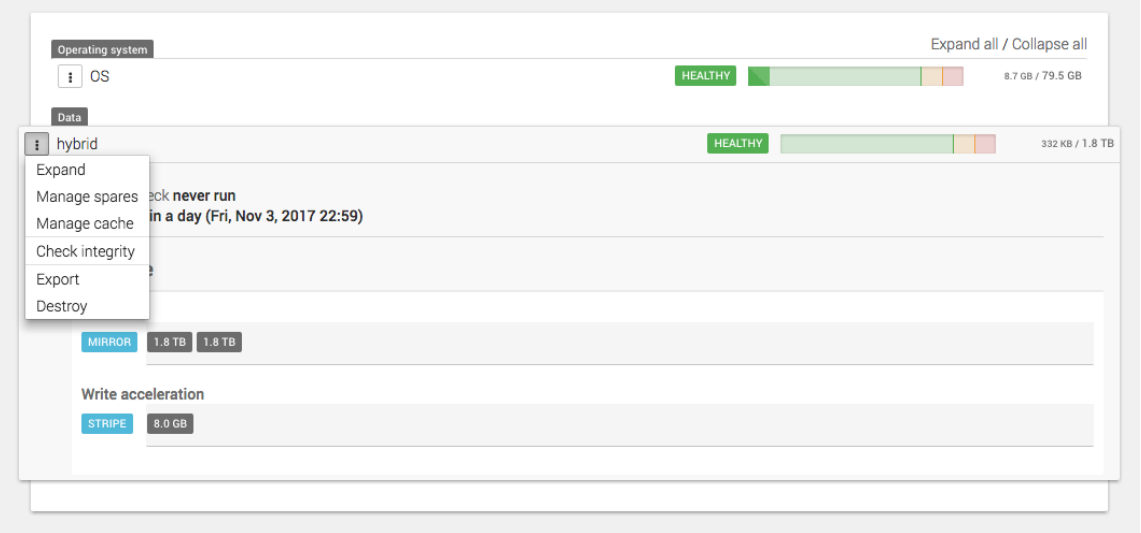

A dialog will open with different configuration options for integrity checking.

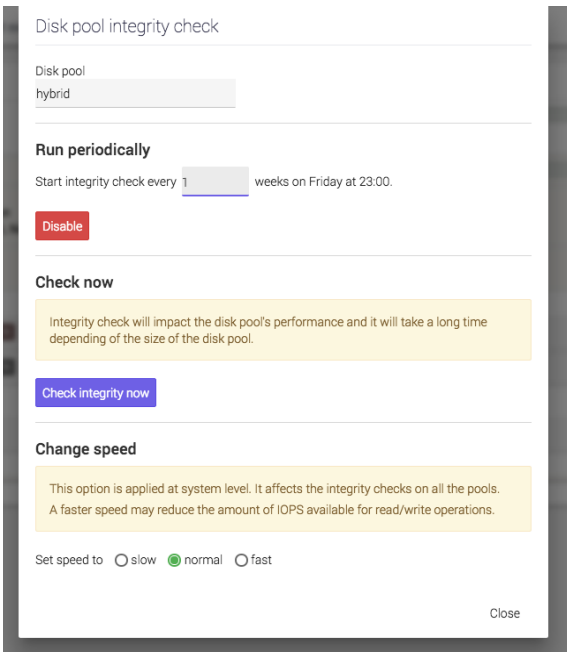

Run periodically section allows you to set a schedule. By default, it is enabled every week on Friday, at 23:00. You can change the week interval. First, you have to *Disable* the current schedule.

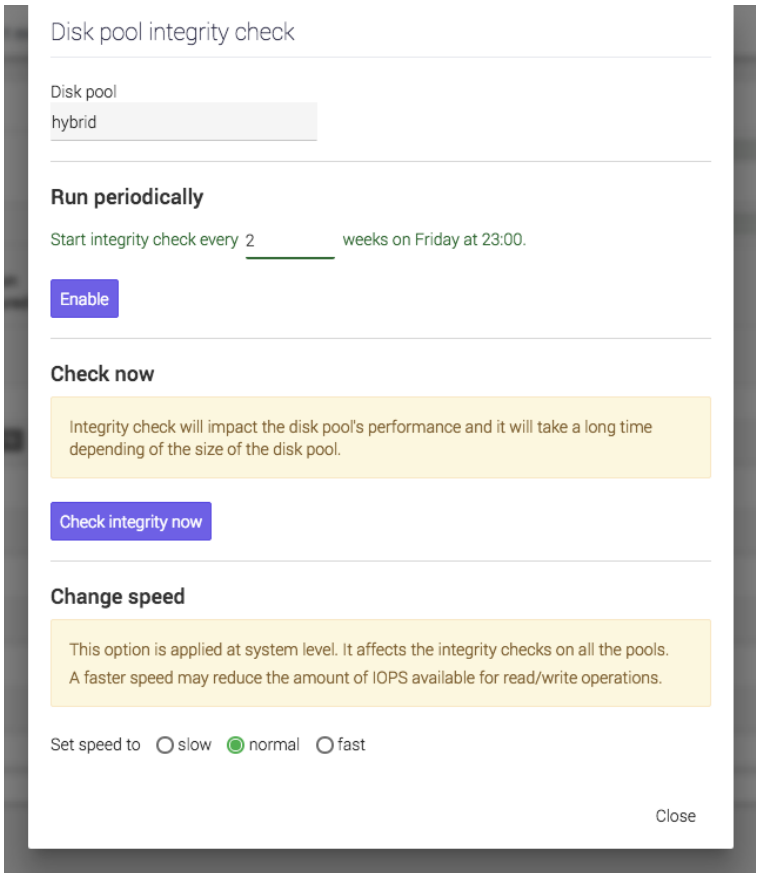

Write the new value you want in the field, between *every* and *weeks*. Click *Enable*.

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 disk is changed.

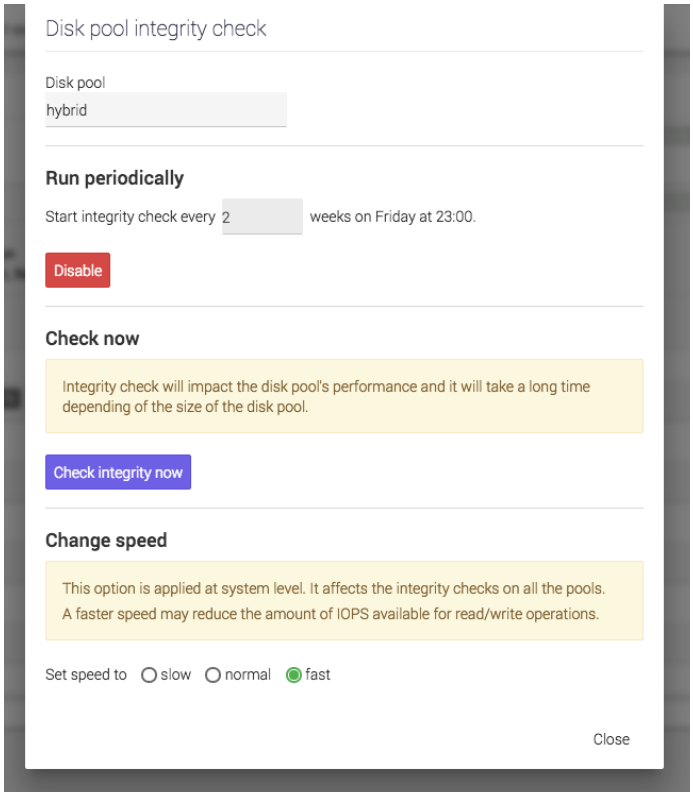

If you want to run an integrity check in this moment, click on *Check integrity now* button. This will perform with the selected speed.

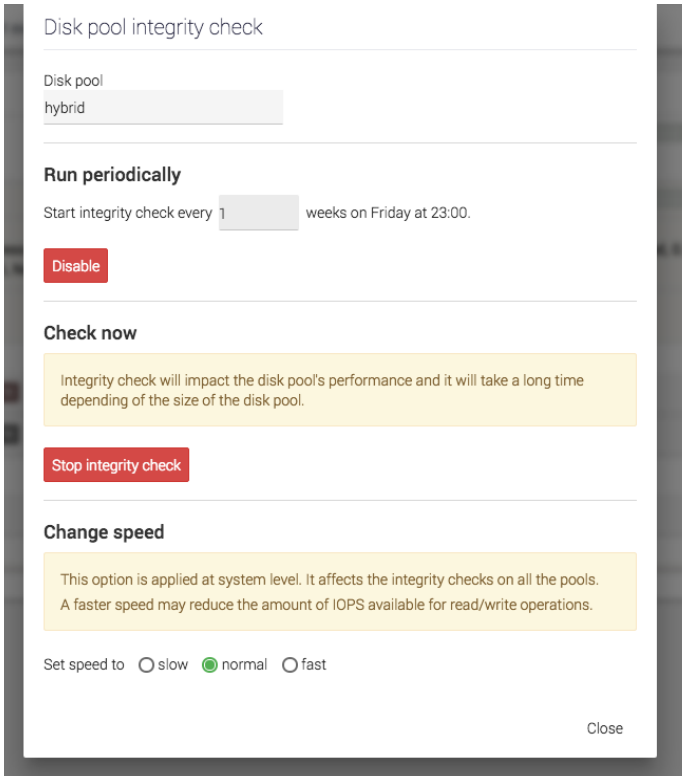

DATA PROTECTION AND REPLICATION

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. This can be done by replicating the snapshots to another machine. To replicate to a machine, you will have to add it in the *replication target* list.

Add a replication target

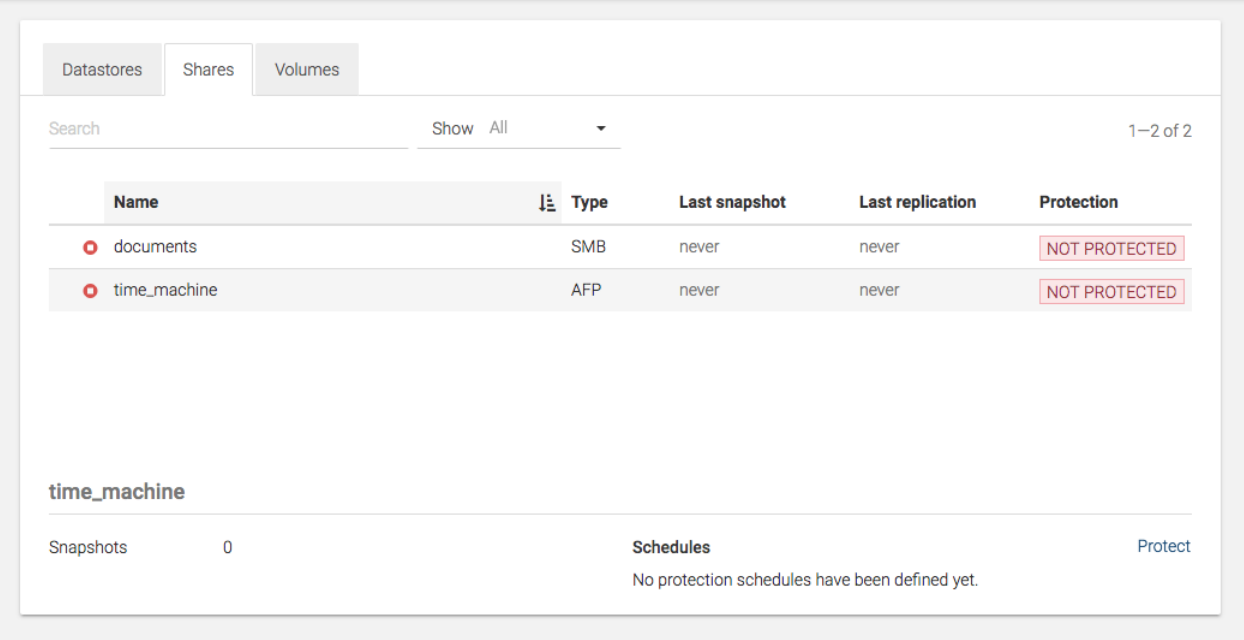

On the *Protect and replicate* page, click *Replication targets* in the action bar. A dialog will open listing the available replication targets.

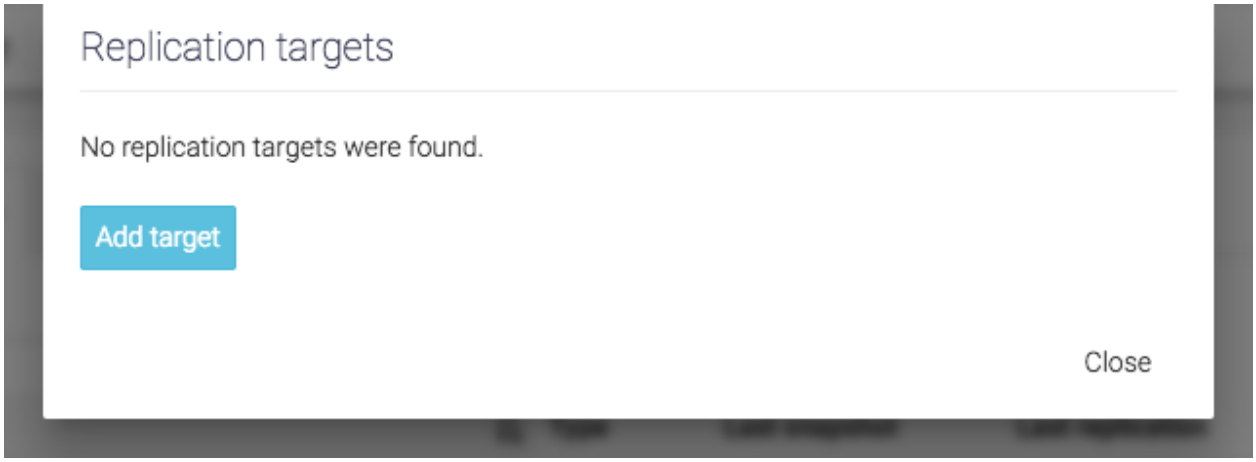

Click *Add target*. A dialog will open where you can enter the configuration for the *new replication target*. In the field *name* write the name you want to give to the target. In the field *Hostname / IP* write the hostname or IP of the replication target. In the *Password* field write the admin's password.

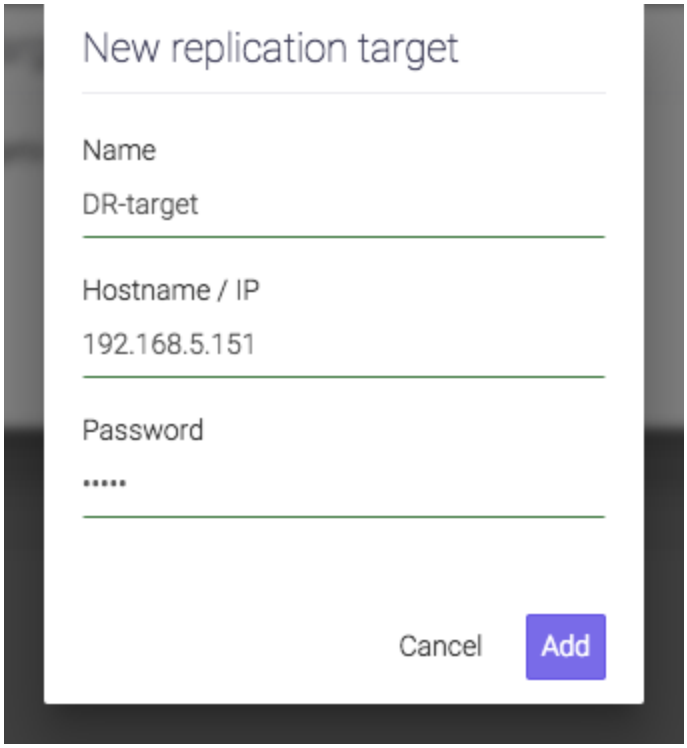

After filling the form and clicking *Add*, the list should be updated with the replication target.

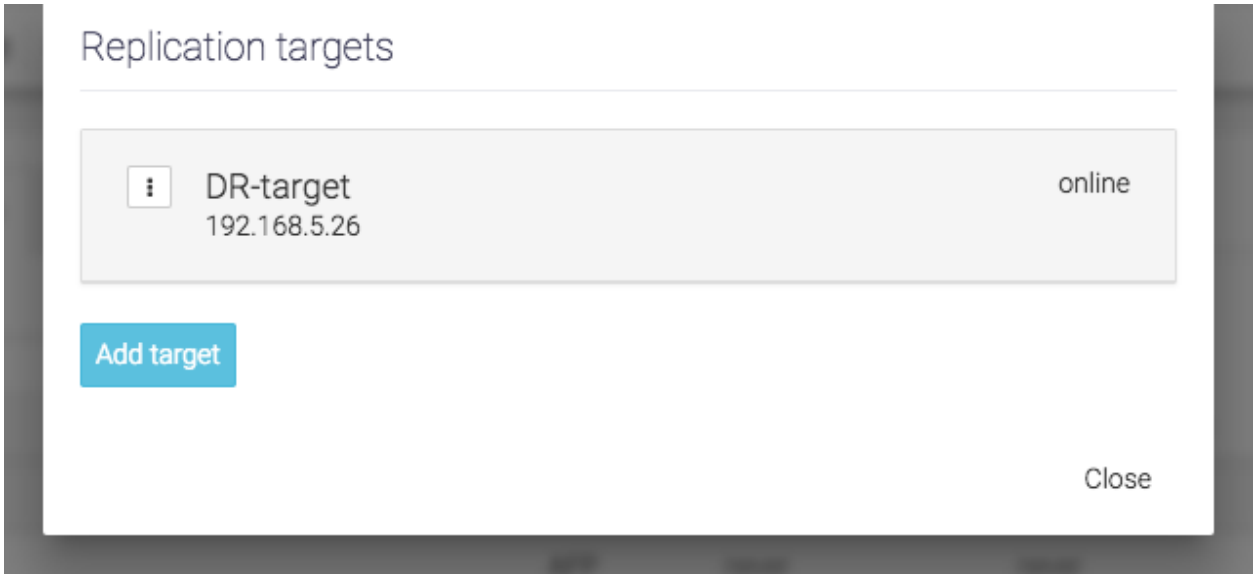

WARNING: This feature is not available on HYPERseries 2000. It will also fail when adding as replication target a HYPERseries 3000 machine with an OS below ver. 3.2.8.

Remove 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.

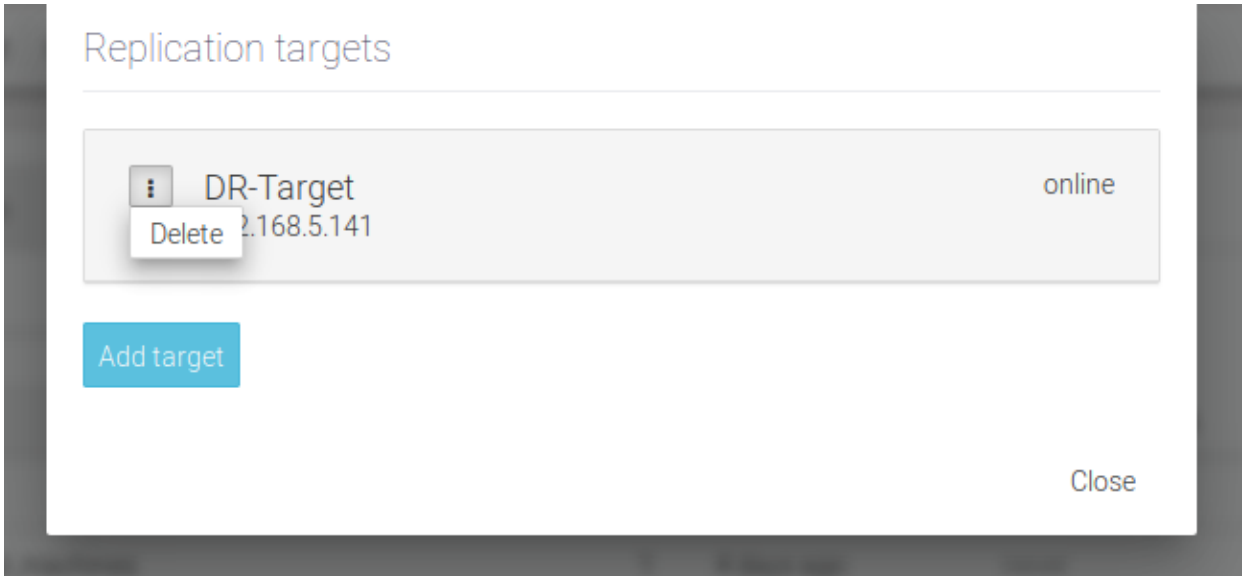

Protecting a datastore

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

The *Protect and replicate* page contains 3 tabs: *Datastores*, *Shares* and *Volumes*. Select *Datastores*.

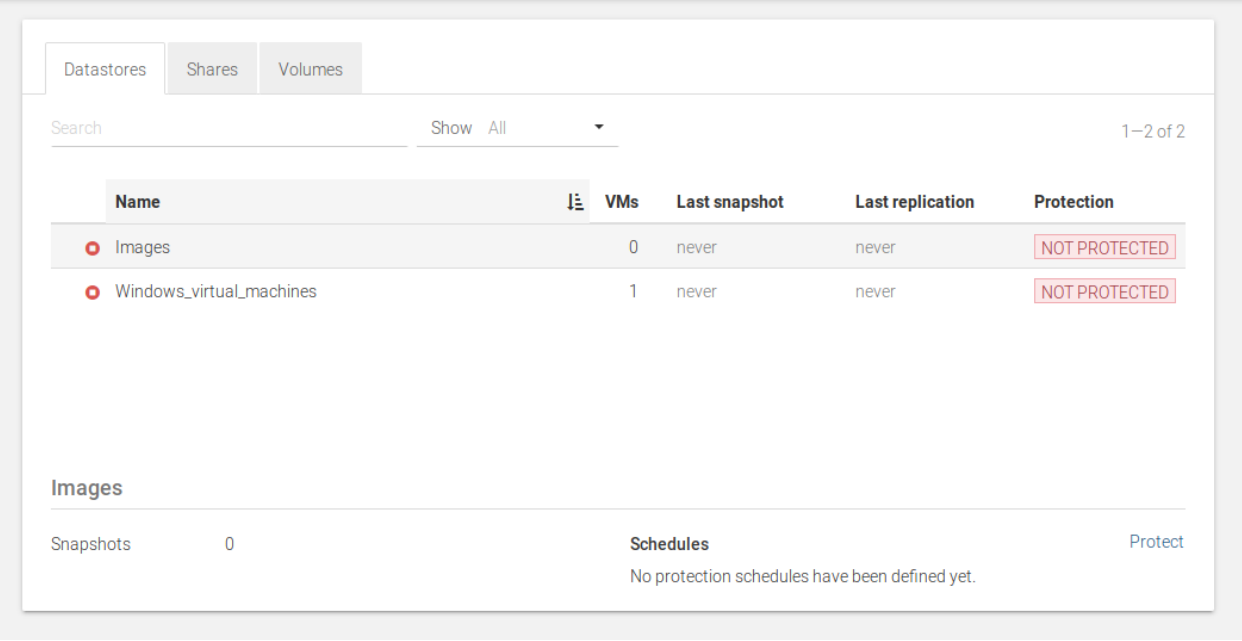

The table lists all the datastores. There is a *Search* bar, useful for searching for a specific datastore by specifying its name.

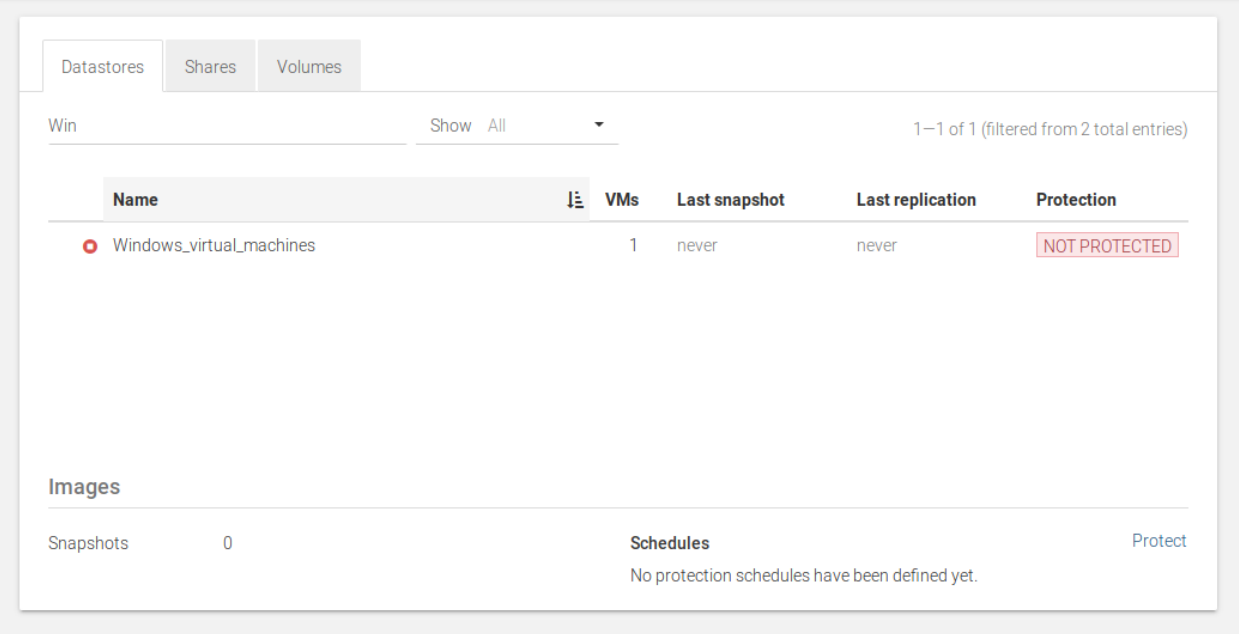

If you click on a datastore, more details about its protection will be shown.

