



Proxmox Backup Documentation

Release 0.9.2-2

Proxmox Support Team

Wednesday, 28 October 2020

TABLE OF CONTENTS

1	Introduction	3
1.1	What is Proxmox Backup Server?	3
1.2	Architecture	3
1.3	Main Features	3
1.4	Reasons for Data Backup?	4
1.5	Software Stack	4
1.6	Getting Help	5
1.6.1	Community Support Forum	5
1.6.2	Mailing Lists	5
1.6.3	Bug Tracker	5
1.7	License	5
1.8	History	5
2	Installation	7
2.1	System Requirements	7
2.1.1	Minimum Server Requirements, for Evaluation	7
2.1.2	Recommended Server System Requirements	7
2.1.3	Supported Web Browsers for Accessing the Web Interface	8
2.2	Debian Package Repositories	8
2.2.1	SecureApt	8
2.2.2	Proxmox Backup Beta Repository	9
2.3	Server installation	9
2.3.1	Install Proxmox Backup with the Installer	9
2.3.2	Install Proxmox Backup server on Debian	10
2.3.3	Install Proxmox Backup server on Proxmox VE	10
2.4	Client installation	10
2.4.1	Install Proxmox Backup client on Debian	10
2.4.2	Installing from source	11
2.4.3	Installing statically linked binary	11
3	Terminology	13
3.1	Backup Content	13
3.1.1	Image Archives: <name>.img	13
3.1.2	File Archives: <name>.pxar	13
3.1.3	Binary Data (BLOBs)	13
3.1.4	Catalog File: catalog.pcat1	13
3.1.5	The Manifest: index.json	14
3.2	Backup Type	14
3.3	Backup ID	14
3.4	Backup Time	14
3.5	Backup Group	14
3.6	Backup Snapshot	14
4	Graphical User Interface	15

4.1	Features	15
4.2	Login	15
4.3	GUI Overview	16
4.4	Sidebar	16
4.4.1	Dashboard	16
4.4.2	Configuration	16
4.4.3	Administration	17
4.4.4	Datastore	18
5	Storage	19
5.1	Disk Management	19
5.2	DataStore	21
5.2.1	Datastore Configuration	22
6	Network Management	25
7	User Management	27
7.1	User Configuration	27
7.2	Access Control	28
8	Managing Remotes	31
8.1	Remote	31
8.2	Sync Jobs	32
9	Maintenance Tasks	33
9.1	Garbage Collection	33
10	Backup Client Usage	35
10.1	Repository Locations	35
10.2	Environment Variables	35
10.3	Output Format	36
10.4	Creating Backups	36
10.4.1	Excluding files/folders from a backup	37
10.5	Encryption	38
10.5.1	Using a master key to store and recover encryption keys	38
10.6	Restoring Data	40
10.6.1	Interactive Restores	40
10.6.2	Mounting of Archives via FUSE	41
10.7	Login and Logout	41
10.8	Pruning and Removing Backups	42
10.9	Garbage Collection	43
10.10	Benchmarking	43
11	Proxmox VE Integration	45
12	pxar Command Line Tool	47
12.1	Description	47
12.2	Creating an Archive	47
12.3	Extracting an Archive	48
12.4	List the Contents of an Archive	48
12.5	Mounting an Archive	48
13	Host System Administration	49
13.1	ZFS on Linux	49
13.1.1	Hardware	50
13.1.2	ZFS Administration	50
14	FAQ	57
14.1	What distribution is Proxmox Backup Server (PBS) based on?	57
14.2	Which platforms are supported as a backup source (client)?	57

14.3	Will Proxmox Backup Server run on a 32-bit processor?	57
14.4	How long will my Proxmox Backup Server version be supported?	57
14.5	Can I copy or synchronize my datastore to another location?	57
14.6	Can Proxmox Backup Server verify data integrity of a backup archive?	58
14.7	When backing up to remote servers, do I have to trust the remote server?	58
14.8	Is the backup incremental/deduplicated?	58
A	Command Syntax	59
A.1	proxmox-backup-client	59
A.1.1	Catalog Shell Commands	65
A.2	proxmox-backup-manager	67
A.3	pxar	77
B	File Formats	79
B.1	Proxmox File Archive Format (.pxar)	79
C	Backup Protocol	81
C.1	Backup Protocol API	81
C.2	Reader Protocol API	83
D	Calendar Events	87
D.1	Introduction and Format	87
D.2	Differences to systemd	88
D.3	Notes on scheduling	88
E	Glossary	89
F	GNU Free Documentation License	91
	Index	97