Configuring the snapshot schedule for the datastore

In the details section, there is *Schedules*. Click on the *Protect* button. A dialog will appear where you can configure the snapshot policy for the selected datastore.

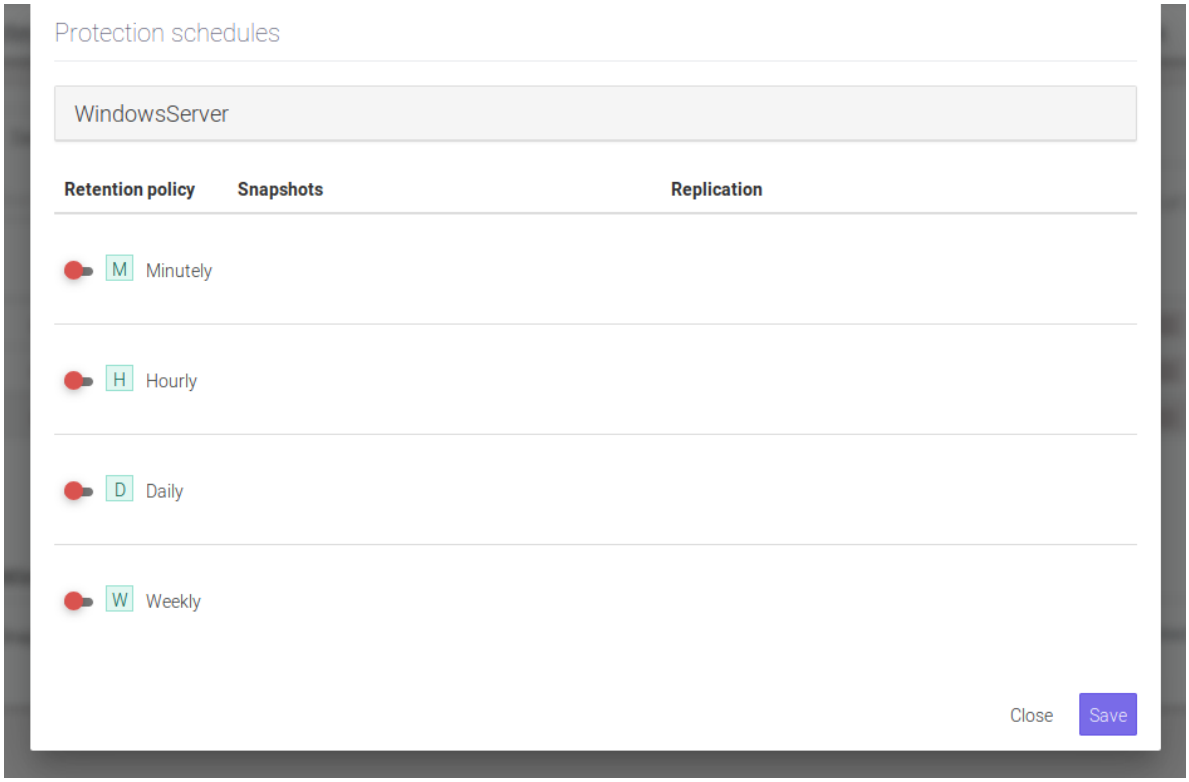

The first column shows the four frequencies a user can choose: *Minutely*, *Hourly*, *Daily* and *Weekly*. By default, they are disabled. To enable one schedule type, switch of the desired interval. After clicking, it will turn green.

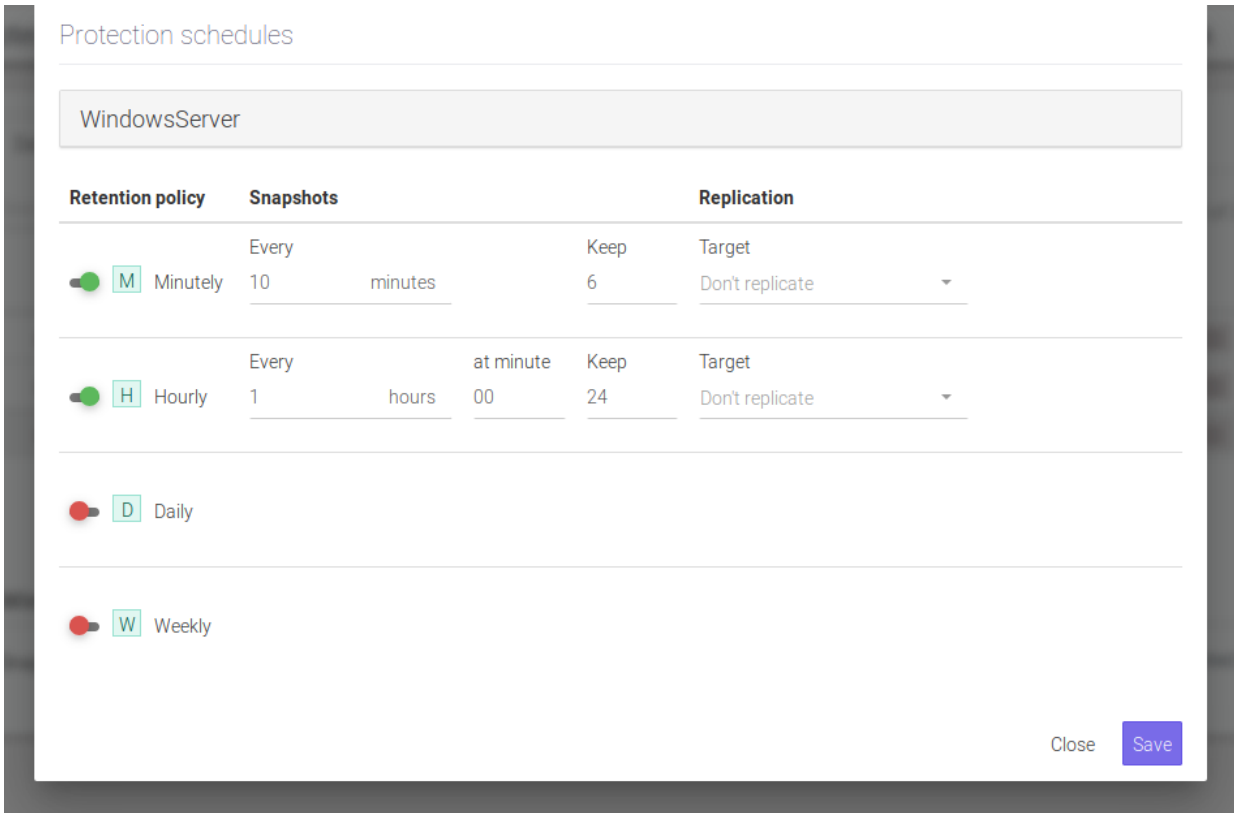

The fields from the *Snapshots* column will appear. Here you can specify how often you want the schedule to perform by changing the *Every* field.

To the right of the Every field, 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.

The third column, *Replication*, allows you to select to which replication target to replicate to. By default no replication is set. If you click on the *Target* field, a list of the defined replication targets will be shown. After you select the desired replication target, two more options will appear.

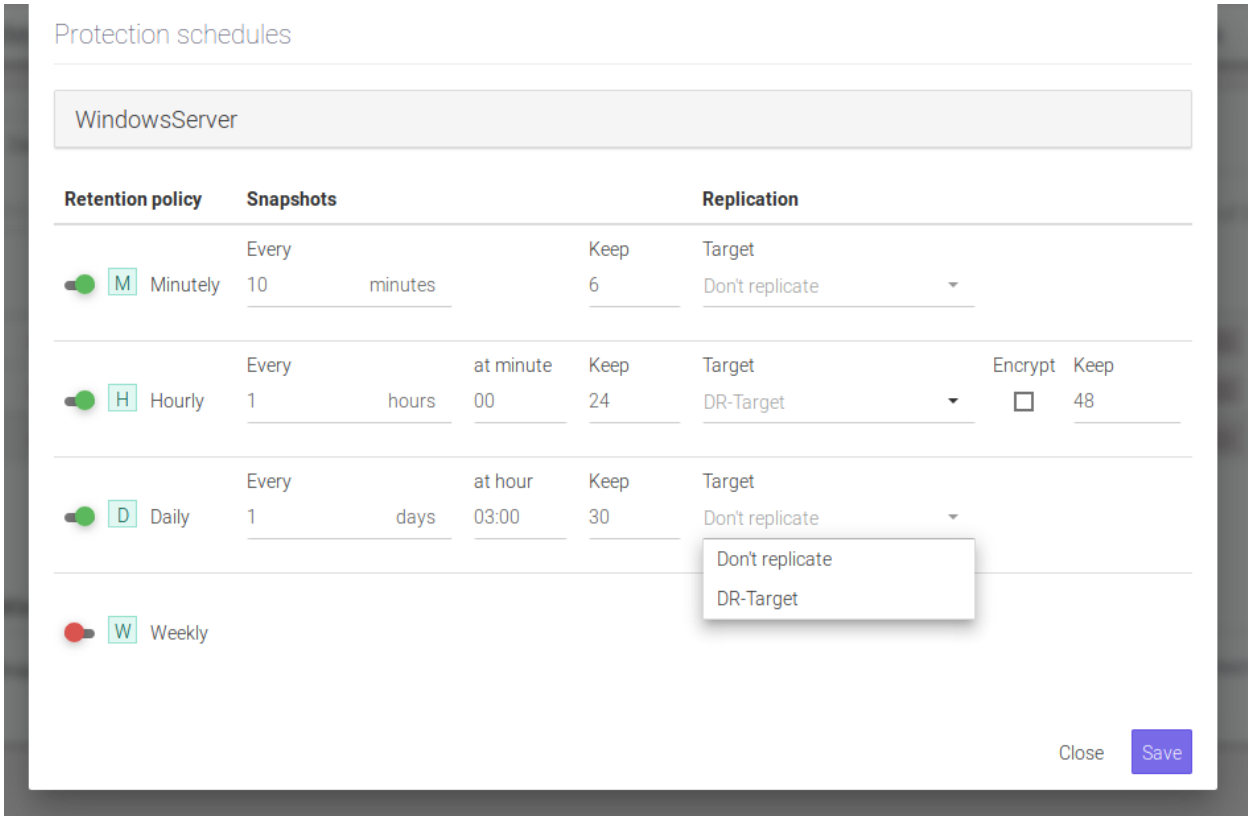

If you want the replication to be encrypted, click on the box below the *Encrypt* field.

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 finish all the configuration, click *Save*. Your schedules should appear in the details section, in *Schedules*. You have two buttons for each schedule type, *Pause* and *Run*.

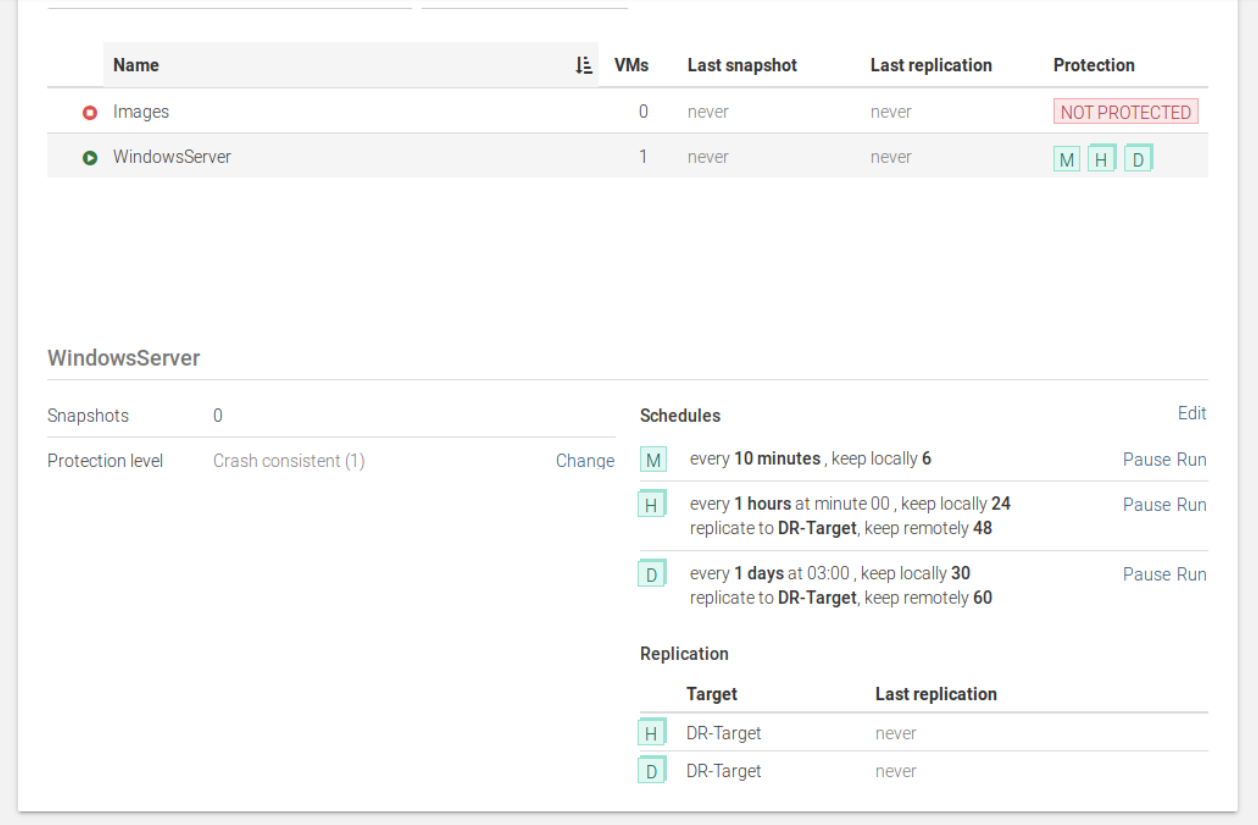

A schedule, that is not currently replicating, can be paused. This will also pause the replication. If you want to resume it, just click on the *Resume* button.

If you want to run a schedule now, click on *Run*. This will start the selected schedule. It will also replicate if the schedule was configured so.

NOTE: First replication will always copy all the data found in the dataset, share, or volume. Subsequent replications will transfer only the differences between the last successful replication and current snapshot.

Canceling a running replication

A schedule that is currently replicating can be canceled.As a consequence, the local snapshot will not be taken and the replica will not be created on the destination.

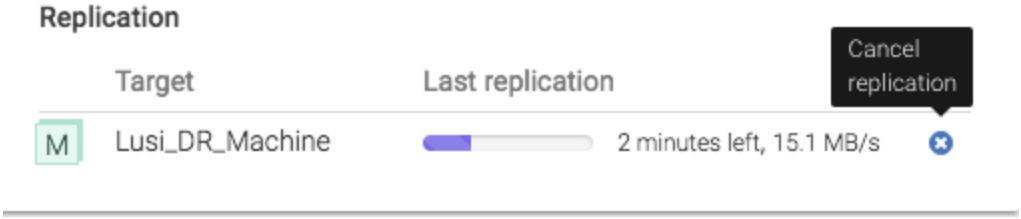

After you click on the cancel replication icon, you will receive a confirmation saying that the Snapshot replication was cancelled.

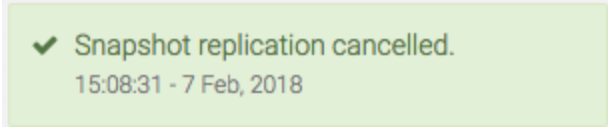

Configuring the type of snapshots for the virtual machines

If a datastore has virtual machines on it, on the left of the details view there is a section *Protection Level*.

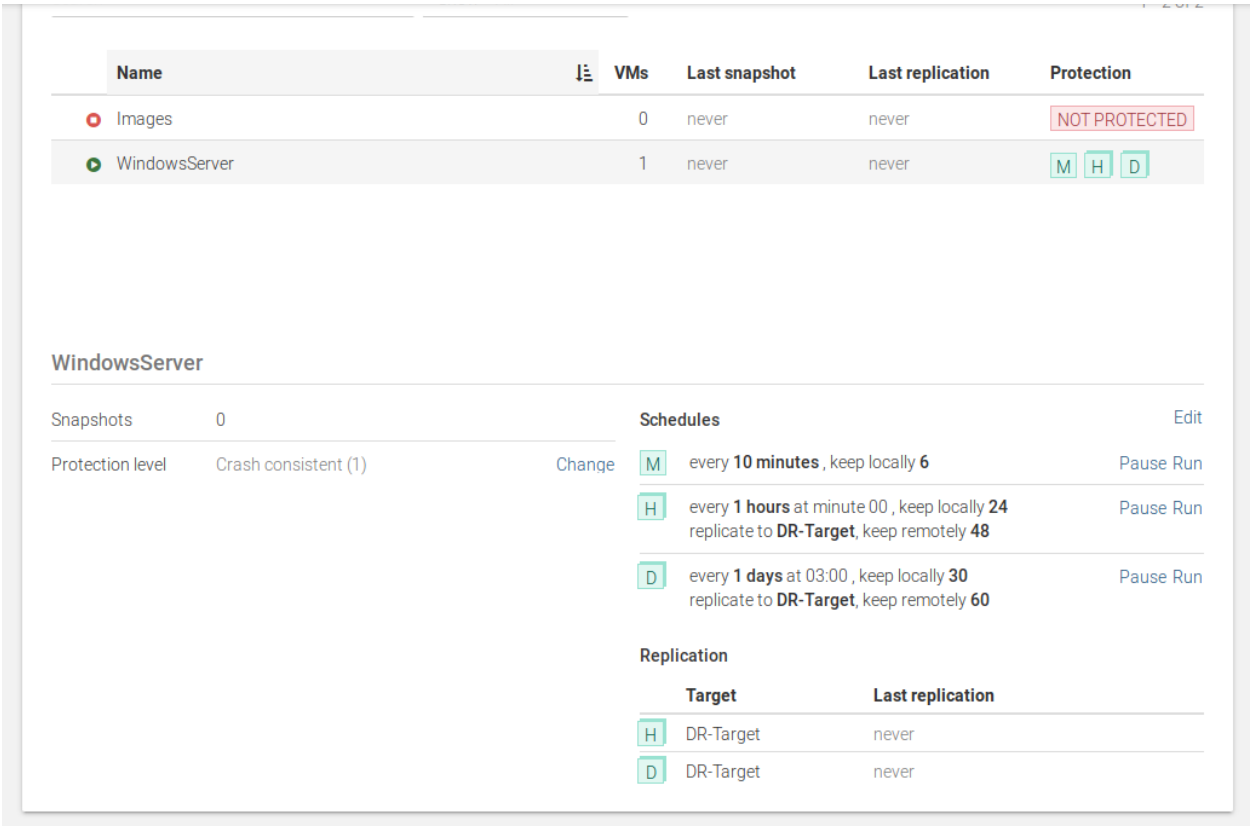

There are three types of protection levels for a virtual machine:

- *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* - 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 disk. The snapshot will be taken after the flush operation finishes. Disk 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)

- *Live snapshot* - Before taking a snapshot all operations on the virtual machine will be suspended and saved to the disk. This includes waiting for disk 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 *Change* button. A dialog will open where you can select the protection level you want for the virtual machine.

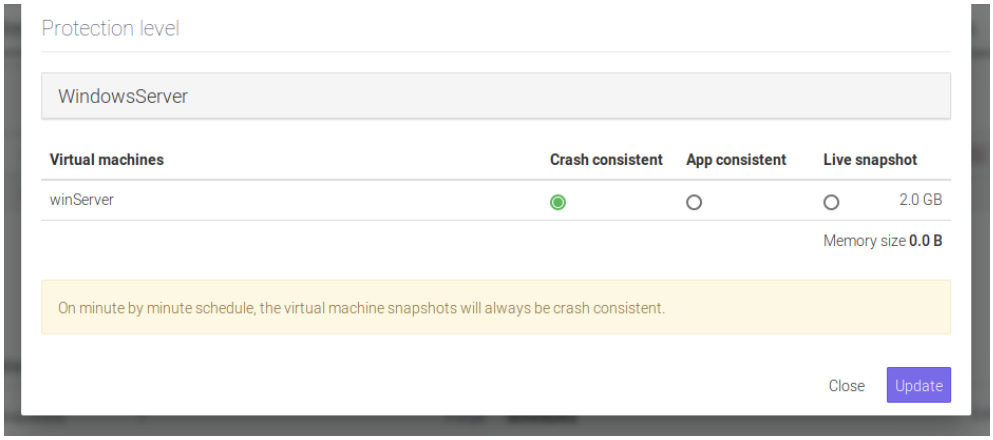

It will contain a list with all the virtual machines contained by the selected datastore. Each virtual machine can have one of the three available protection levels: *Crash consistent*, *App consistent* and *Live snapshot*. Select your desired options and then click on *Update*.

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

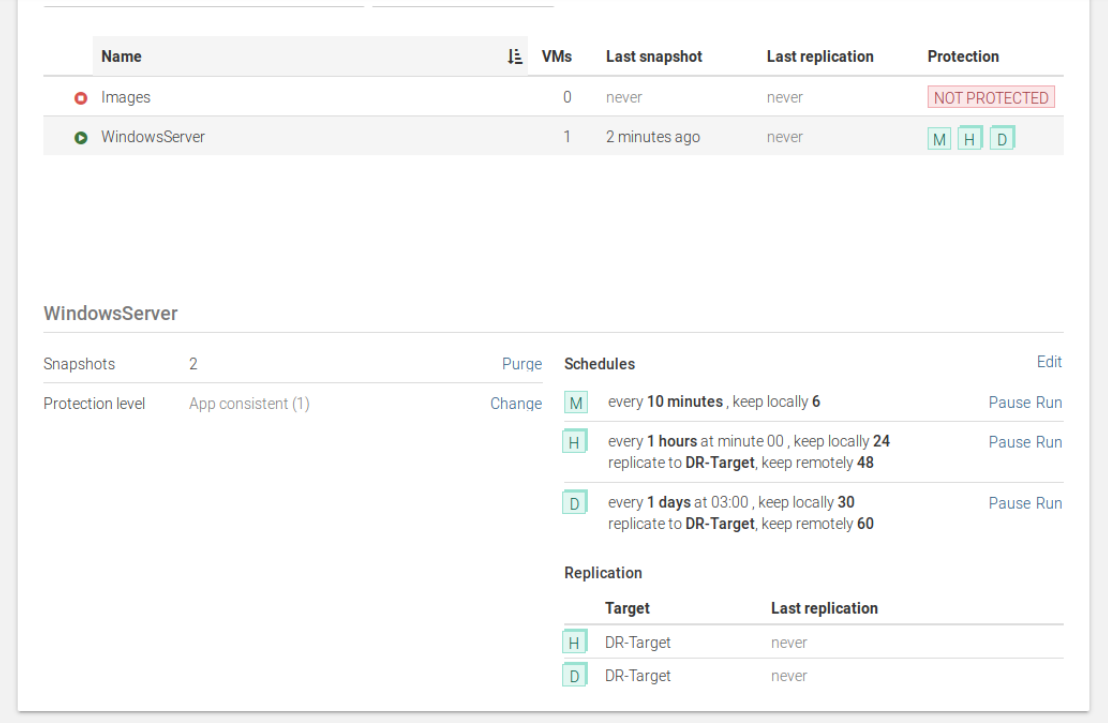

The details section will be updated with the new information. The following snapshots will take the virtual machine snapshots with the new configuration.

NOTE: For the minutely schedule, the protection level will always be crash consistent.

Protecting a share

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

The *Protect and replicate* page contains three tabs: *Datastores*, *Shares* and *Volumes*. Select *Shares*.

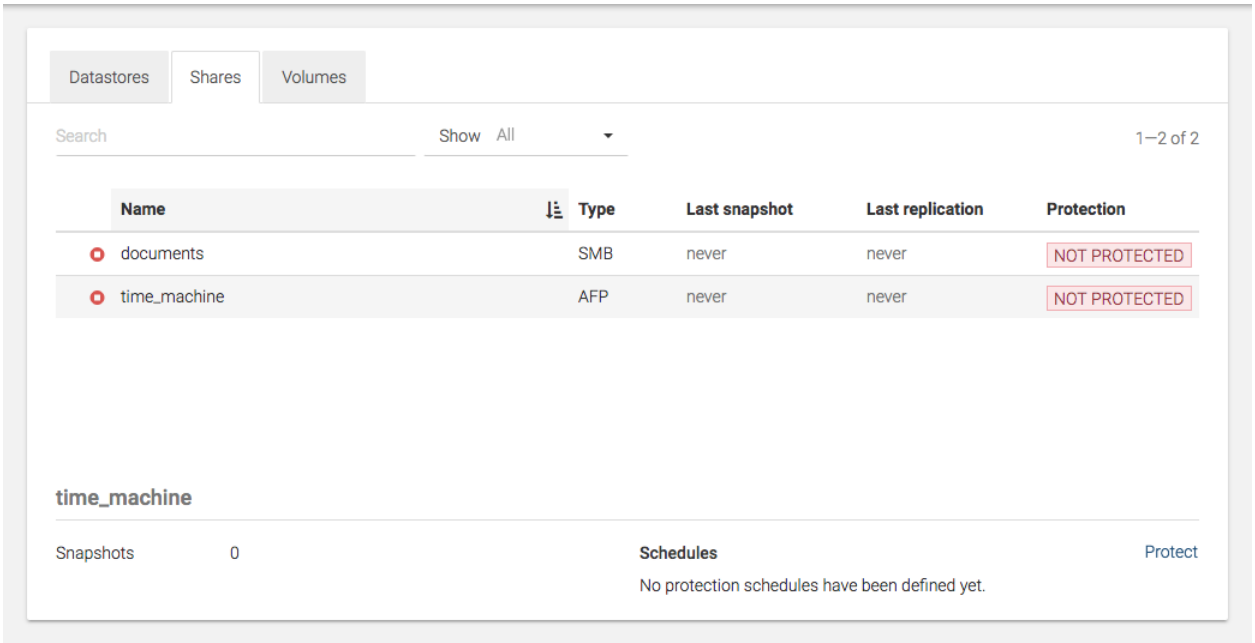

The table lists all the shares. There is a *Search* bar, useful for searching for a specific share by specifying its name.