Copyright (C) 2019-2020 Proxmox Server Solutions GmbH
Version 0.9.2 – Wednesday, 28 October 2020

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

INTRODUCTION

1.1 What is Proxmox Backup Server?

Proxmox Backup Server is an enterprise-class, client-server backup software package that backs up *virtual machines*, *containers*, and physical hosts. It is specially optimized for the *Proxmox Virtual Environment* platform and allows you to back up your data securely, even between remote sites, providing easy management with a web-based user interface.

It supports deduplication, compression, and authenticated encryption (AE). Using *Rust* as the implementation language guarantees high performance, low resource usage, and a safe, high-quality codebase.

Proxmox Backup uses state of the art cryptography for client communication and backup content *encryption*. Encryption is done on the client side, making it safer to back up data to targets that are not fully trusted.

1.2 Architecture

Proxmox Backup Server uses a *client-server model*. The server stores the backup data and provides an API to create and manage datastores. With the API, it's also possible to manage disks and other server-side resources.

The backup client uses this API to access the backed up data. With the command line tool `proxmox-backup-client` you can create backups and restore data. For *QEMU* with *Proxmox Virtual Environment* we deliver an integrated client.

A single backup is allowed to contain several archives. For example, when you backup a *virtual machine*, each disk is stored as a separate archive inside that backup. The VM configuration itself is stored as an extra file. This way, it's easy to access and restore only important parts of the backup, without the need to scan the whole backup.

1.3 Main Features

Support for Proxmox VE The *Proxmox Virtual Environment* is fully supported and you can easily backup *virtual machines* and *containers*.

Performance The whole software stack is written in *Rust*, in order to provide high speed and memory efficiency.

Deduplication Periodic backups produce large amounts of duplicate data. The deduplication layer avoids redundancy and minimizes the storage space used.

Incremental backups Changes between backups are typically low. Reading and sending only the delta reduces the storage and network impact of backups.

Data Integrity The built-in [SHA-256](#) checksum algorithm ensures accuracy and consistency in your backups.

Remote Sync It is possible to efficiently synchronize data to remote sites. Only deltas containing new data are transferred.

Compression The ultra-fast [Zstandard](#) compression is able to compress several gigabytes of data per second.

Encryption Backups can be encrypted on the client-side, using AES-256 in Galois/Counter Mode ([GCM](#)) mode. This authenticated encryption ([AE](#)) mode provides very high performance on modern hardware.

Web interface Manage the Proxmox Backup Server with the integrated, web-based user interface.

Open Source No secrets. Proxmox Backup Server is free and open-source software. The source code is licensed under AGPL, v3.

Support Enterprise support will be available from [Proxmox](#) once the beta phase is over.

1.4 Reasons for Data Backup?

The main purpose of a backup is to protect against data loss. Data loss can be caused by both faulty hardware and human error.

A common mistake is to accidentally delete a file or folder which is still required. Virtualization can even amplify this problem, as deleting a whole virtual machine can be as easy as pressing a single button.

For administrators, backups can serve as a useful toolkit for temporarily storing data. For example, it is common practice to perform full backups before installing major software updates. If something goes wrong, you can easily restore the previous state.

Another reason for backups are legal requirements. Some data, especially business records, must be kept in a safe place for several years by law, so that they can be accessed if required.

In general, data loss is very costly as it can severely damage your business. Therefore, ensure that you perform regular backups and run restore tests.

1.5 Software Stack

Proxmox Backup Server consists of multiple components:

- A server-daemon providing, among other things, a RESTful API, super-fast asynchronous tasks, lightweight usage statistic collection, scheduling events, strict separation of privileged and unprivileged execution environments
- A JavaScript management web interface
- A management CLI tool for the server (*proxmox-backup-manager*)
- A client CLI tool (*proxmox-backup-client*) to access the server easily from any *Linux amd64* environment

Aside from the web interface, everything is written in the Rust programming language.

“The Rust programming language helps you write faster, more reliable software. High-level ergonomics and low-level control are often at odds in programming language design; Rust challenges that conflict. Through balancing powerful technical capacity and a great developer experience, Rust gives you the option to control low-level details (such as memory usage) without all the hassle traditionally associated with such control.”

1.6 Getting Help

1.6.1 Community Support Forum

We always encourage our users to discuss and share their knowledge using the [Proxmox Community Forum](#). The forum is moderated by the Proxmox support team. The large user base is spread out all over the world. Needless to say that such a large forum is a great place to get information.

1.6.2 Mailing Lists

Proxmox Backup Server is fully open-source and contributions are welcome! Here is the primary communication channel for developers:

Mailing list for developers [PBS Development List](#)

1.6.3 Bug Tracker

Proxmox runs a public bug tracker at <https://bugzilla.proxmox.com>. If an issue appears, file your report there. An issue can be a bug as well as a request for a new feature or enhancement. The bug tracker helps to keep track of the issue and will send a notification once it has been solved.

1.7 License

Copyright (C) 2019-2020 Proxmox Server Solutions GmbH

This software is written by Proxmox Server Solutions GmbH <support@proxmox.com>

Proxmox Backup Server is free and open source software: you can use it, redistribute it, and/or modify it under the terms of the GNU Affero General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Affero General Public License for more details.

You should have received a copy of the GNU Affero General Public License along with this program. If not, see [AGPL3](#).

1.8 History

Backup is, and always has been, a central aspect of IT administration. The need to recover from data loss is fundamental and only increases with virtualization.

For this reason, we've been shipping a backup tool with Proxmox VE, from the beginning. This tool is called `vzdump` and is able to make consistent snapshots of running LXC containers and KVM virtual machines.

However, `vzdump` only allows for full backups. While this is fine for small backups, it becomes a burden for users with large VMs. Both backup duration and storage usage are too high for this case,

especially for users who want to keep many backups of the same VMs. To solve these problems, we needed to offer deduplication and incremental backups.

Back in October 2018, development started. We investigated several technologies and frameworks and finally decided to use [Rust](#) as the implementation language, in order to provide high speed and memory efficiency. The 2018-edition of Rust seemed promising for our requirements.

In July 2020, we released the first beta version of Proxmox Backup Server, followed by the first stable version in November 2020. With support for incremental, fully deduplicated backups, Proxmox Backup significantly reduces network load and saves valuable storage space.

INSTALLATION

Proxmox Backup is split into a server and client part. The server part can either be installed with a graphical installer or on top of Debian from the provided package repository.

2.1 System Requirements

We recommend using high quality server hardware when running Proxmox Backup in production. To further decrease the impact of a failed host, you can set up periodic, efficient, incremental *data-store synchronization* from other Proxmox Backup Server instances.

2.1.1 Minimum Server Requirements, for Evaluation

These minimum requirements are for evaluation purposes only and should not be used in production.

- CPU: 64bit (x86-64 or AMD64), 2+ Cores
- Memory (RAM): 2 GB RAM
- Hard drive: more than 8GB of space.
- Network card (NIC)

2.1.2 Recommended Server System Requirements

- CPU: Modern AMD or Intel 64-bit based CPU, with at least 4 cores
- Memory: minimum 4 GiB for the OS, filesystem cache and Proxmox Backup Server daemons. Add at least another GiB per TiB storage space.
- OS storage:
 - 32 GiB, or more, free storage space
 - Use a hardware RAID with battery protected write cache (*BBU*) or a redundant ZFS setup (ZFS is not compatible with a hardware RAID controller).
- Backup storage:
 - Use only SSDs, for best results
 - If HDDs are used: Using a metadata cache is highly recommended, for example, add a ZFS *special device mirror*.
- Redundant Multi-GBit/s network interface cards (NICs)

2.1.3 Supported Web Browsers for Accessing the Web Interface

To access the server's web-based user interface, we recommend using one of the following browsers:

- Firefox, a release from the current year, or the latest Extended Support Release
- Chrome, a release from the current year
- Microsoft's currently supported version of Edge
- Safari, a release from the current year

2.2 Debian Package Repositories

All Debian based systems use [APT](#) as a package management tool. The lists of repositories are defined in `/etc/apt/sources.list` and the `.list` files found in the `/etc/apt/sources.d/` directory. Updates can be installed directly with the `apt` command line tool, or via the GUI.

`APT sources.list` files list one package repository per line, with the most preferred source listed first. Empty lines are ignored and a `#` character anywhere on a line marks the remainder of that line as a comment. The information available from the configured sources is acquired by `apt update`.