Select the share you want to protect by clicking on it. In the details section there is *Schedules*. Click on the *Protect* button.

A dialog will appear where you can configure the snapshot policy for the selected share.

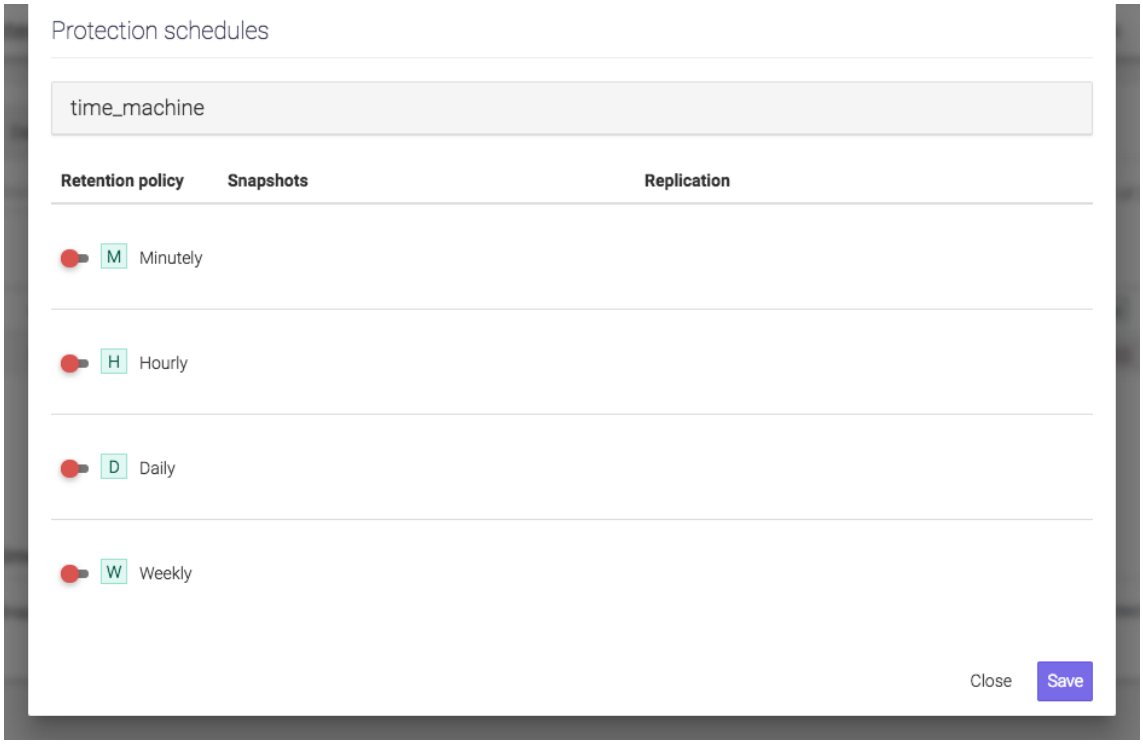

The first column shows the four frequencies a user can choose: *Minutely*, *Hourly*, *Daily* and *Weekly*. By default, they are disabled. To enable one schedule type, click on the switch on the left of the desired interval.

After clicking, it will turn green.

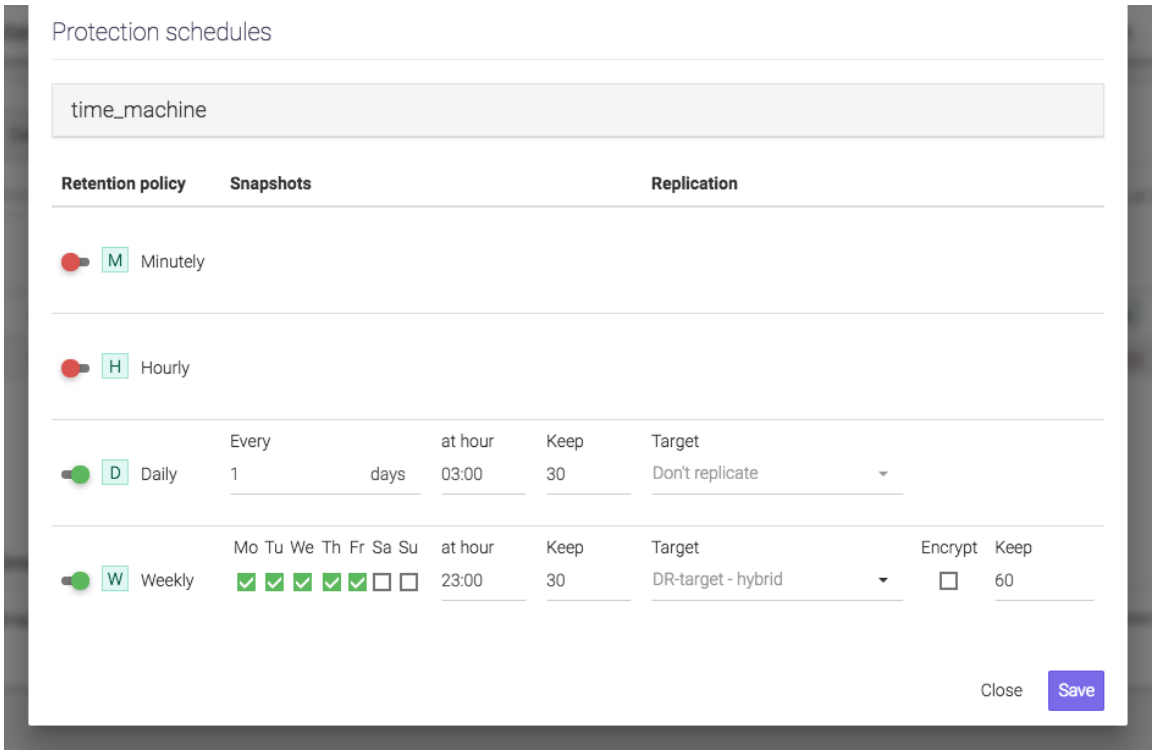

The fields from the Snapshots column will appear. Here you can specify how often you want the schedule to perform under the Every field.

NOTE: We recommend spacing out snapshot taking time at least 5 minutes. If all schedules run at the same time, and then replicate, it can lead to high network traffic and high disk IOPS on the replication target.

The second field of the Snapshots column is the Keep field, to the right of the Every field. Here you can specify how many snapshots should be kept on the machine. SynetoOS will keep the most recent ones and will delete the older ones, if there are more snapshots than you specified.

NOTE: The retention policy applies to a type of interval only and for the selected share.

The third column, *Replication*, 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 replication targets will be shown. After you select the desired replication target, two more options will appear.

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.

After you finish all the configuration, click on *Save*. Your schedules should appear in the details section, in *Schedules*. You have two buttons for each schedule type, *Pause* and *Run*.

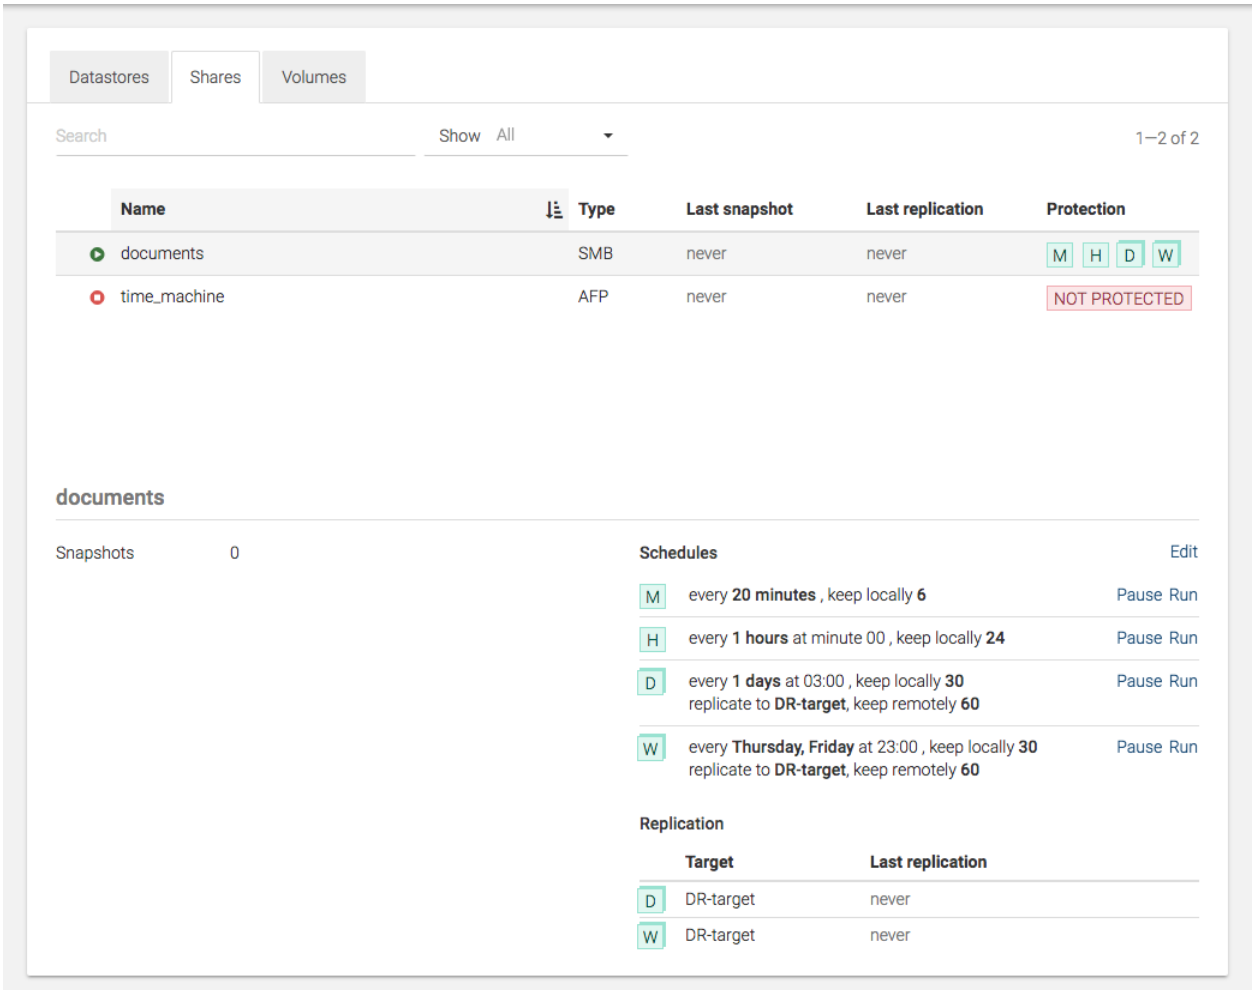

A schedule, that is not replicating right now, can be paused. This will also pause the replication. If you want to resume it, just click on the *Resume* button.

If you want to run a schedule now, click on *Run*. This will start the selected schedule. It will also replicate if it has replication set.

Protecting a volume

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

The *Protect and replicate* page contains 3 sections: *Datastores*, *Shares* and *Volumes*. Select *Volumes*.

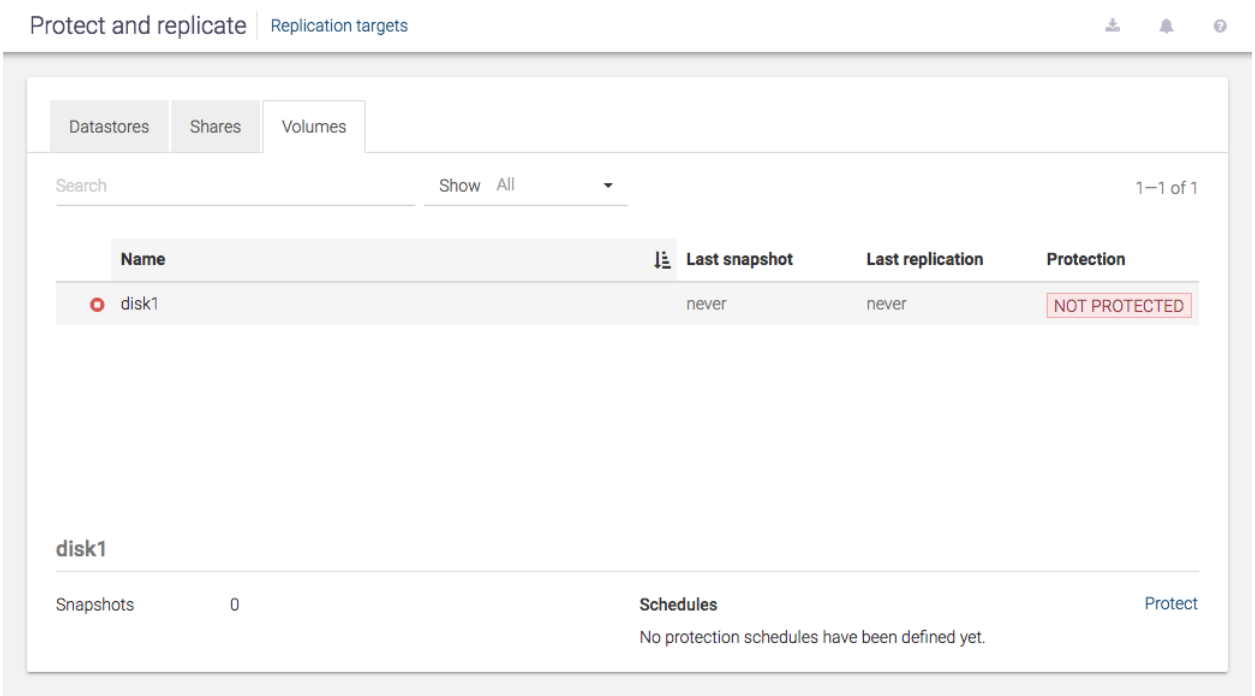

The table lists all the volumes. There is a *Search* bar, useful for searching for a specific volume by specifying its name.

Select the volume you want to protect by clicking on it.

On the details section, there is *Schedules*. Click on the *Protect* button.

A dialog will appear.

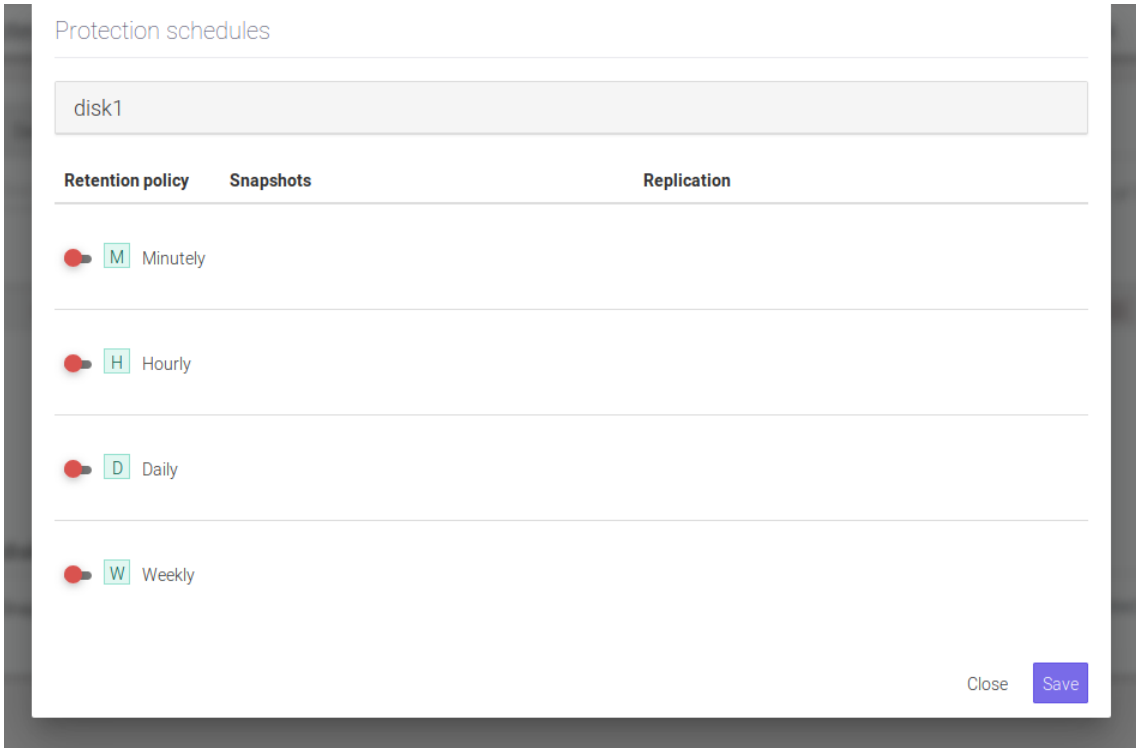

The first column shows the four frequencies a user can choose: *Minutely*, *Hourly*, *Daily* and *Weekly*. By default, they are disabled. To enable one schedule type, click on the switch on the left of the desired interval.

After clicking, it will turn green.

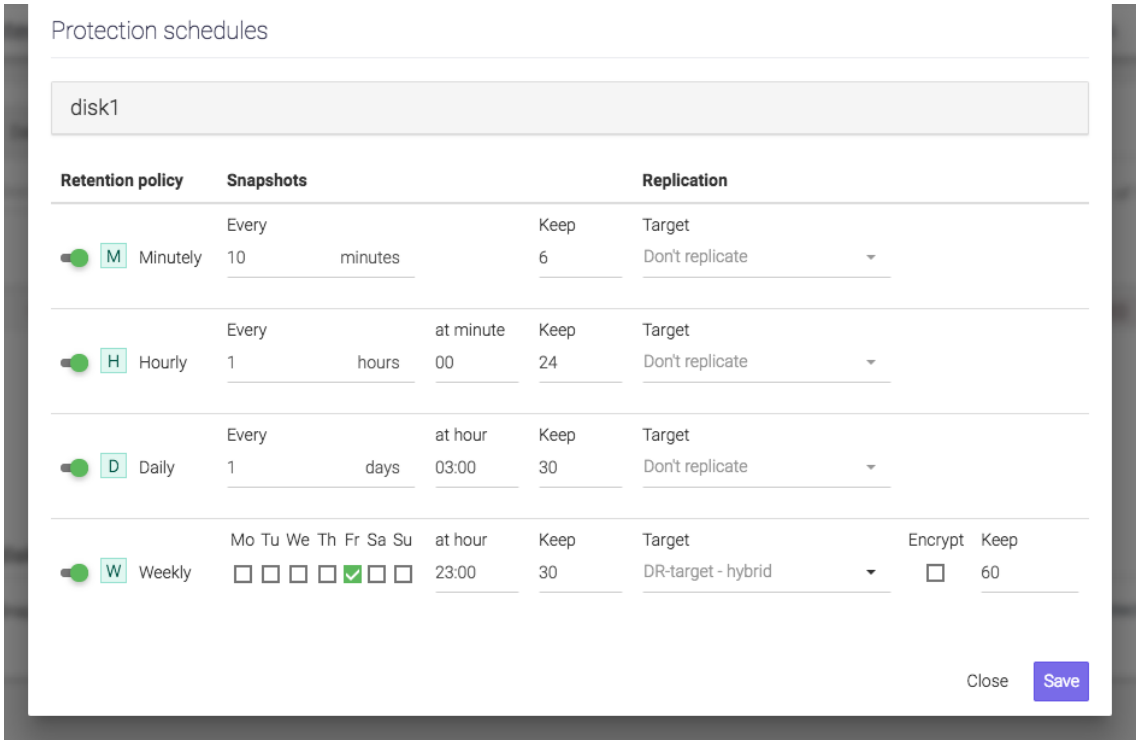

The fields from the Snapshots column will appear. Here you can specify how often you want the schedule to perform under the Every field.

The second field of the Snapshots column is the Keep field, to the right of the Every field. Here you can specify how many snapshots should be kept on the machine. SynetoOS will keep the most recent ones and will delete the older ones, if there are more snapshots than you specified.

NOTE: The retention policy applies to a type of interval only and for the selected volume.

The third column, *Replication*, 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 replication targets will be shown. After you select the desired replication target, two more options will appear.

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.

After you finish all the configuration, click on save. Your schedules should appear in the details section, in *Schedules*. You have two buttons for each schedule type, *Pause* and *Run*.

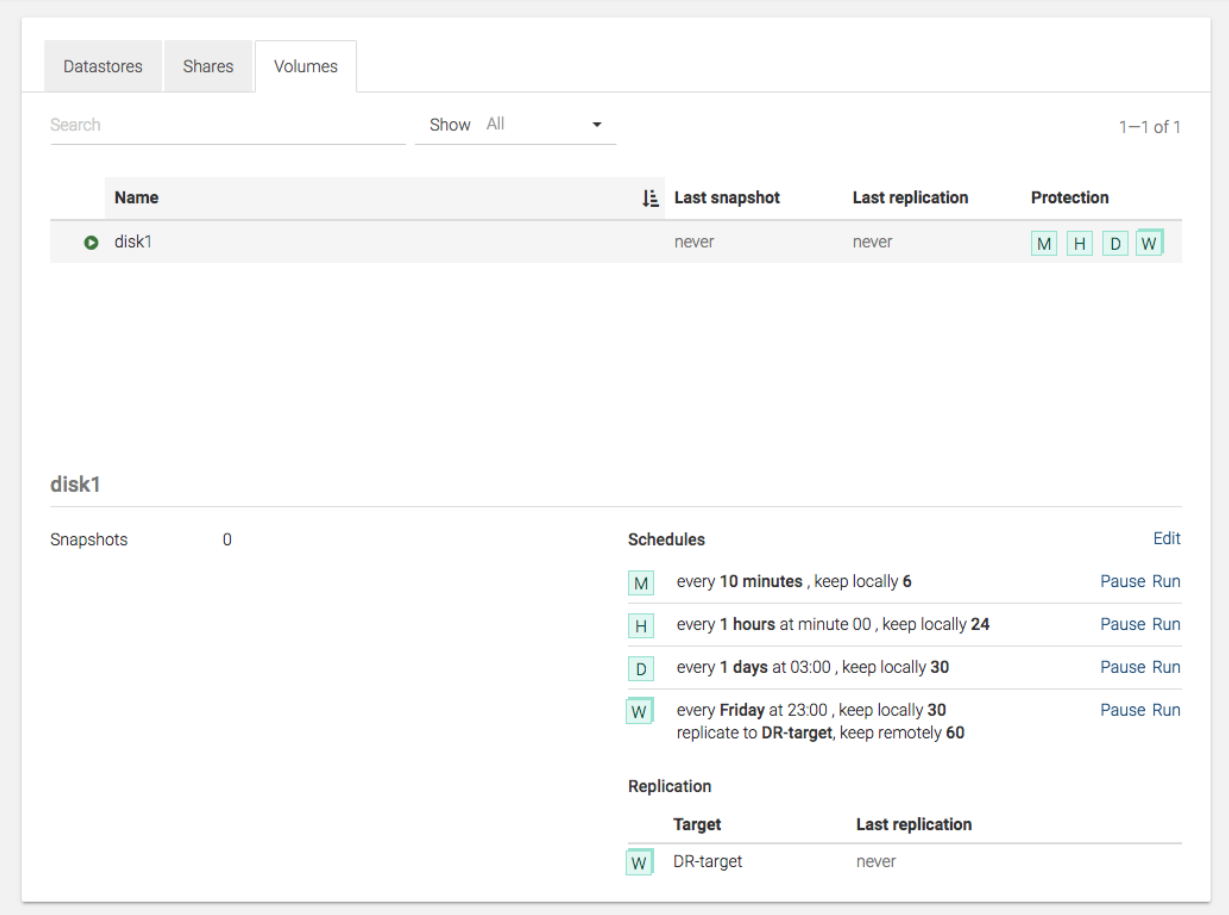

A schedule can be paused. This will also pause the replication. If you want to resume it, just click on the *Resume* button.

If you want to run a schedule now, click on *Run*. This will start the selected schedule. It will also replicate if it has replication set.

Monitoring

To get an overview of the protection status of the datastores, shares or volumes, select the tab you desire on the *Protect and Replicate* page.

The table can be filtered using the *Search* and *Show* bars. *Search* bar lets you search for a specific name or parts of a name.