Listing 1: File: `/etc/apt/sources.list`

```
deb http://ftp.debian.org/debian buster main contrib
deb http://ftp.debian.org/debian buster-updates main contrib

# security updates
deb http://security.debian.org/debian-security buster/updates main contrib
```

In addition, you need a package repository from Proxmox to get Proxmox Backup updates.

During the Proxmox Backup beta phase, only one repository (pbtest) will be available. Once released, an Enterprise repository for production use and a no-subscription repository will be provided.

2.2.1 SecureApt

The *Release* files in the repositories are signed with GnuPG. APT is using these signatures to verify that all packages are from a trusted source.

If you install Proxmox Backup Server from an official ISO image, the verification key is already installed.

If you install Proxmox Backup Server on top of Debian, download and install the key with the following commands:

```
# wget http://download.proxmox.com/debian/proxmox-ve-release-6.x.gpg -O /etc/apt/trusted.gpg.d/proxmox-ve-release-6.x.gpg
```

Verify the SHA512 checksum afterwards with:

```
# sha512sum /etc/apt/trusted.gpg.d/proxmox-ve-release-6.x.gpg
```

The output should be:

```
acca6f416917e8e11490a08a1e2842d500b3a5d9f322c6319db0927b2901c3eae23cfb5cd5df6facf2b57399d3cfa52ad7769ebdd75d9
→/etc/apt/trusted.gpg.d/proxmox-ve-release-6.x.gpg
```

and the md5sum:

```
# md5sum /etc/apt/trusted.gpg.d/proxmox-ve-release-6.x.gpg
```

Here, the output should be:

```
f3f6c5a3a67baf38ad178e5ff1ee270c /etc/apt/trusted.gpg.d/proxmox-ve-release-6.x.gpg
```

2.2.2 Proxmox Backup Beta Repository

During the public beta, there is a repository called `pbstest`. This one contains the latest packages and is heavily used by developers to test new features.

You can access this repository by adding the following line to `/etc/apt/sources.list`:

Listing 2: `sources.list` entry for `pbstest`

```
deb http://download.proxmox.com/debian/pbs buster pbstest
```

If you installed Proxmox Backup Server from the official beta ISO, you should have this repository already configured in `/etc/apt/sources.list.d/pbstest-beta.list`

2.3 Server installation

The backup server stores the actual backed up data and provides a web based GUI for various management tasks such as disk management.

Note: You always need a backup server. It is not possible to use [Proxmox Backup](#) without the server part.

The disk image (ISO file) provided by Proxmox includes a complete Debian system ("buster" for version 1.x) as well as all necessary packages for the [Proxmox Backup](#) server.

The installer will guide you through the setup process and allow you to partition the local disk(s), apply basic system configurations (e.g. timezone, language, network), and install all required packages. The provided ISO will get you started in just a few minutes, and is the recommended method for new and existing users.

Alternatively, [Proxmox Backup](#) server can be installed on top of an existing Debian system.

2.3.1 Install Proxmox Backup with the Installer

Download the ISO from <https://www.proxmox.com/downloads>. It includes the following:

- The [Proxmox Backup](#) server installer, which partitions the local disk(s) with ext4, ext3, xfs or ZFS, and installs the operating system
- Complete operating system (Debian Linux, 64-bit)
- Our Linux kernel with ZFS support
- Complete tool-set to administer backups and all necessary resources
- Web based GUI management interface

Note: During the installation process, the complete server is used by default and all existing data is removed.

2.3.2 Install Proxmox Backup server on Debian

Proxmox ships as a set of Debian packages which can be installed on top of a standard Debian installation. After configuring the `sysadmin_package_repositories`, you need to run:

```
# apt-get update
# apt-get install proxmox-backup-server
```

The commands above keep the current (Debian) kernel and install a minimal set of required packages.

If you want to install the same set of packages as the installer does, please use the following:

```
# apt-get update
# apt-get install proxmox-backup
```

This will install all required packages, the Proxmox kernel with [ZFS](#) support, and a set of common and useful packages.

Caution: Installing [Proxmox Backup](#) on top of an existing [Debian](#) installation looks easy, but it assumes that the base system and local storage have been set up correctly. In general this is not trivial, especially when [LVM](#) or [ZFS](#) is used. The network configuration is completely up to you as well.

Note:

You can access the web interface of the Proxmox Backup Server with your web browser, using HTTPS on port 8007. For example at

`https://<ip-or-dns-name>:8007`

2.3.3 Install Proxmox Backup server on Proxmox VE

After configuring the `sysadmin_package_repositories`, you need to run:

```
# apt-get update
# apt-get install proxmox-backup-server
```

Caution: Installing the backup server directly on the hypervisor is not recommended. It is safer to use a separate physical server to store backups. Should the hypervisor server fail, you can still access the backups.

Note: You can access the web interface of the Proxmox Backup Server with your web browser, using HTTPS on port 8007. For example at `https://<ip-or-dns-name>:8007`

2.4 Client installation

2.4.1 Install Proxmox Backup client on Debian

Proxmox ships as a set of Debian packages to be installed on top of a standard Debian installation. After configuring the `sysadmin_package_repositories`, you need to run:


```
# apt-get update
# apt-get install proxmox-backup-client
```

2.4.2 Installing from source

2.4.3 Installing statically linked binary

TERMINOLOGY

3.1 Backup Content

When doing deduplication, there are different strategies to get optimal results in terms of performance and/or deduplication rates. Depending on the type of data, it can be split into *fixed* or *variable* sized chunks.

Fixed sized chunking requires minimal CPU power, and is used to backup virtual machine images.

Variable sized chunking needs more CPU power, but is essential to get good deduplication rates for file archives.

The Proxmox Backup Server supports both strategies.

3.1.1 Image Archives: <name>.img

This is used for virtual machine images and other large binary data. Content is split into fixed-sized chunks.

3.1.2 File Archives: <name>.pxar

A file archive stores a full directory tree. Content is stored using the *Proxmox File Archive Format (.pxar)*, split into variable-sized chunks. The format is optimized to achieve good deduplication rates.

3.1.3 Binary Data (BLOBs)

This type is used to store smaller (< 16MB) binary data such as configuration files. Larger files should be stored as image archive.

Caution: Please do not store all files as BLOBs. Instead, use the file archive to store whole directory trees.

3.1.4 Catalog File: catalog.pcat1

The catalog file is an index for file archives. It contains the list of files and is used to speed up search operations.

3.1.5 The Manifest: `index.json`

The manifest contains the list of all backup files, their sizes and checksums. It is used to verify the consistency of a backup.

3.2 Backup Type

The backup server groups backups by *type*, where *type* is one of:

vm This type is used for *virtual machines*. Typically consists of the virtual machine's configuration file and an image archive for each disk.

ct This type is used for *containers*. Consists of the container's configuration and a single file archive for the filesystem content.

host This type is used for backups created from within the backed up machine. Typically this would be a physical host but could also be a virtual machine or container. Such backups may contain file and image archives, there are no restrictions in this regard.

3.3 Backup ID

A unique ID. Usually the virtual machine or container ID. `host` type backups normally use the host-name.

3.4 Backup Time

The time when the backup was made.

3.5 Backup Group

The tuple `<type>/<ID>` is called a backup group. Such a group may contain one or more backup snapshots.

3.6 Backup Snapshot

The triplet `<type>/<ID>/<time>` is called a backup snapshot. It uniquely identifies a specific backup within a datastore.

Listing 1: Backup Snapshot Examples

```
vm/104/2019-10-09T08:01:06Z  
host/elsa/2019-11-08T09:48:14Z
```

As you can see, the time format is [RFC3399](#) with Coordinated Universal Time (UTC, identified by the trailing Z).

GRAPHICAL USER INTERFACE


Proxmox Backup Server offers an integrated, web-based interface to manage the server. This means that you can carry out all administration tasks through your web browser, and that you don't have to worry about installing extra management tools. The web interface also provides a built in console, so if you prefer the command line or need some extra control, you have this option.

The web interface can be accessed via <https://youripaddress:8007>. The default login is *root*, and the password is the one specified during the installation process.