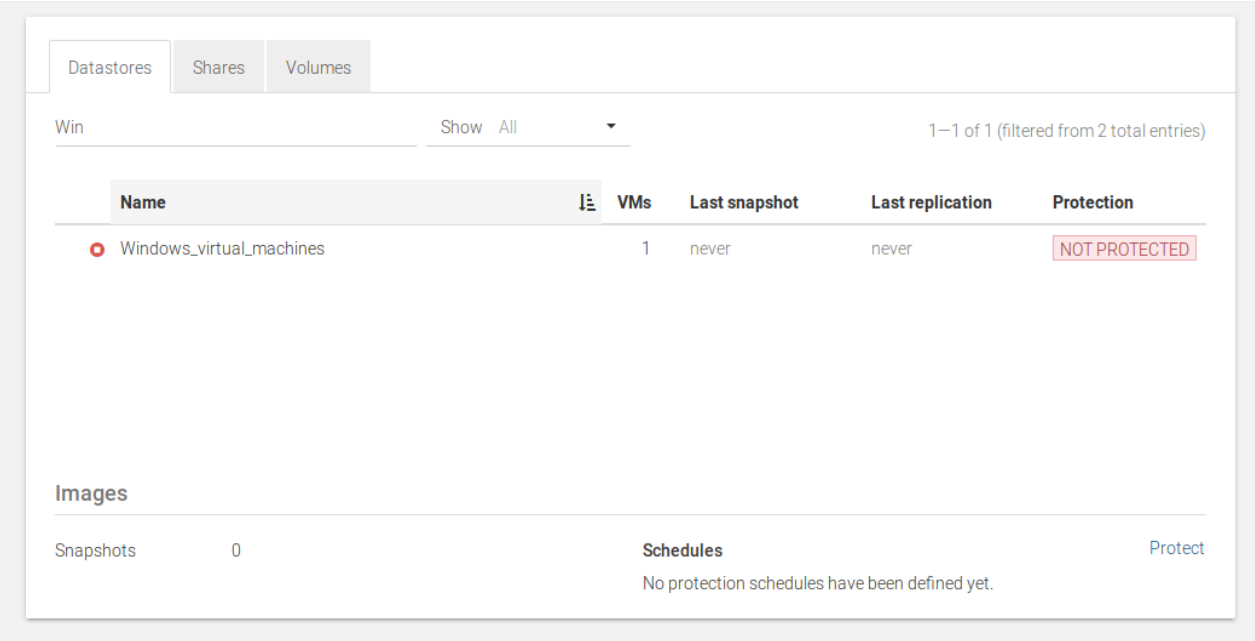

Show bar has three options: *Protected*, *Unprotected* and *All*. Protected option will show you only the ones which have at least one schedule set. Unprotected option will show only the ones that have no protection schedule set. The *All* option shows both the protected and unprotected entities.

The left icon in a row represents the state of protection. It can be *ON*, *OFF* or *paused*. *ON* means there is at least one schedule set for it which is not paused. *OFF* means there is no schedule set for it. *Paused* means it has schedules set on it, but they are all paused.

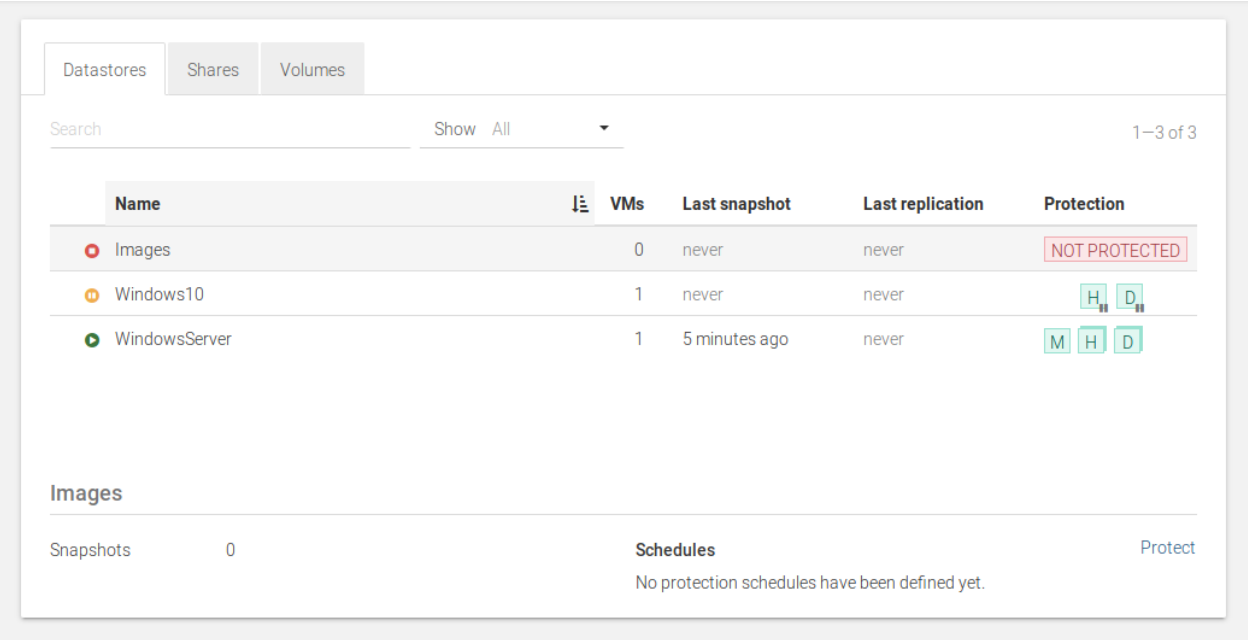

In the *Last Snapshot* column we can see how much time has passed since the last successful local snapshot. If it is colored red, it means the last scheduled snapshot failed.

In the *Last replication* column we can see how much time has passed since the last successful replication. If it is colored red, it means the last scheduled replication failed.

In the *Protection* tab, we can see what schedules have been defined for that datastore.

For more details of a specific entity protection status, click on the entity in the table.

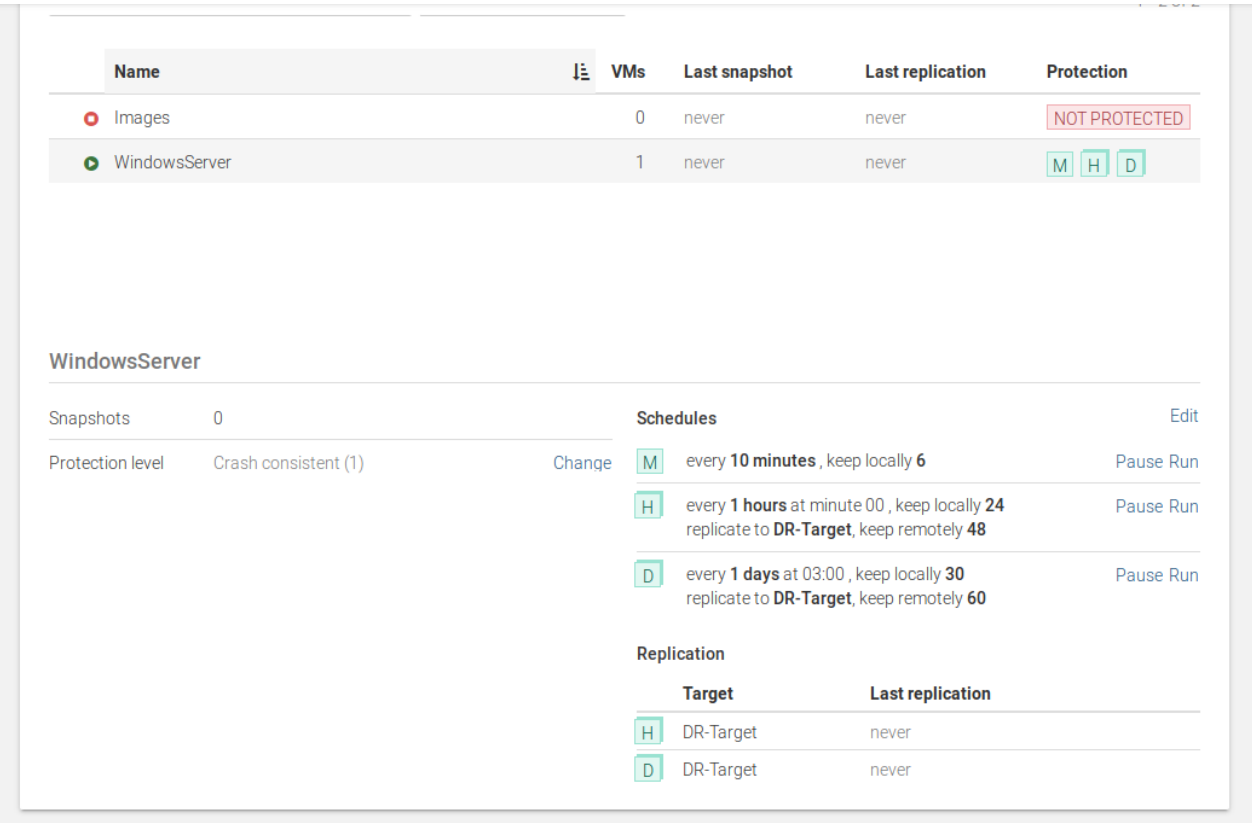

In the details section, there is *Schedules* on the right. There you can find info about all the snapshot schedules set on that entity. Below it, there is the *Replication* section. For each schedule with replication, there will be a row showing the *Target* to which it replicates to and when was the *Last replication*. If the last replication failed, it will be highlighted with red.

If the replication is currently in progress, a progress bar with relevant information will be shown.

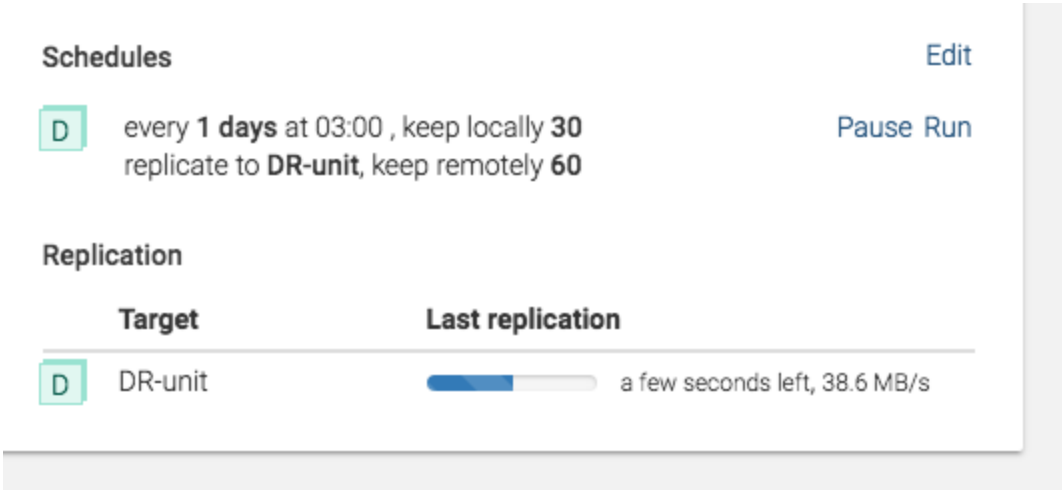

SYSTEM CONFIGURATION

Networking

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

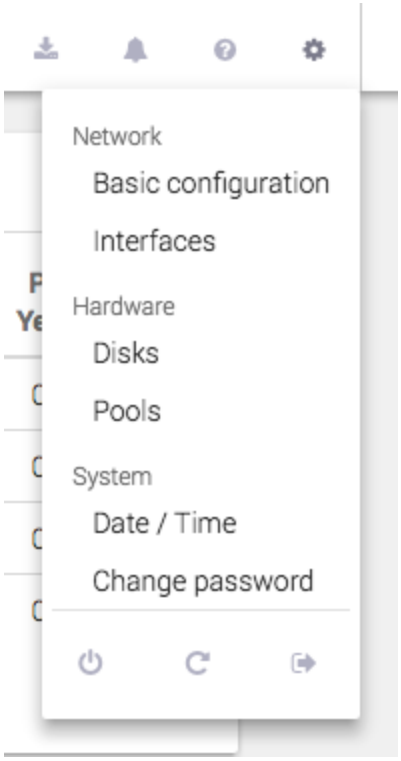

The network page allows you to configure the following:

- *Hostname*

WARNING: This will require a reboot.

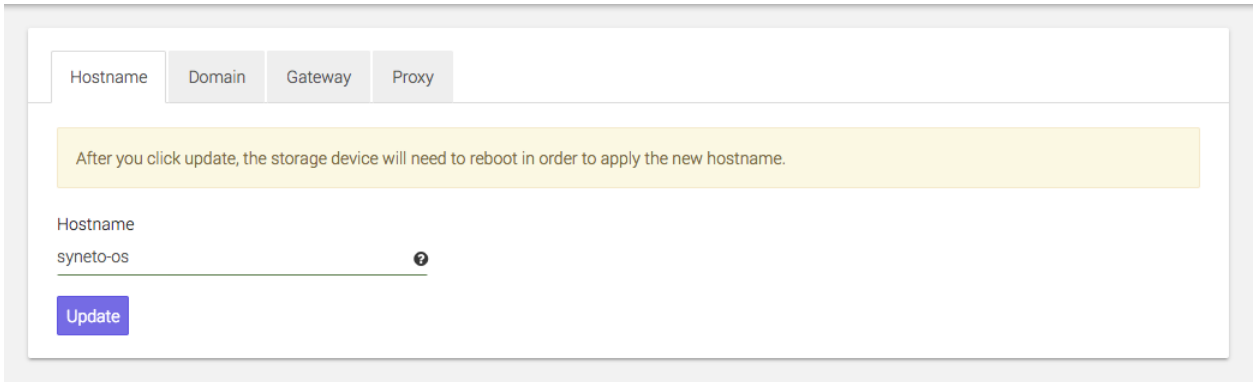

- *Domain* - Up to three domain name servers to be used when resolving internet names and a domain name to append to hostnames during DNS name resolution

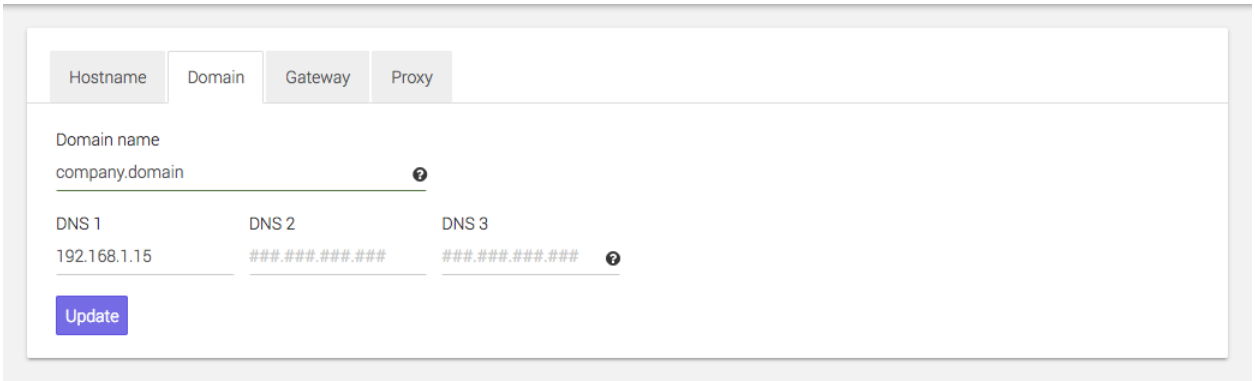

- *Gateway* - a router used by the storage to connect to the internet

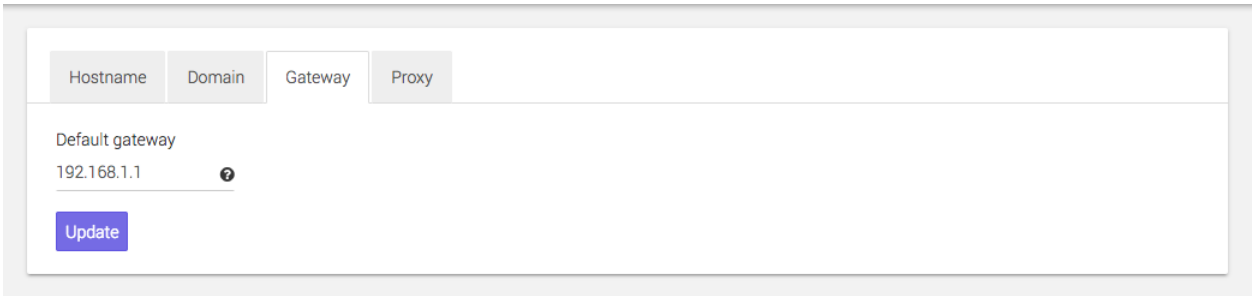

- *Proxy* - the host and port of the HTTP proxy server

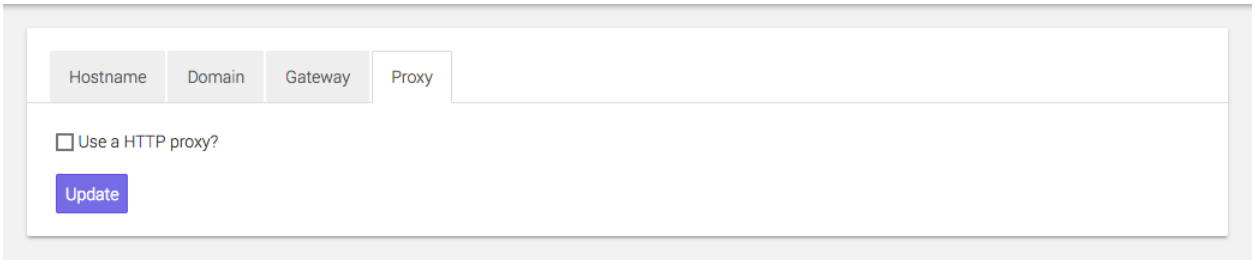

Ethernet interfaces

Select *Secondary menu* → *Interfaces*, in the *Network* section.

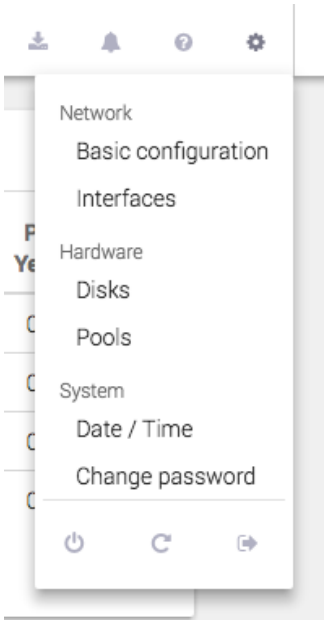

Syneto HYPER comes preconfigured with two network interfaces. The first one is the management and the second one is used for interaction with the ESX host. These are not physical interfaces, they are virtualized interfaces provided by the local ESX.

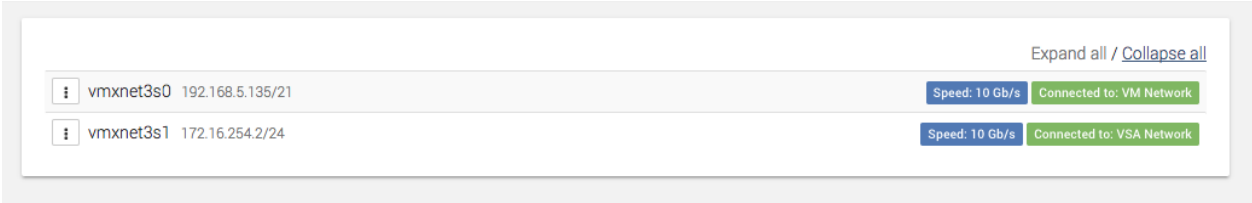

The *Network interfaces* page allows you to configure the available ethernet interfaces.

Click the *Edit* option in the context menu.

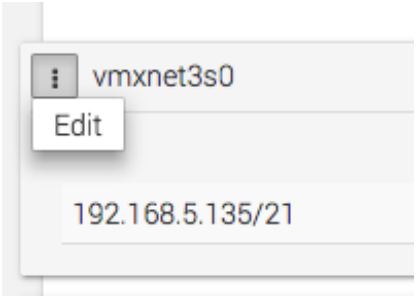

A dialog will open where you can configure the network interface. You can have a *dynamic* or *static* IP. Also, you can configure a limit for the data being transmitted through that interface in the *MTU* field.

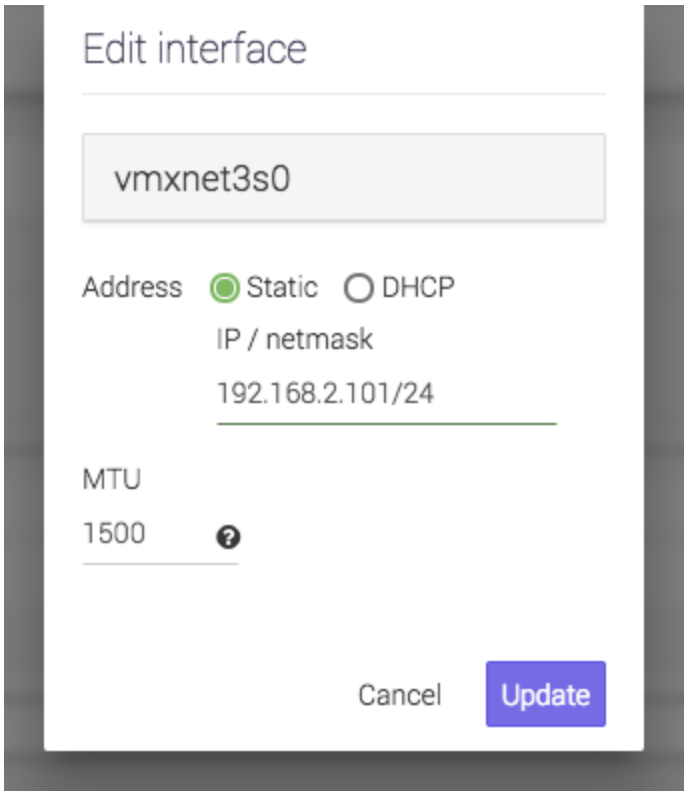

NOTE: SynetoOS 4 does not offer support for aggregates or VLANs.

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.

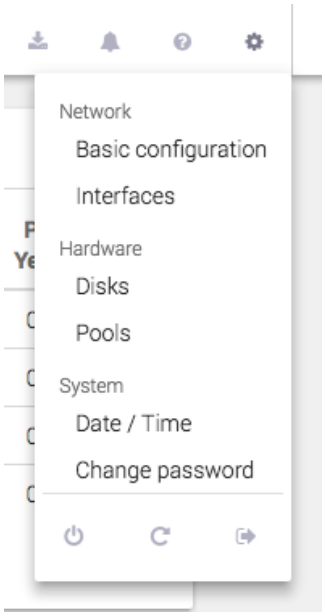

The time page allows you to configure the following:

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

  Recommended: Enable *automatic time synchronization*.

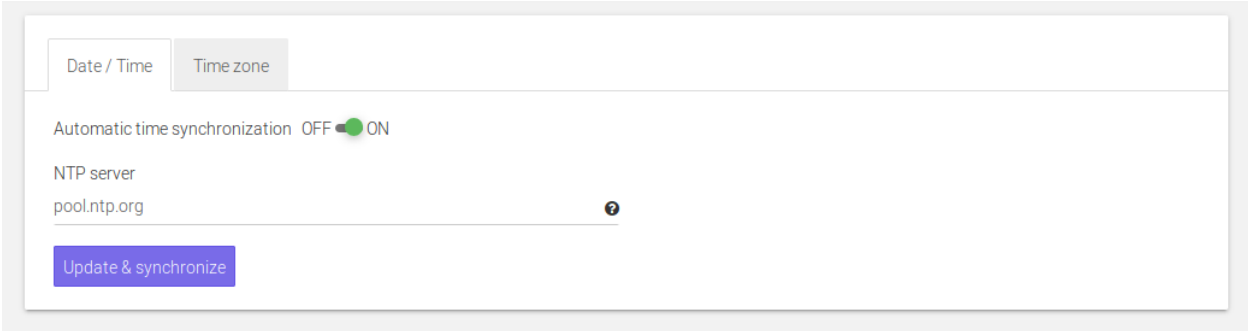

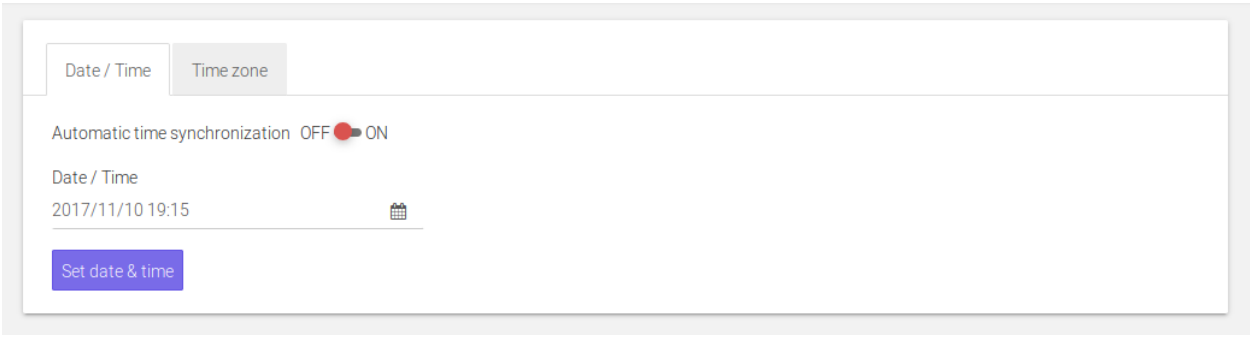

- *Timezone*

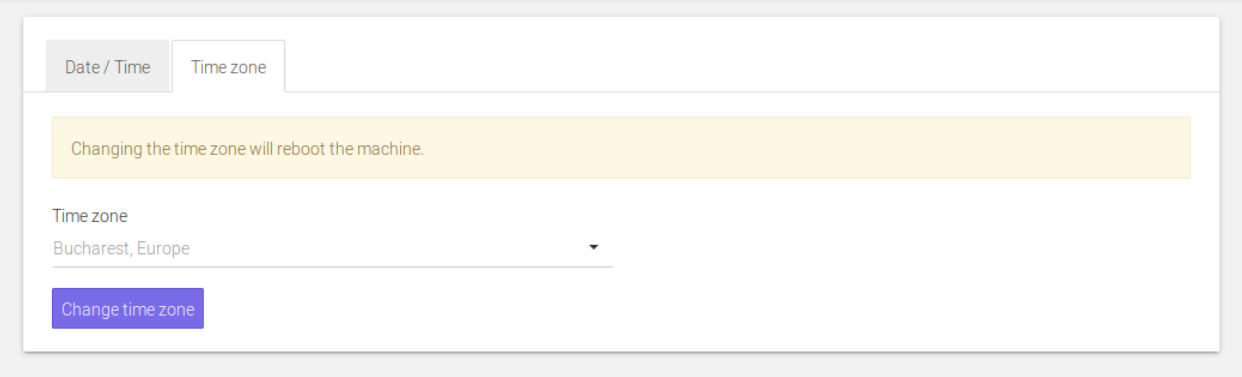

WARNING: This will require a reboot.

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 *Seconday (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.

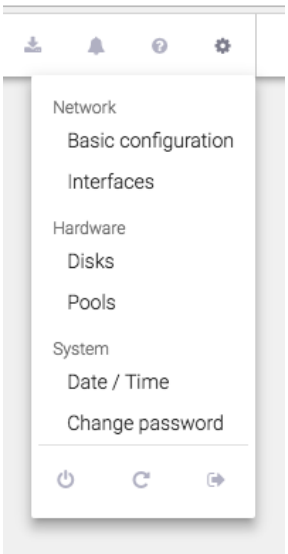

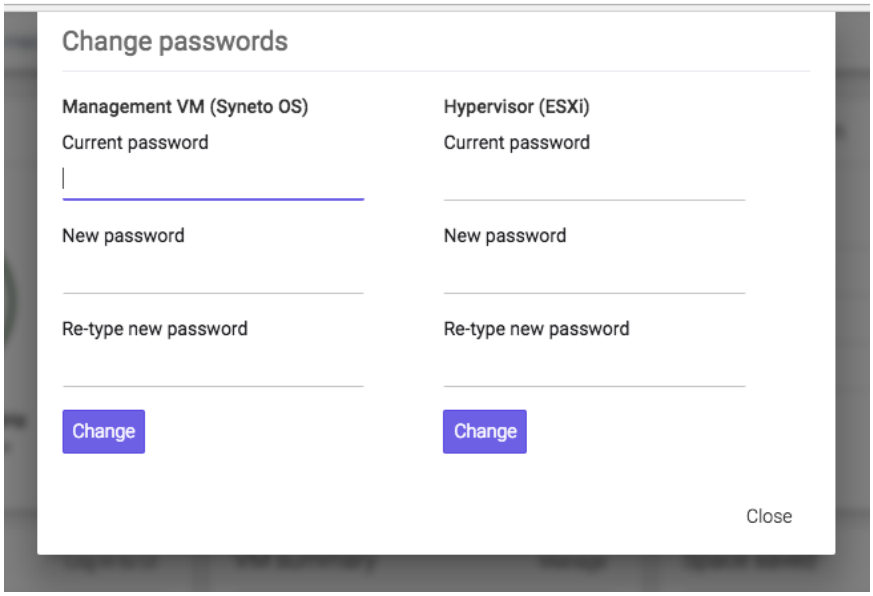

Changing External ESXi Host Password

If an external ESXi password is changed from the vSphere management interface the password should be updated in SynetoOS too. In SynetoOS go to *Datastores → Hosts* and select *Edit* from contextual menu.

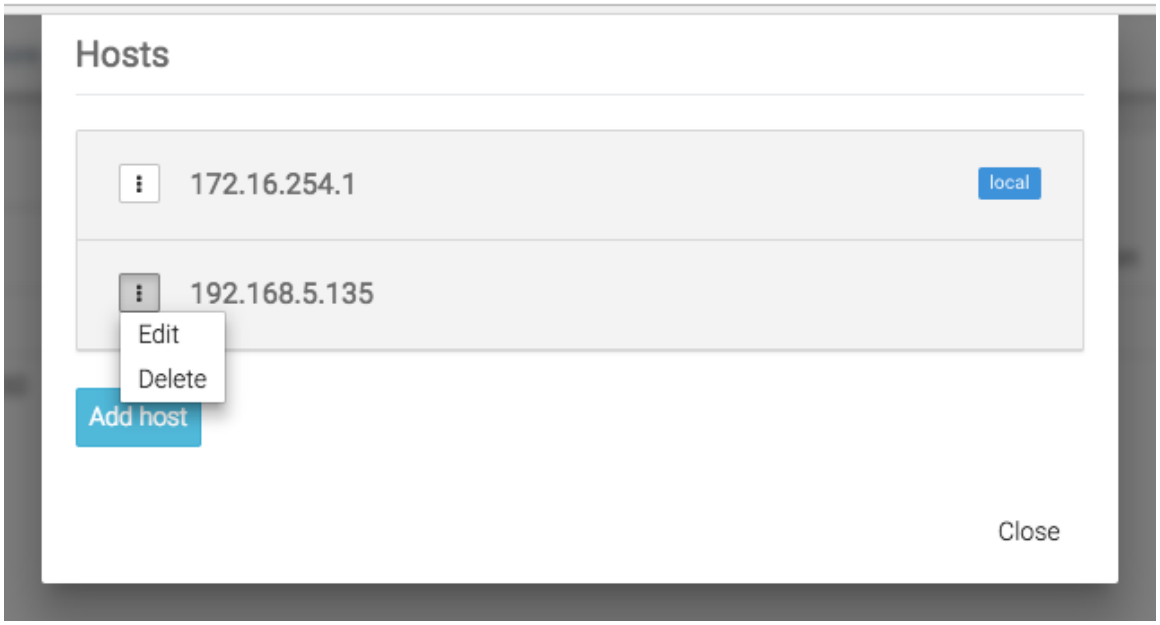

Then a dialog will open, where the new password can be introduced.

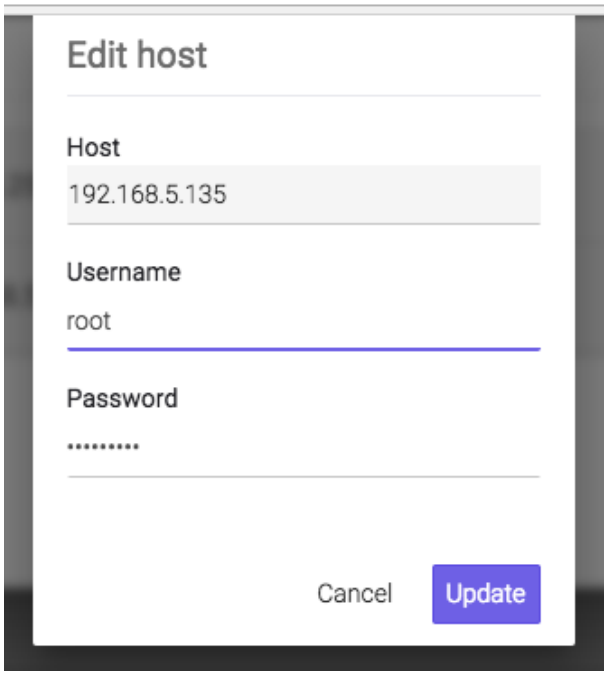

After hypervisor password is updated successfully, a confirmation message is displayed.

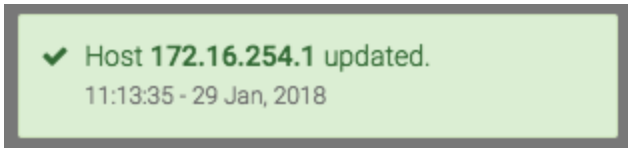

Alerts and notifications

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.

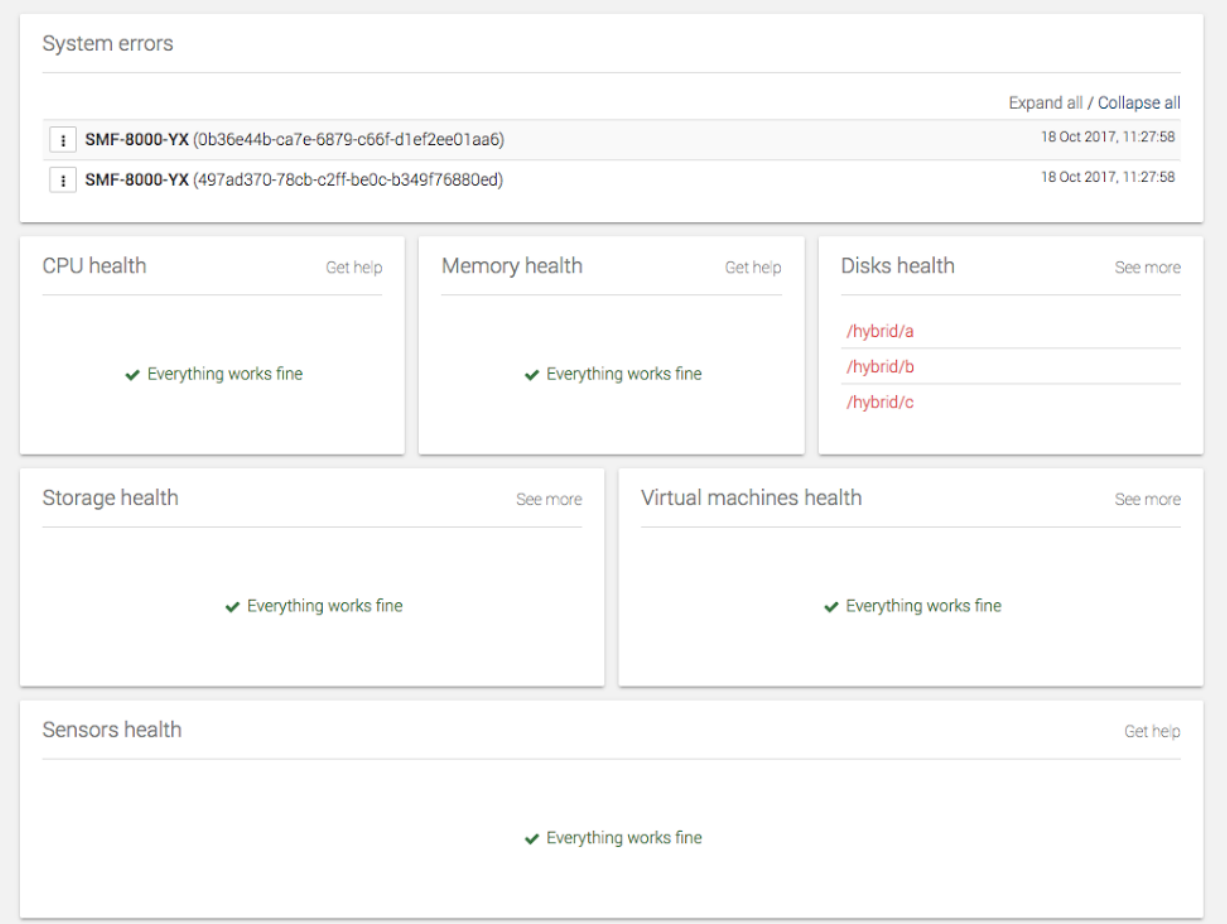

To see the *Alerts* page, go to action bar and select the *bell*

.

Here, you will find information about the *System errors*, *CPU health, Memory health, Disks health, Storage health, Virtual machines health and Sensors health*.

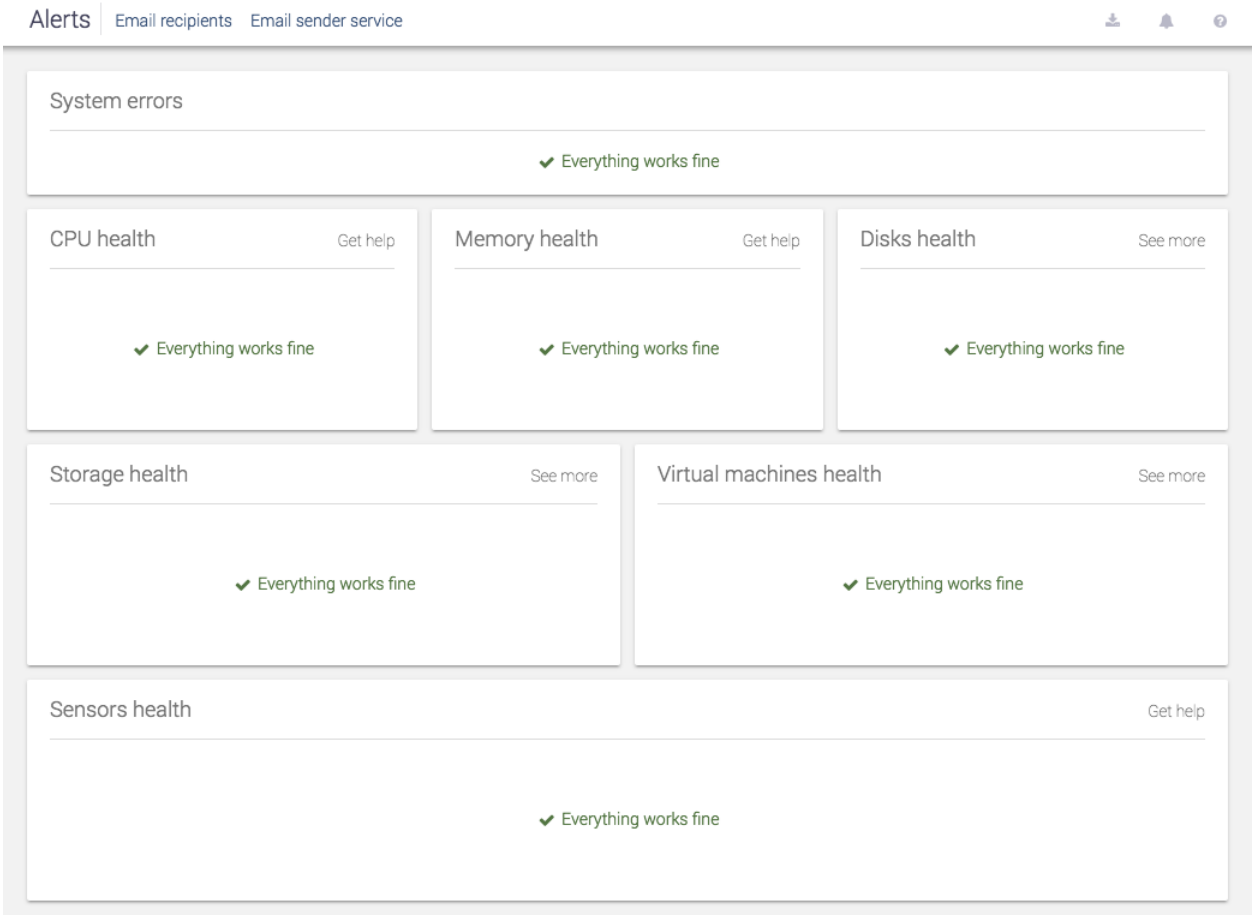

You can access support by following the *Get help* links , or you can go directly to the advised page by following the *See more* links.

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.

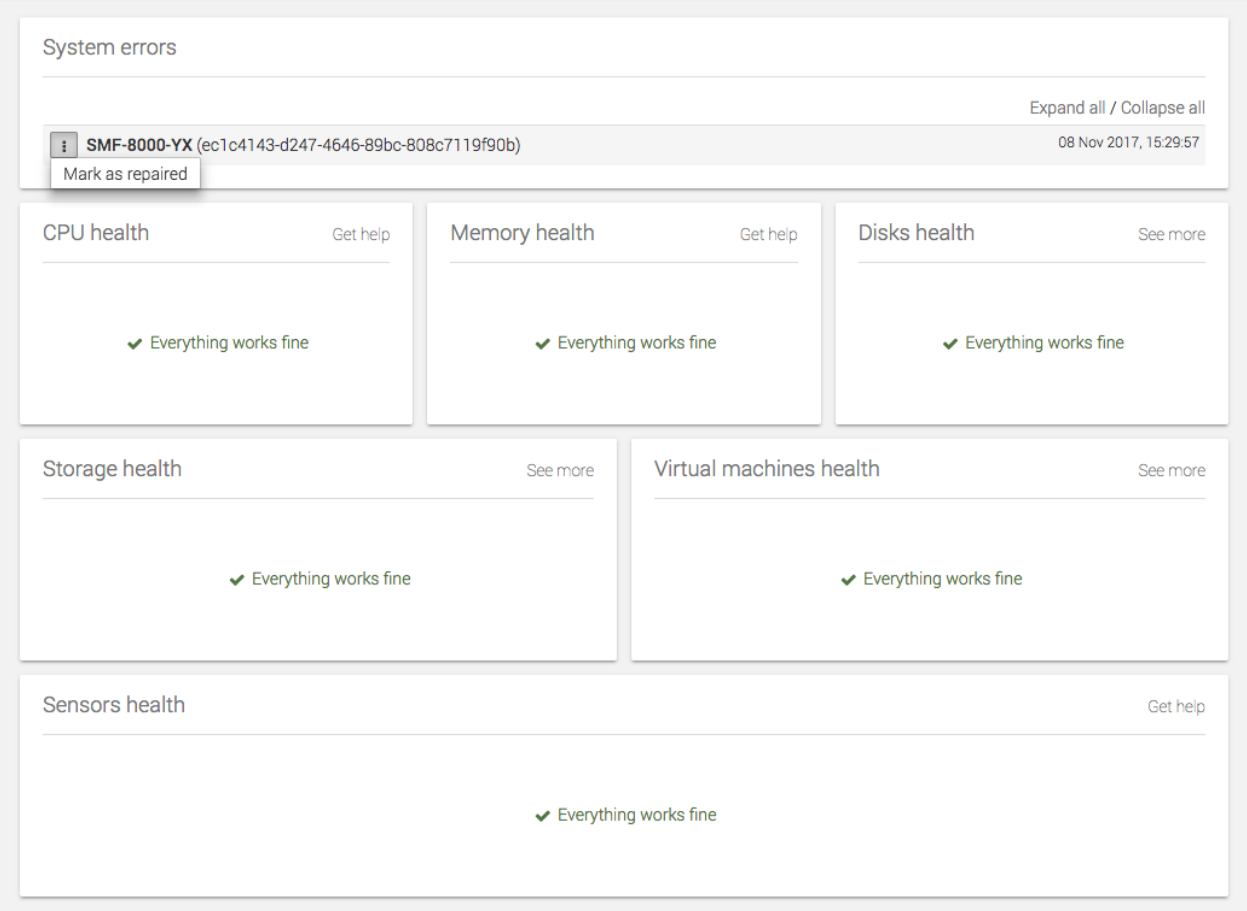

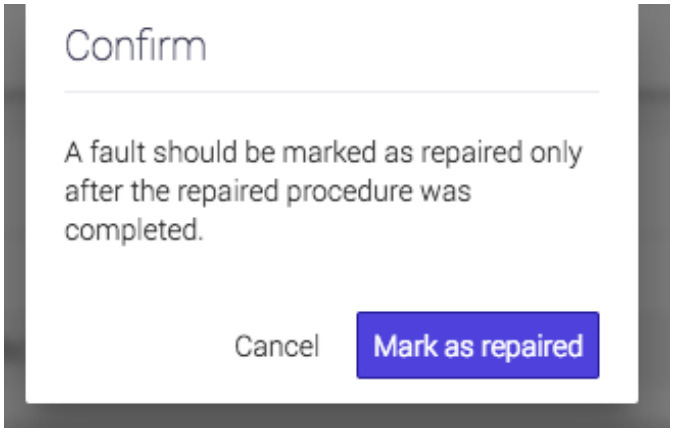

NOTE: It is highly recommended that you introduce at least one email address in order to receive fast and accurate emails in case there are any issues on your machine.

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.

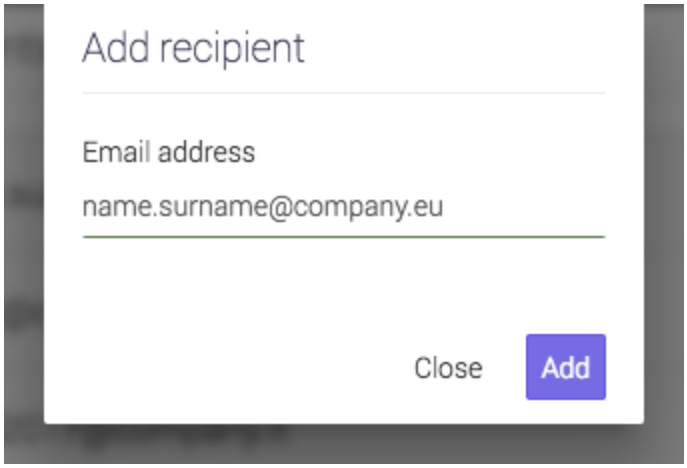

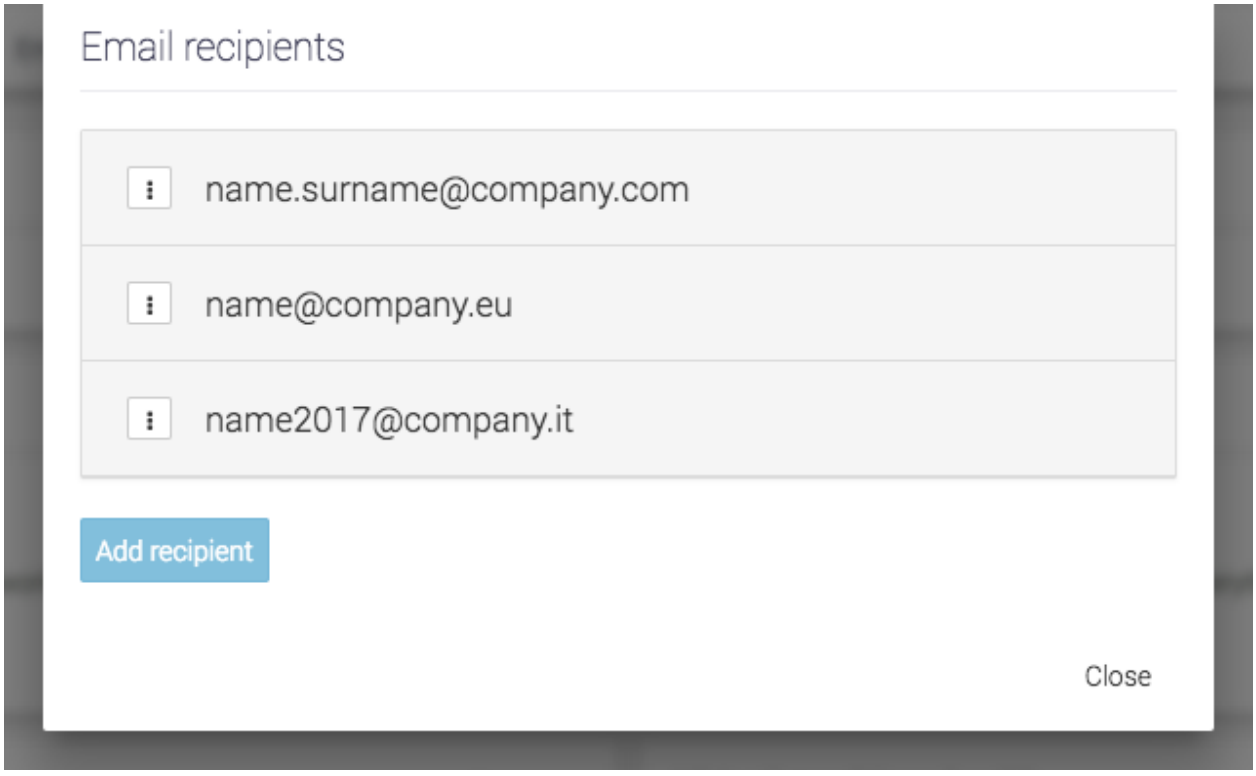

In order to check that your email can be reached, go to *Email sender service* → *Send test email*.

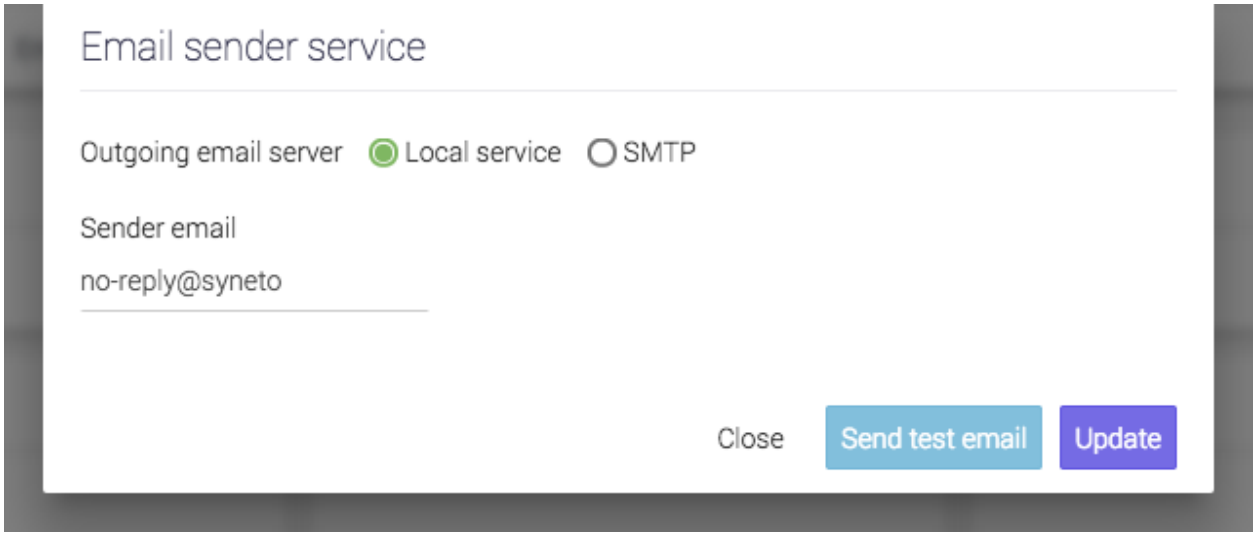

You will receive an email on all added addresses.

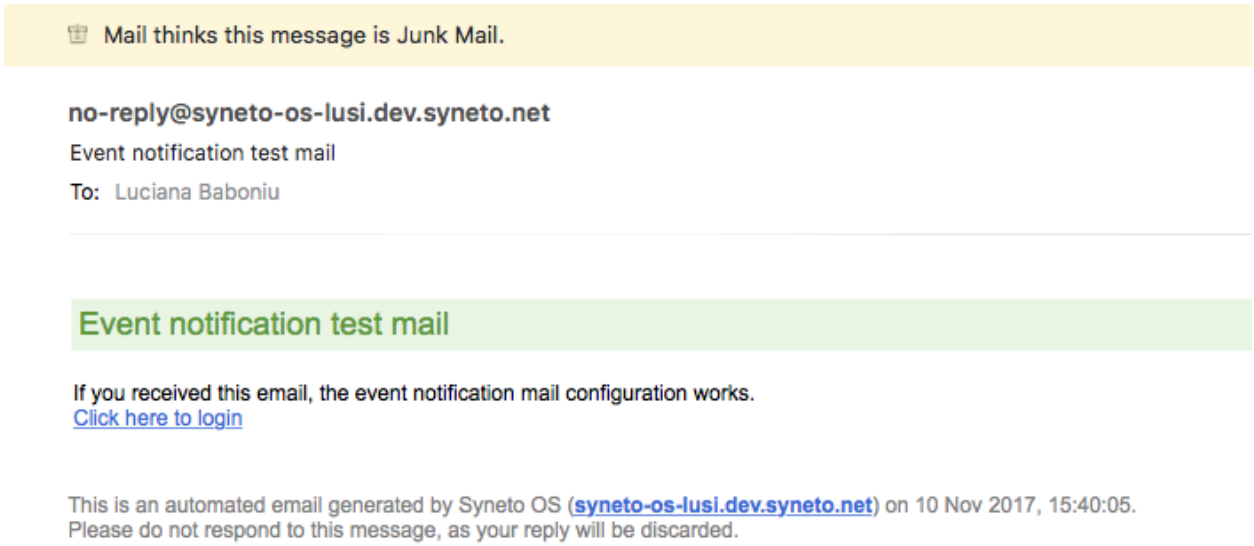

New Software Update

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.

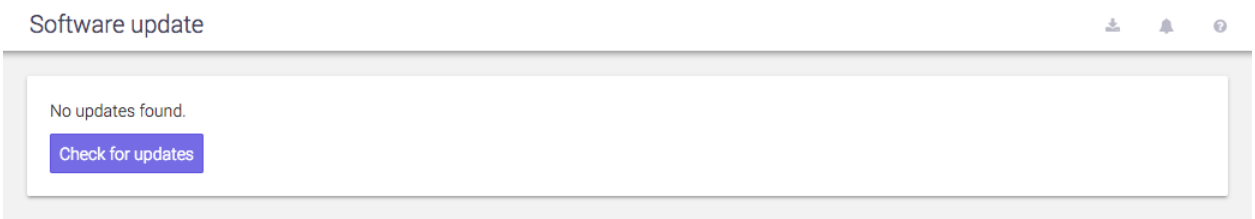

You can read more information about New Software Update in Software updates Chapter.

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.

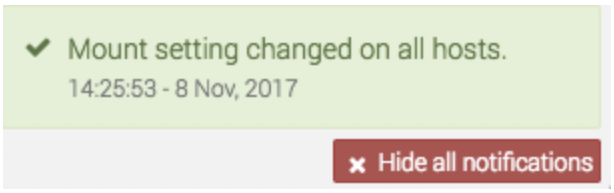

- *Red boxes* - when an action is not successful.

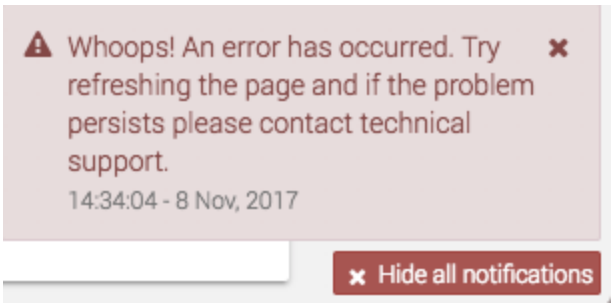

- *Yellow boxes* - when there is a connectivity issue or when the session expired, and the user will be redirected to the authentication page.

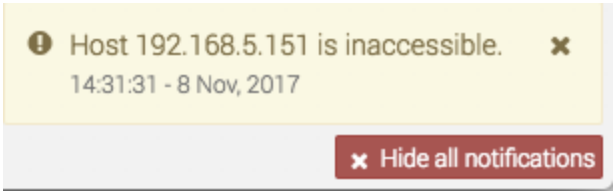

- Blue boxes - when an action is received and underway.

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 disk 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 *Disks → Average service time*. Then add the *Disks → IOPS* graphics. Then add the *Disks → Transactions* graphics.

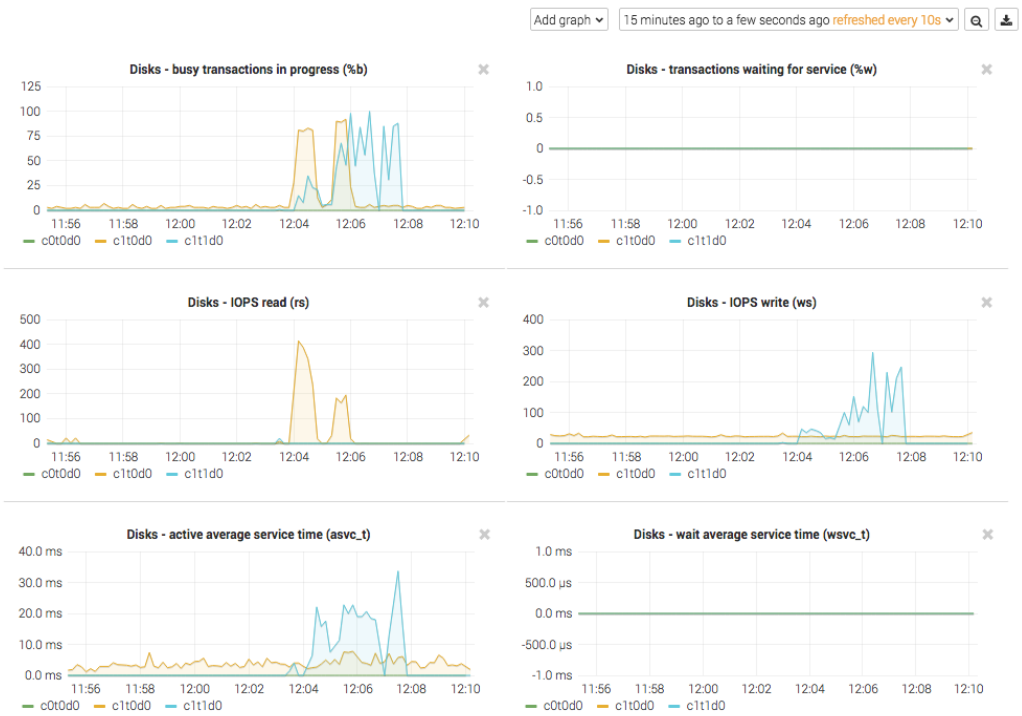

The most important indicator is the *Disks - active average service time (asvc_t)*. This represents the time required to process active transactions. But looking at it in isolation is not enough. There are legitimate cases for this value to be high. That is why we need other graphs as well. We need context.

In the images above we are interested in the blue lines on the graphs. We need to correlate *asvc_t* with *Disks - IOPS write (ws)* and with *Disks - 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 disk 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 *Disks - 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 disk is busy, but it is not doing a large amount of reads or writes, something seems to be wrong. The disk may need to be replaced.

Mid level analytics

Disks 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 disks, you can analyze pool performance with several graphics and correlate between them.

The most important indicators are similar to disks: service time, IOPS, and transaction.

But pools are more complex than simple disks, 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 Technical Support.

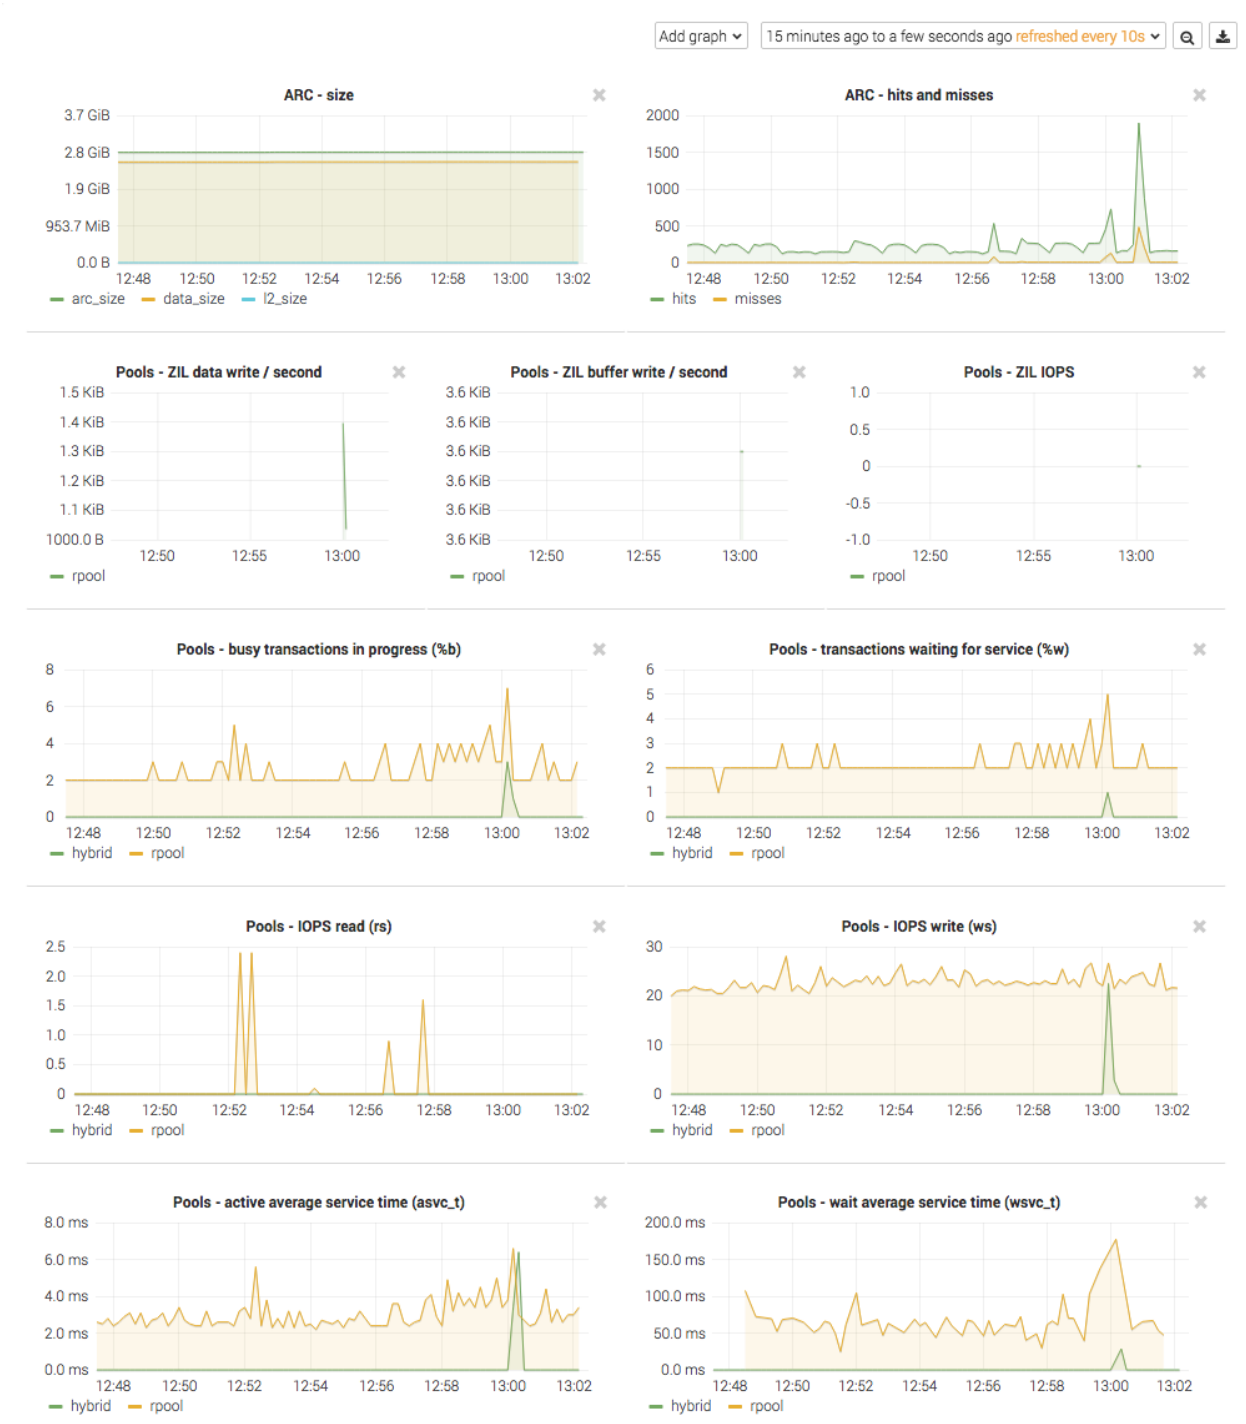
Add graph
15 minutes ago to a few seconds ago refreshed every 10s
ARC - size
3.7 GiB
2.8 GiB
1.9 GiB
953.7 MiB
0.0 B
12:48 12:50 12:52 12:54 12:56 12:58 13:00 13:02
arc_size data_size l2_size
ARC - hits and misses
2000
1500
1000
500
0
12:48 12:50 12:52 12:54 12:56 12:58 13:00 13:02
hits misses
Pools - ZIL data write / second
1.5 KiB
1.4 KiB
1.3 KiB
1.2 KiB
1.1 KiB
1000.0 B
12:50 12:55 13:00
rpool
Pools - ZIL buffer write / second
3.6 KiB
3.6 KiB
3.6 KiB
3.6 KiB
3.6 KiB
3.6 KiB
12:50 12:55 13:00
rpool
Pools - ZIL IOPS
1.0
0.5
0
-0.5
-1.0
12:50 12:55 13:00
rpool
Pools - busy transactions in progress (%b)
8
6
4
2
0
12:48 12:50 12:52 12:54 12:56 12:58 13:00 13:02
hybrid rpool
Pools - transactions waiting for service (%w)
6
5
4
3
2
1
0
12:48 12:50 12:52 12:54 12:56 12:58 13:00 13:02
hybrid rpool
Pools - IOPS read (rs)
2.5
2.0
1.5
1.0
0.5
0
12:48 12:50 12:52 12:54 12:56 12:58 13:00 13:02
hybrid rpool
Pools - IOPS write (ws)
30
20
10
0
12:48 12:50 12:52 12:54 12:56 12:58 13:00 13:02
hybrid rpool
Pools - active average service time (asvc_t)
8.0 ms
6.0 ms
4.0 ms
2.0 ms
0.0 ms
12:48 12:50 12:52 12:54 12:56 12:58 13:00 13:02
hybrid rpool
Pools - wait average service time (wsvc_t)
200.0 ms
150.0 ms
100.0 ms
50.0 ms
0.0 ms
12:48 12:50 12:52 12:54 12:56 12:58 13:00 13:02
hybrid rpool

Case 1 - I read a lot of data, but I don't see any read activity in analytics

Syneto HYPER employes 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 disk 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 disks

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 disk 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 disk activity.

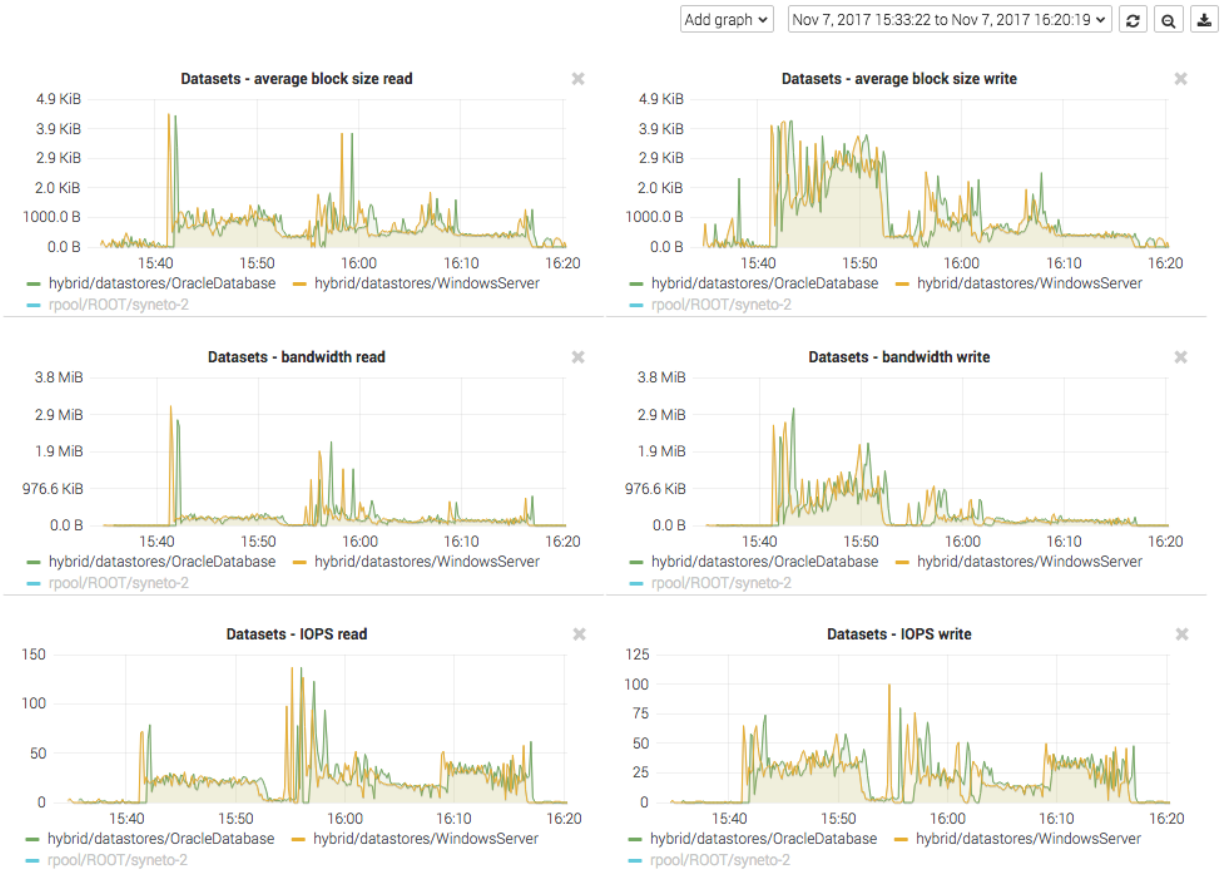

In the image you can see two virtual machines competing for dataset IOPS and bandwidth. If one of those virtual machine would take all IOPS, the other would have no resources left to operate properly.

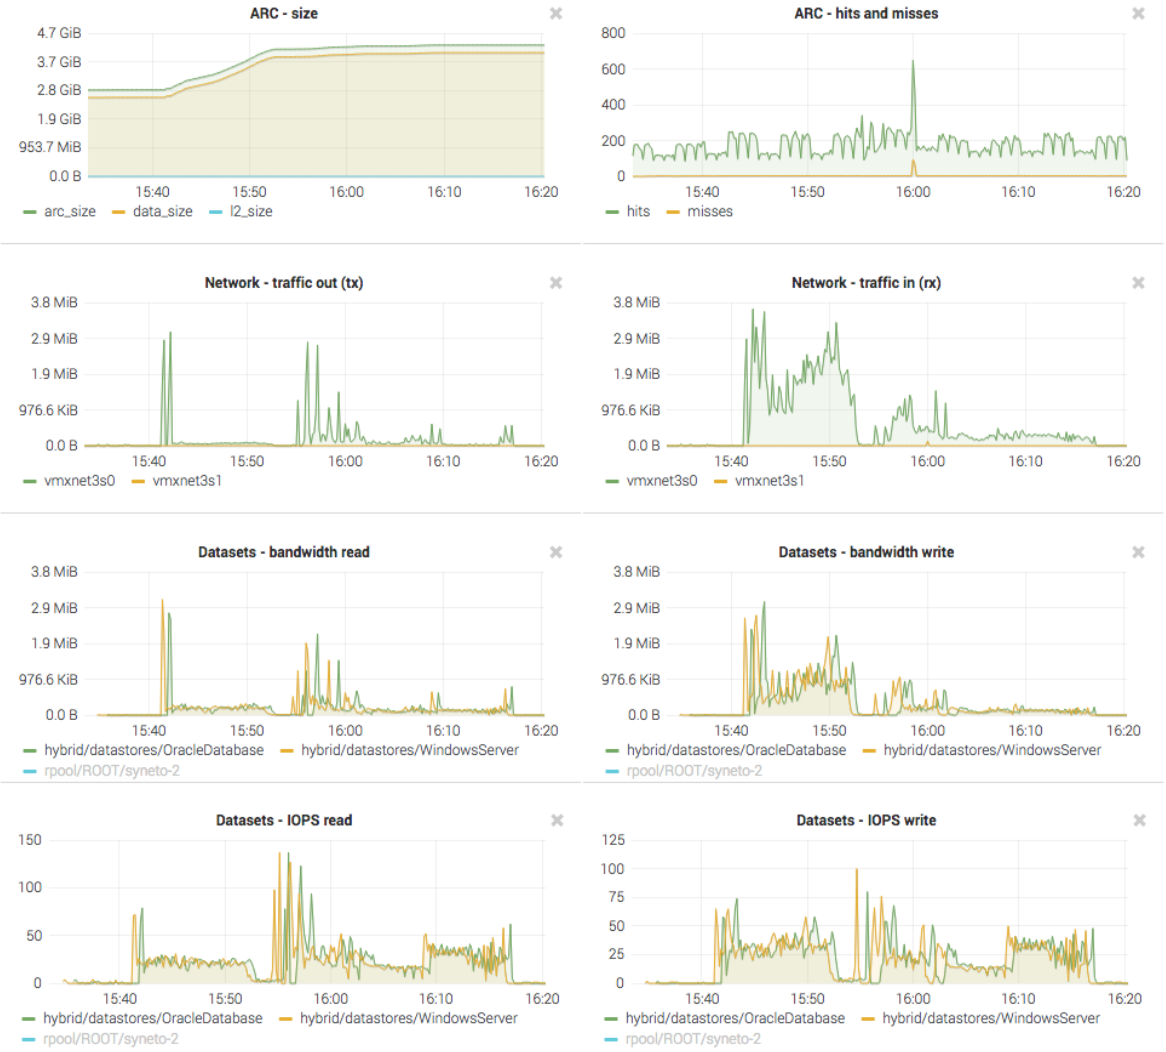

Additionally you can add networking graphs and correlate virtual machine activity with network activity and ARC size and hits graphs.

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*.

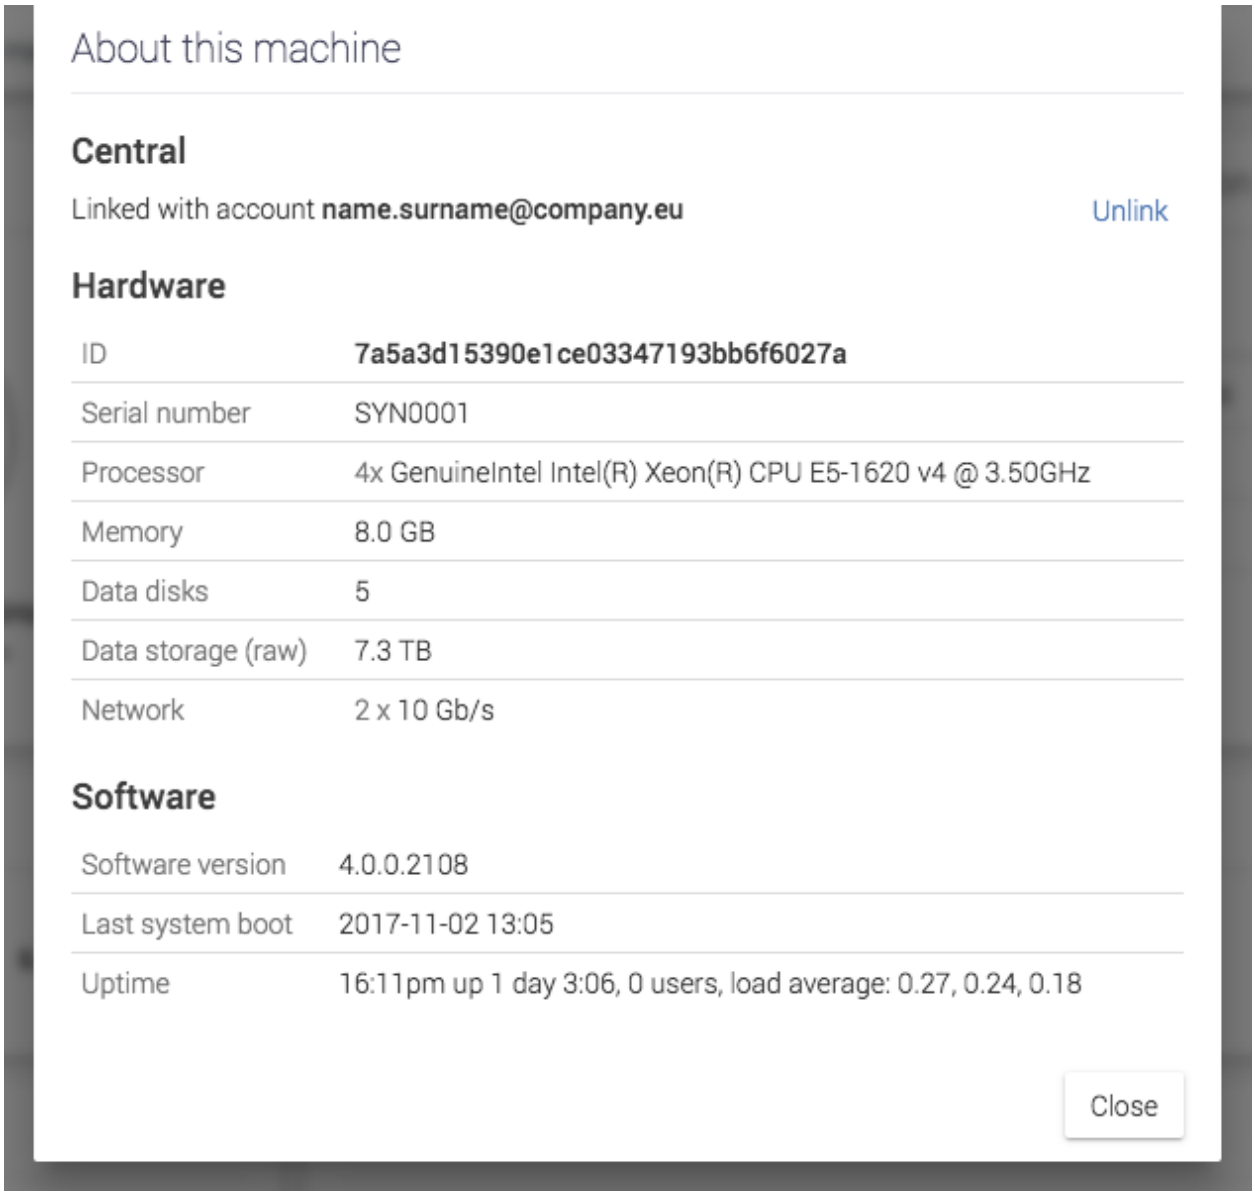

Internet access is required to check for updates. If the storage needs a *HTTP proxy* to access the internet, configure the proxy on the *Network → Basic configuration* page.

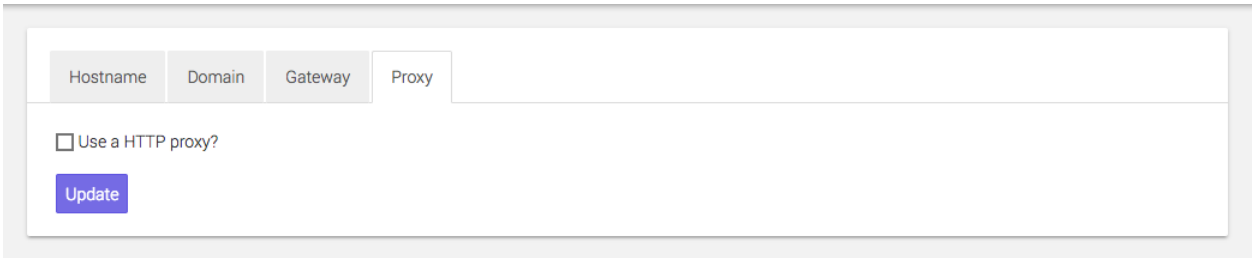

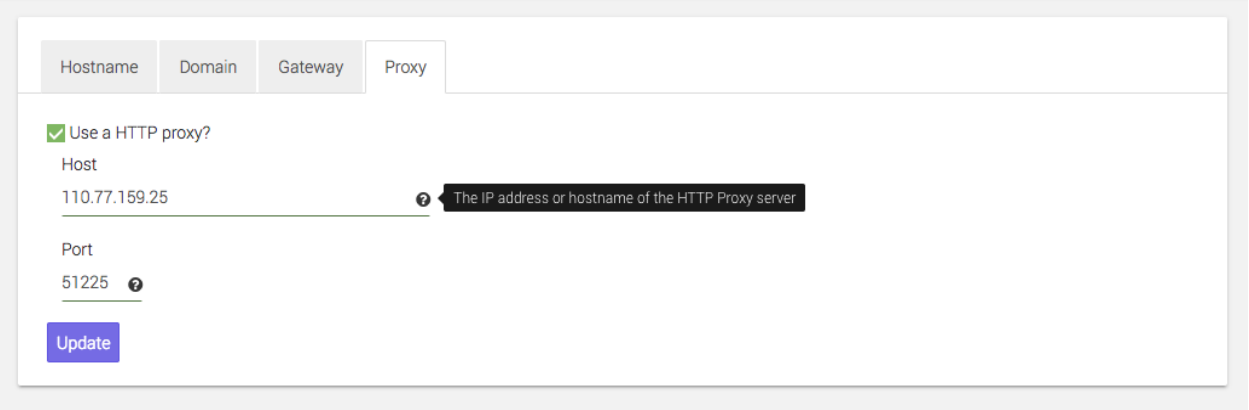

When updates are available, the icon in the top right menu will be highlighted . Click it in order to start the update process. If a reboot is required, you will be told so by a tag next to the update option.

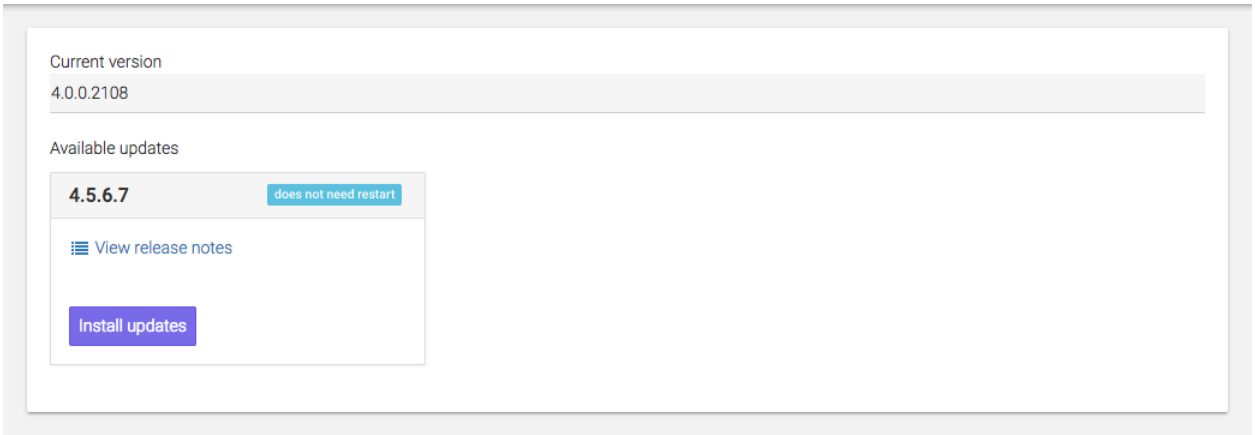

Update without restart

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.

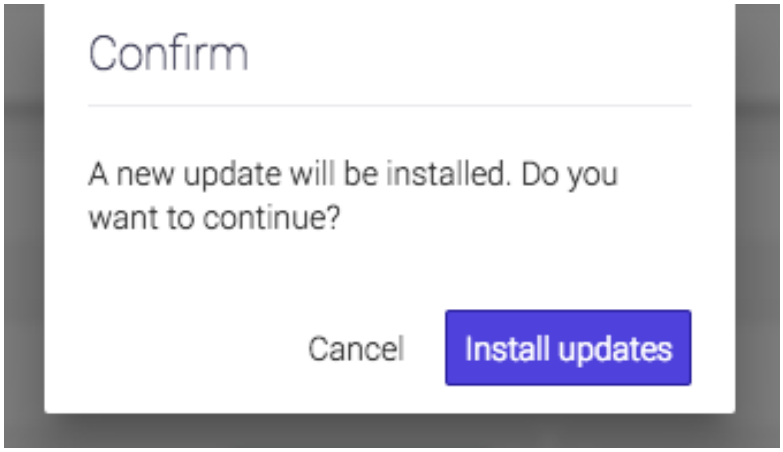

3. The update to the new version starts. Based on your internet speed, downloading the new version can take up to tens of minutes.

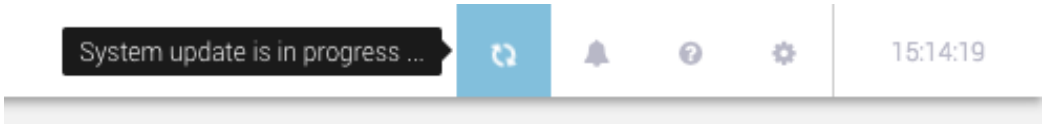

4. Wait for the update to finish.

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.

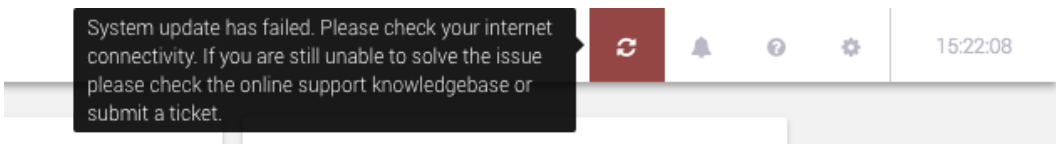

Update with restart

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.

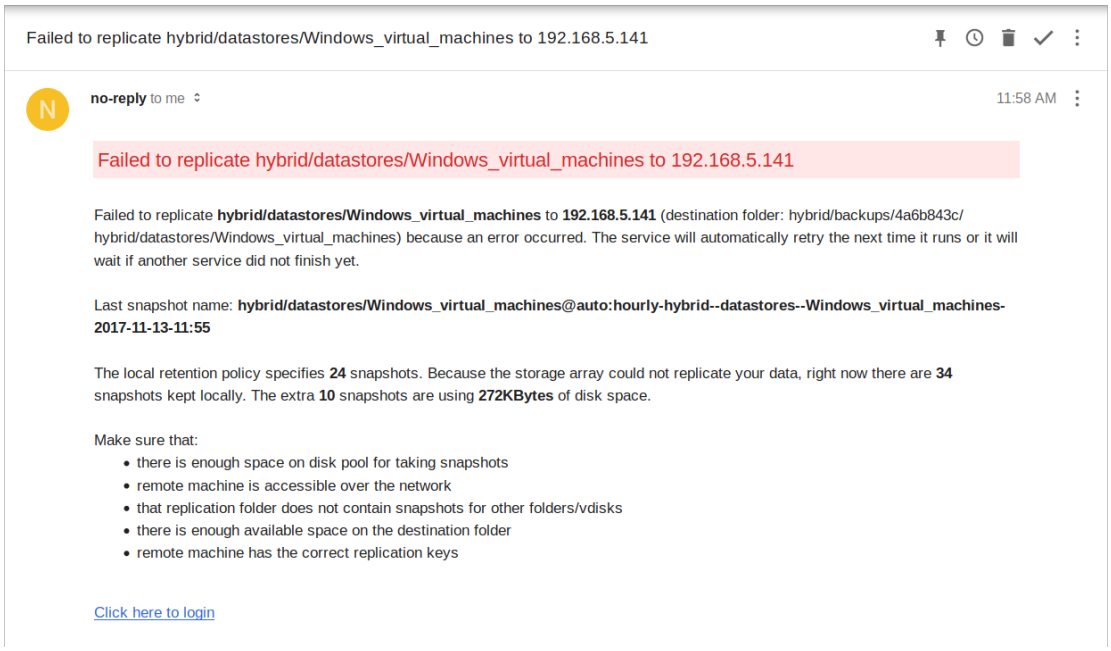

Failed to replicate hybrid/datastores/Windows_virtual_machines to 192.168.5.141

no-reply to me 11:58 AM

Failed to replicate hybrid/datastores/Windows_virtual_machines to 192.168.5.141

Failed to replicate **hybrid/datastores/Windows_virtual_machines** to **192.168.5.141** (destination folder: hybrid/backups/4a6b843c/hybrid/datastores/Windows_virtual_machines) because an error occurred. The service will automatically retry the next time it runs or it will wait if another service did not finish yet.

Last snapshot name: **hybrid/datastores/Windows_virtual_machines@auto:hourly-hybrid--datastores--Windows_virtual_machines-2017-11-13-11:55**

The local retention policy specifies **24** snapshots. Because the storage array could not replicate your data, right now there are **34** snapshots kept locally. The extra **10** snapshots are using **272KBytes** of disk space.

Make sure that:

- there is enough space on disk pool for taking snapshots
- remote machine is accessible over the network
- that replication folder does not contain snapshots for other folders/vdisks
- there is enough available space on the destination folder
- remote machine has the correct replication keys

Click here to login

First of all, check the replication target is up. It may be offline (power outage, hardware failure, system error).

If it is up, check that the machine which replicates can connect to it. You can 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 2 machines.

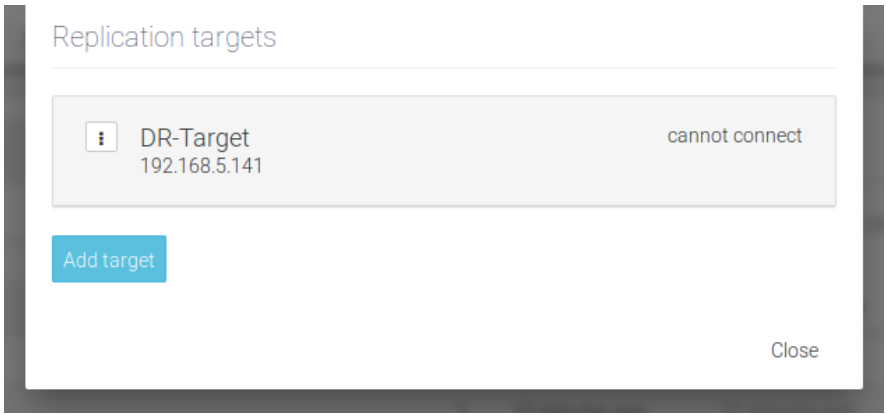

If it can connect to it, check the available space on the machines (both the primary and disaster-recovery). This can be done by checking the *Pools* page.

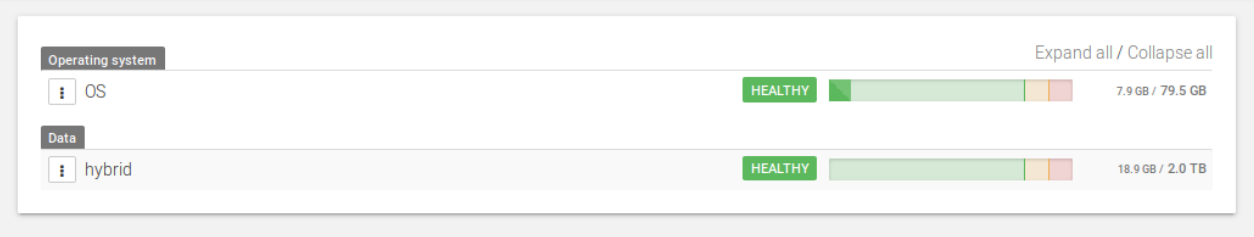

If it has enough space, remove the replication target and try adding it again. Wait for the next scheduled replication or click on *Run* for the desired schedule to start now.

If it still fails, please contact support.

Cannot enable technical support

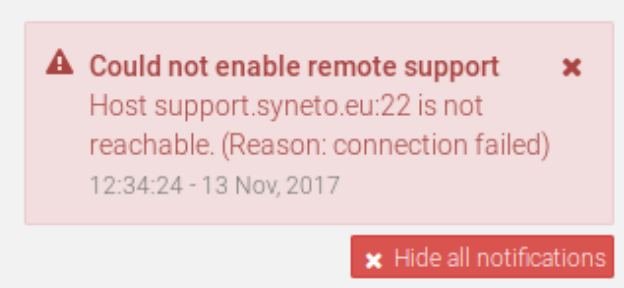

If you can't enable remote access for Syneto support, it means there is a network connectivity problem. First of all, check the internet connection is up. If you are on the same network as the machine, you can simply try to connect to a website.

If the internet connection is up, check the network configuration of the SynetoOS. Verify 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

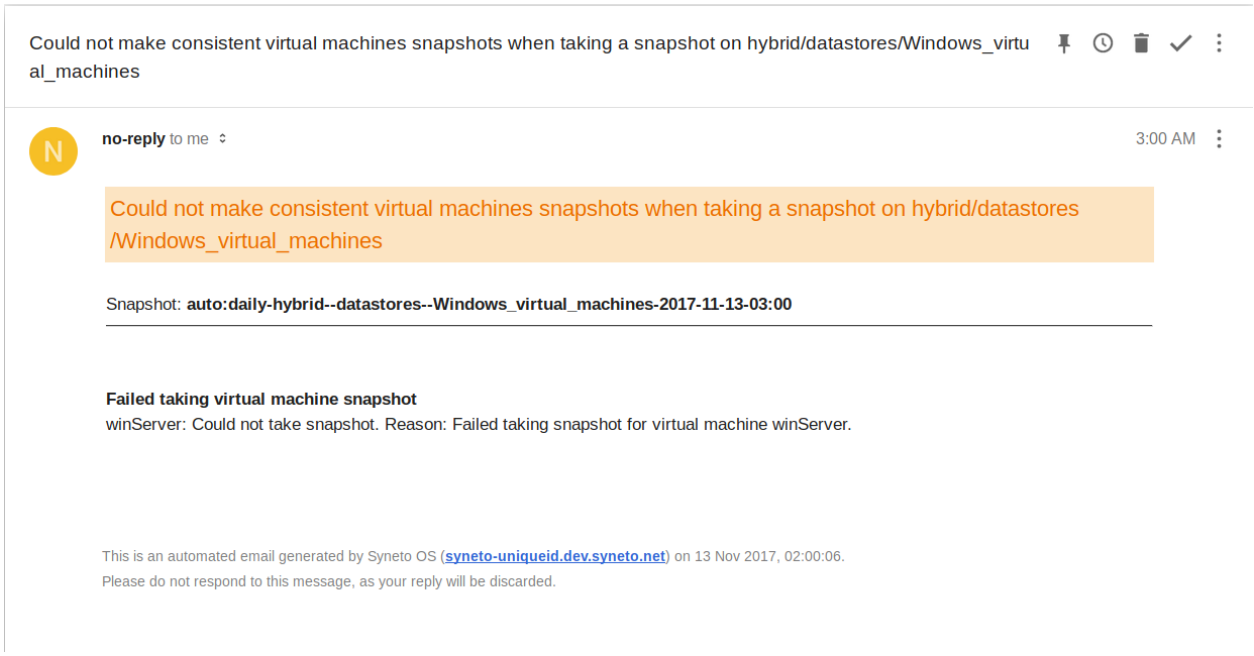

There are two situations that can cause this problem.

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.

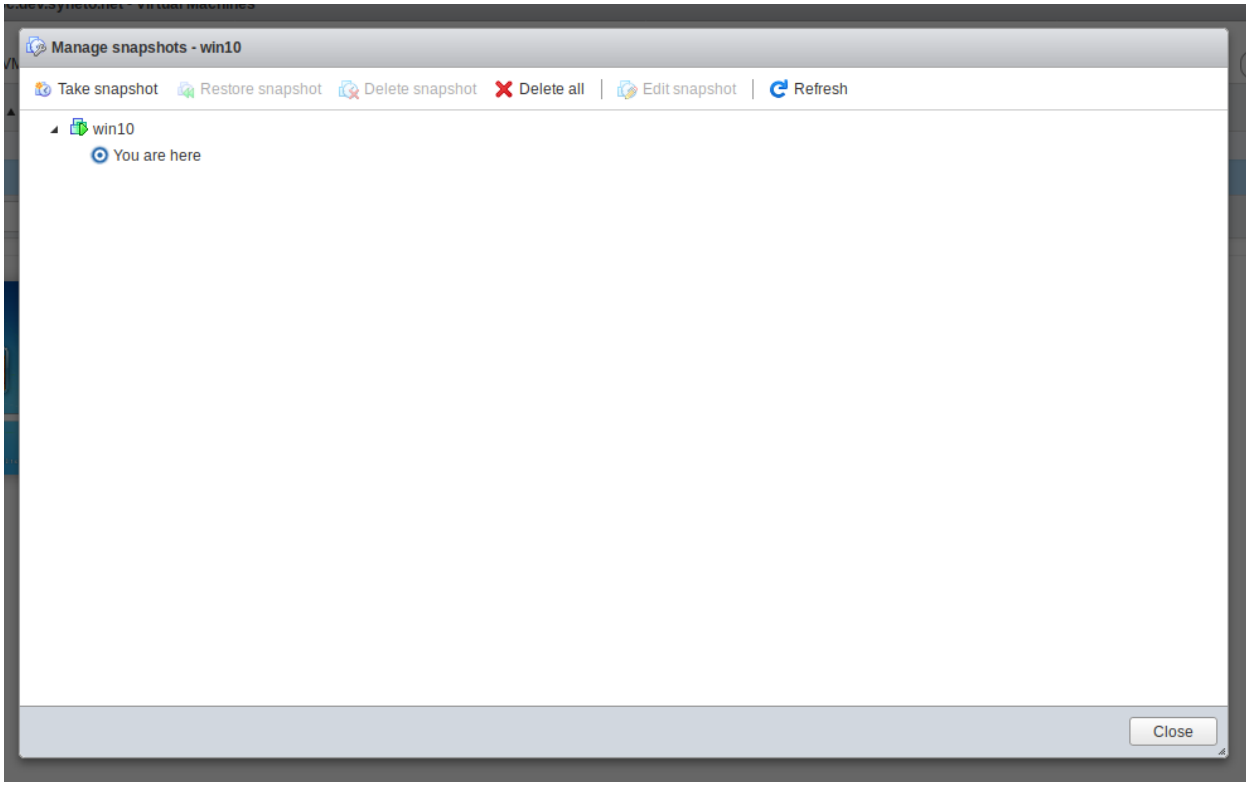

After reinstall data pool cannot be imported

This affects **Syneto HYPER 2000 Series** and **Syneto HYPER DR Play Series**. The write acceleration disk is a virtual one, and after a reinstall it will be recreated. As a consequence, the pool cannot find the old disk.

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). Click the context menu button and select *Manage caches*. From the *Write acceleration disk group* remove all

disks. Click *Update*. Go to *Manage caches* again and select the disk that has 8GB (usually c1t1d0) and then click *Update*.

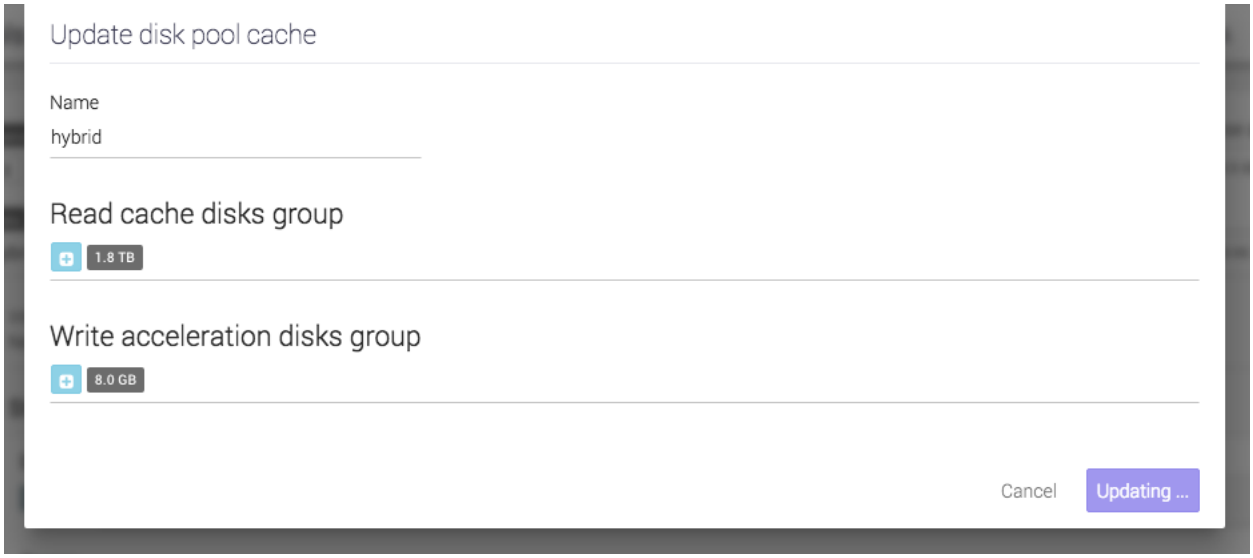

Error when creating a datastore, share, or volume

When you try to create a datastore, share, or volume, you may see an error similar to this:

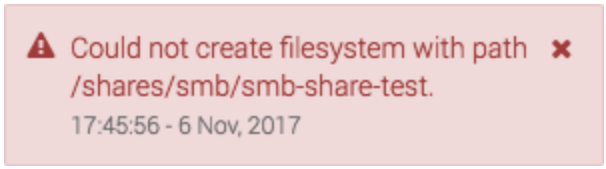

Make sure there is a data pool imported. Go to *Secondary menu → Pools* and confirm you can see data pools (other than the OS pools). If no data pool is present, please create a new one or import an existing one.

If the problem persists, contact technical 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 ESX 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.

Dialog appears asking for the new local ESXi password

If an 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 Syneto OS web interface.

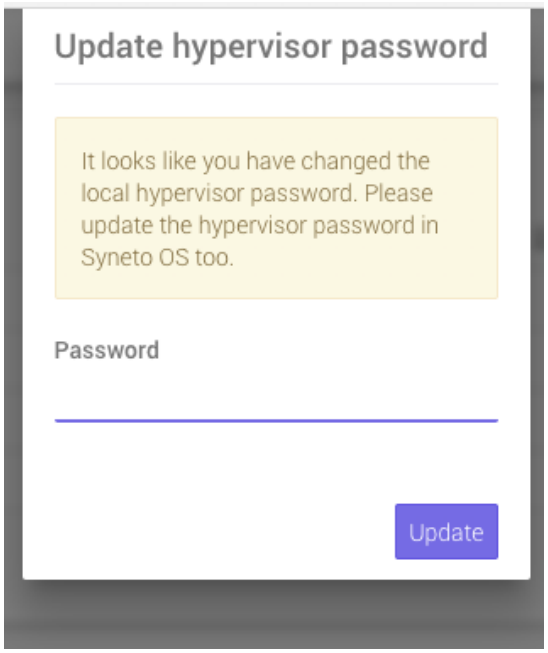

APPENDIX A: CREATING A CENTRAL ACCOUNT

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.

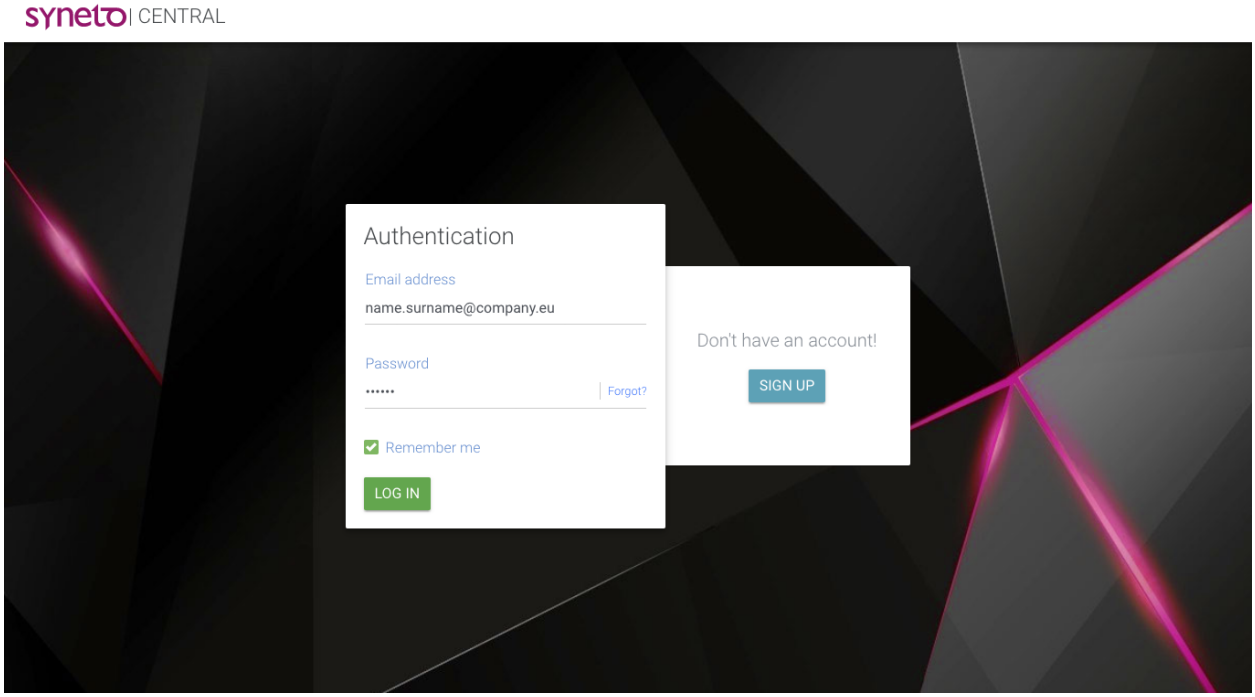

Type in the email address of the person who will be using the product and choose a password. Also insert a backup email address. Complete the name and phone number for the person in charge with this account, but also complete the information fields for your company.

WARNING: Please make sure that you completed all fields correctly. Failure to do so might prevent you from obtaining full support on your product.

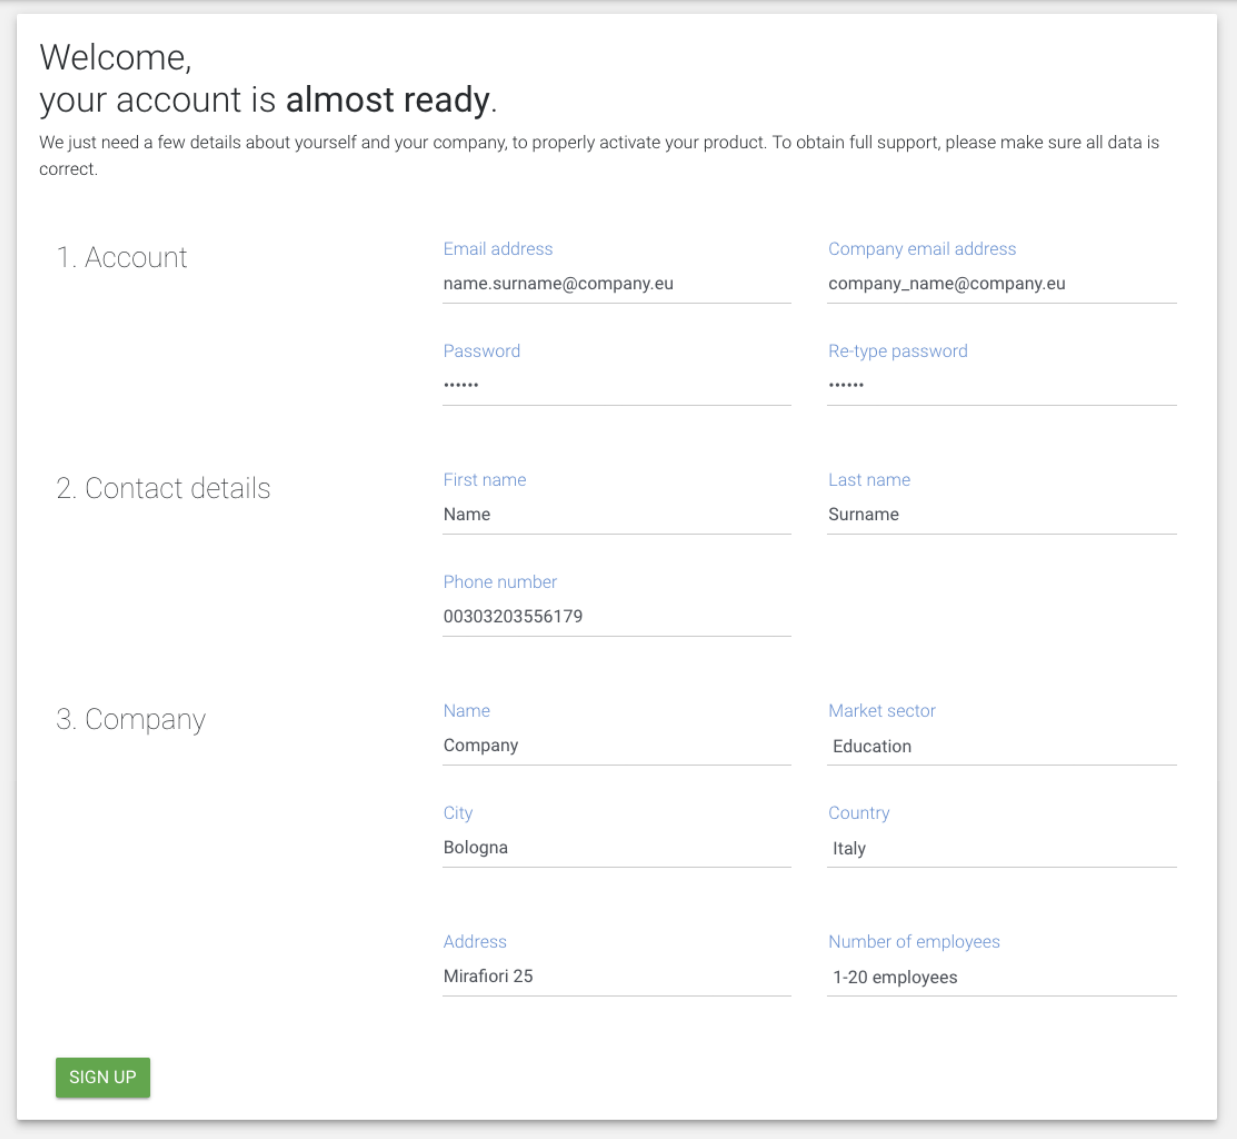

After all the required fields are completed, press *Sign up*. Your Central account is now ready. You can start activating your products.

For more information about activating a product with Central, please go to Activate using a Central account chapter.

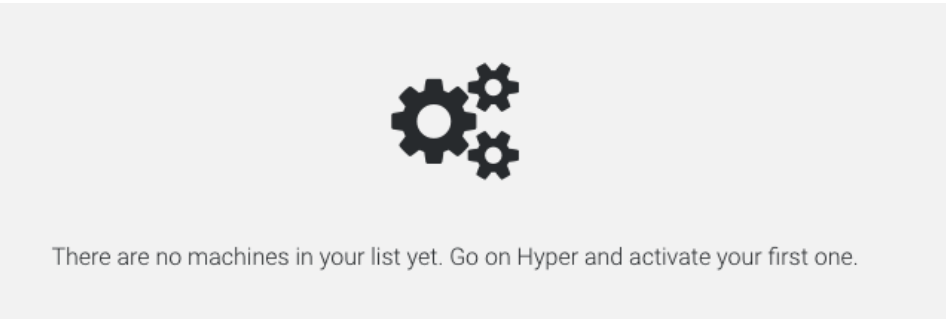

After you start activating products, they will all show up on the page with information about their *Name, Support type, Expiration date, Hardware ID and Model* .

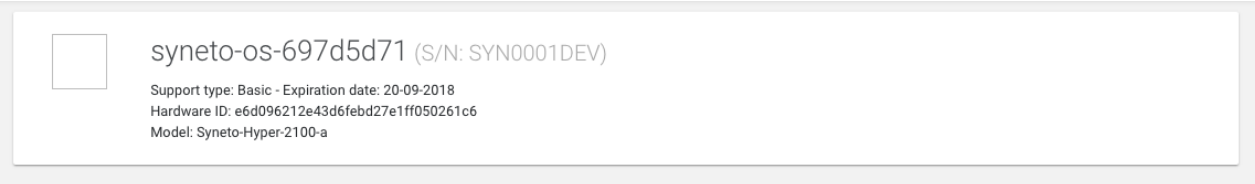

If you forgot your password to the Central account, from the login page, go to *Forgot? Link . A new page will open , where you insert your email address and click on SEND RECOVERY INSTRUCTIONS.*

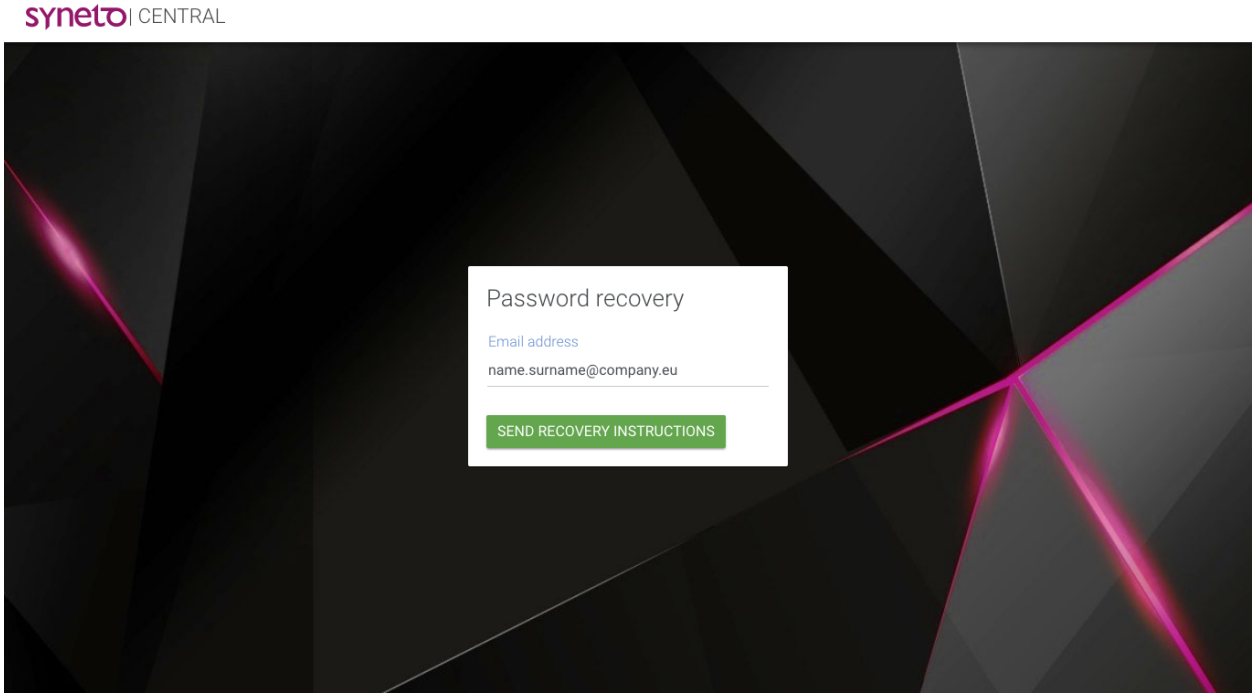

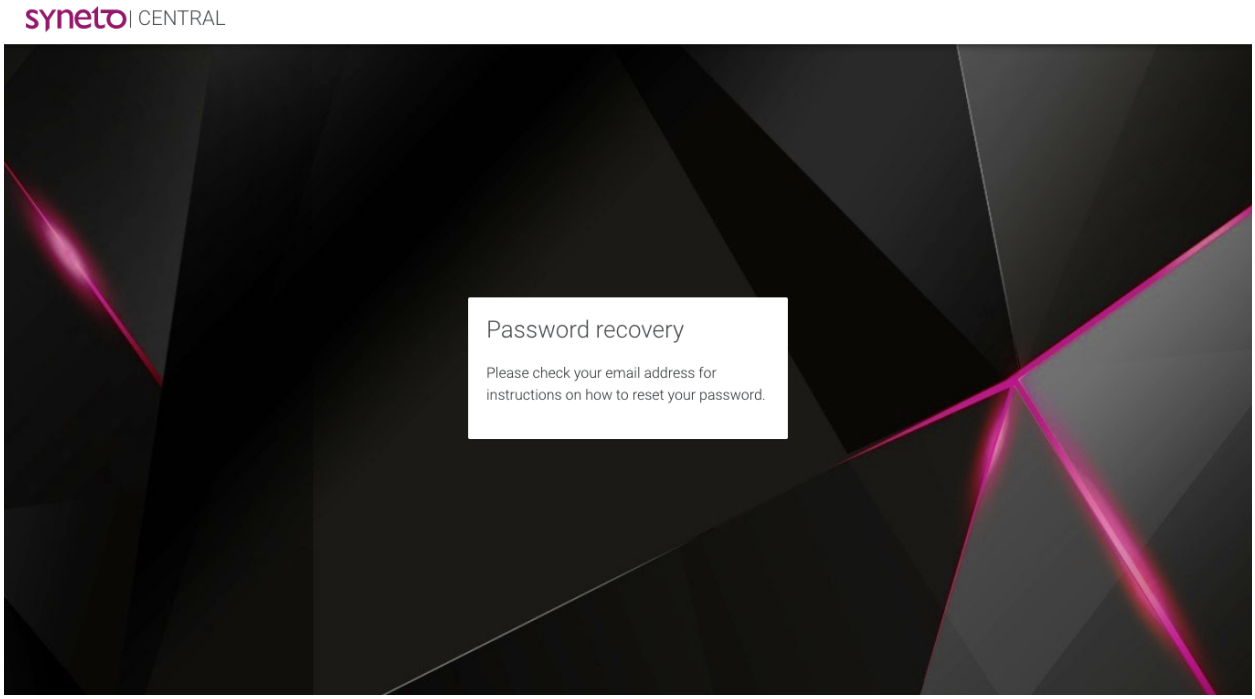

An email with password recovery instructions will be sent to the provided email address.

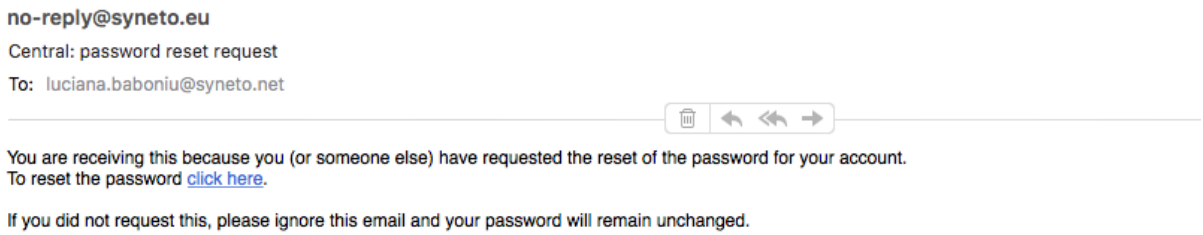

Follow the *click here* link in the email, and you will be redirected to a new page, here you can insert a new password. Make sure you confirm it by inserting it twice and then click on CHANGE PASSWORD. You can now log in to your Central account with the same username and the newly changed password.

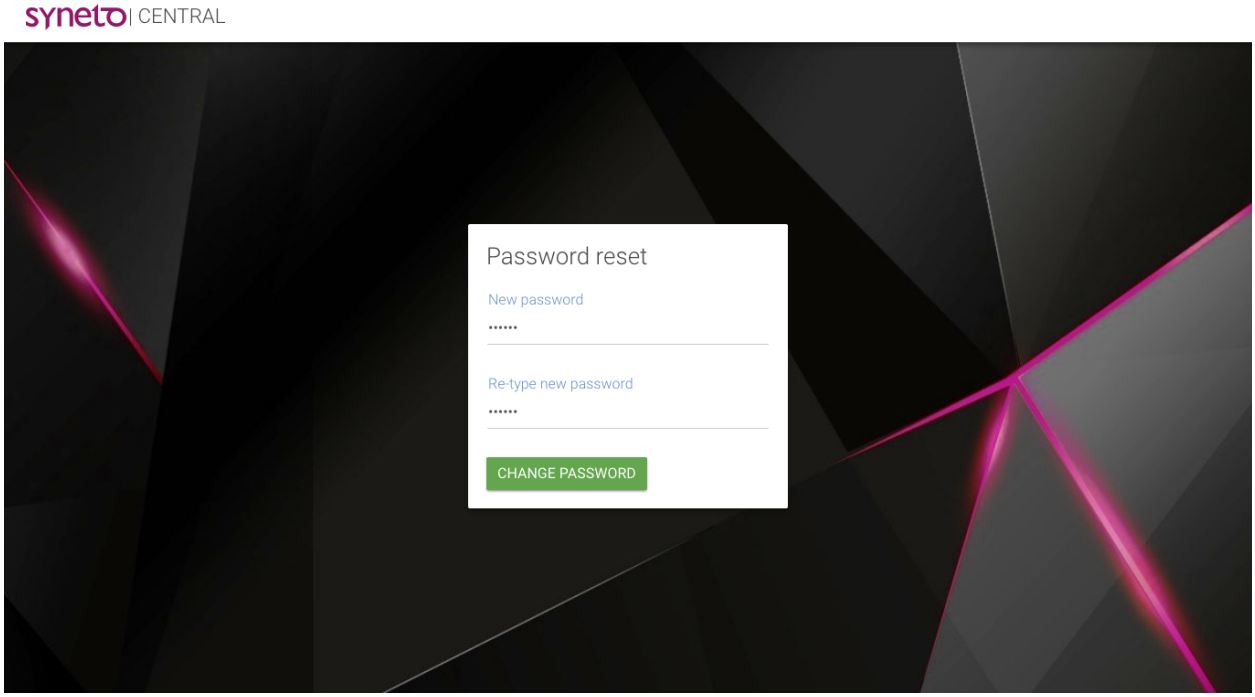
CENTRAL
Password reset
New password
Re-type new password
CHANGE PASSWORD

APPENDIX B: ACCESSING TECHNICAL 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.*

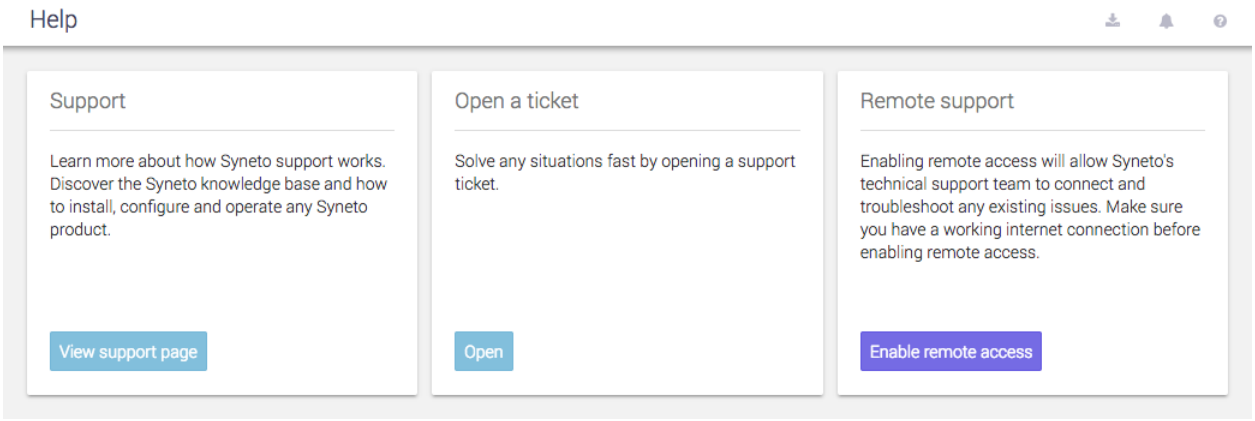

Because your product is registered with Central, when opening a ticket, the data will be already filled with the information provided in your *Central profile*. All you have to do now, is *Describe your request*.

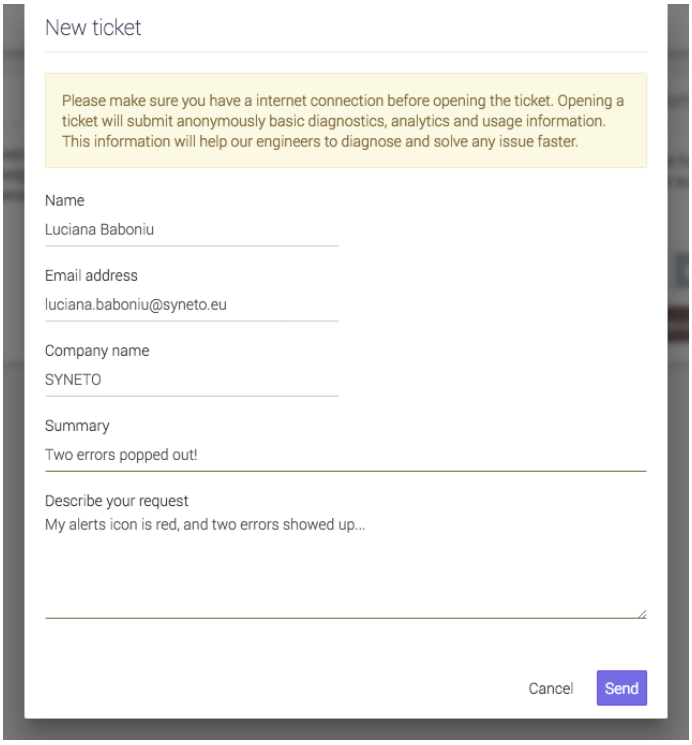

Enabling remote access, will allow our support engineers to connect and troubleshoot any existing issues.Make sure you have a working internet connection before enabling remote access. You can disable the remote access anytime you want.

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.

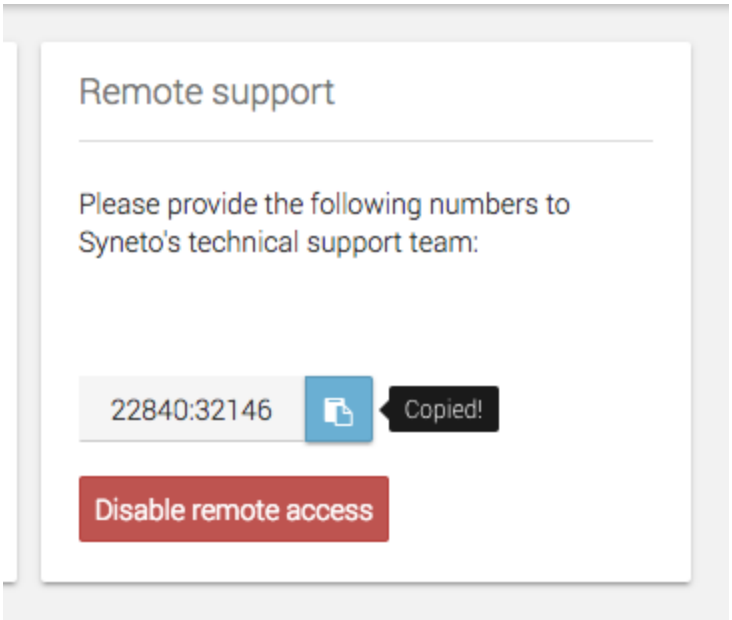

You can also access support when you see *Alerts*, by simply following the *Get help* links you can find on the *Alerts page*.

NOTE: Please note that the Support services differ with the Maintenance type you have.

Contacting Syneto Technical 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.