4.1 Features

- Simple management interface for Proxmox Backup Server
- Monitoring of tasks, logs and resource usage
- Management of users, permissions, datastores, etc.
- Secure HTML5 console
- Support for multiple authentication sources
- Support for multiple languages
- Based on ExtJS 6.x JavaScript framework

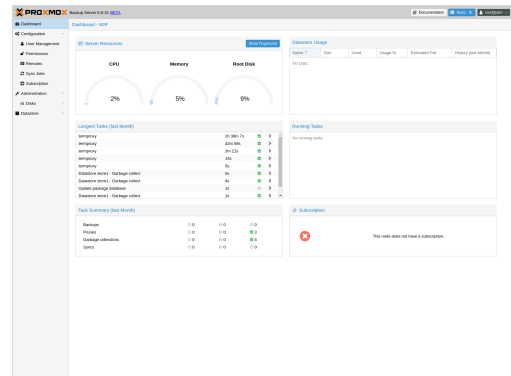
4.2 Login



When you connect to the web interface, you will first see the login window. Proxmox Backup Server supports various languages and authentication back ends (*Realms*), both of which can be selected here.

Note: For convenience, you can save the username on the client side, by selecting the “Save User name” checkbox at the bottom of the window.

4.3 GUI Overview



The Proxmox Backup Server web interface consists of 3 main sections:

- **Header:** At the top. This shows version information, and contains buttons to view documentation, monitor running tasks, and logout.
- **Sidebar:** On the left. This contains the configuration options for the server.
- **Configuration Panel:** In the center. This contains the control interface for the configuration options in the *Sidebar*.

4.4 Sidebar

In the sidebar, on the left side of the page, you can see various items relating to specific management activities.

4.4.1 Dashboard

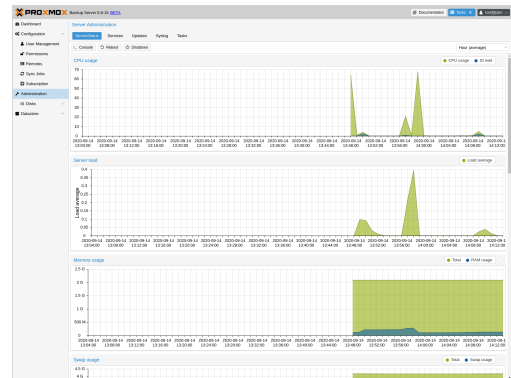
The Dashboard shows a summary of activity and resource usage on the server. Specifically, this displays hardware usage, a summary of previous and currently running tasks, and subscription information.

4.4.2 Configuration

The Configuration section contains some system configuration options, such as time and network configuration. It also contains the following subsections:

- **User Management:** Add users and manage accounts
- **Permissions:** Manage permissions for various users
- **Remotes:** Add, edit and remove remotes (see [Remote](#))
- **Sync Jobs:** Manage and run sync jobs to remotes
- **Subscription:** Upload a subscription key and view subscription status

4.4.3 Administration



The Administration section contains a top panel, with further administration tasks and information. These are:

- **ServerStatus:** Provides access to the console, power options, and various resource usage statistics
- **Services:** Manage and monitor system services
- **Updates:** An interface for upgrading packages
- **Syslog:** View log messages from the server
- **Tasks:** Task history with multiple filter options

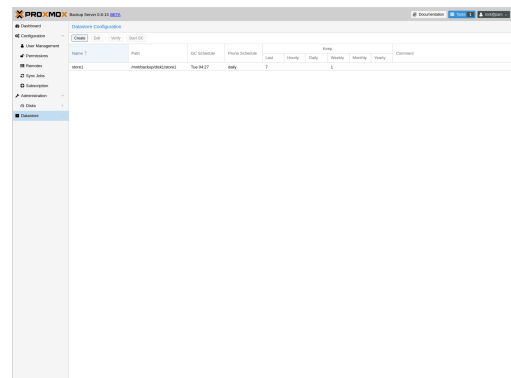
The screenshot displays the 'Disks' section of the Proxmox Backup web interface. It shows a table of available disks with columns: Name, Type, Format, Size, LUN, Serial, and Status. The table lists four disks, all of which are 'Free'.

Name	Type	Format	Size	LUN	Serial	Status
disk1	Hard Disk	ext4	10.00 GB	0	QEMU_HARDDISK	Free
disk2	Hard Disk	ext4	10.00 GB	1	QEMU_HARDDISK	Free
disk3	Hard Disk	ext4	10.00 GB	2	QEMU_HARDDISK	Free
disk4	Hard Disk	ext4	10.00 GB	3	QEMU_HARDDISK	Free

The administration menu item also contains a disk management subsection:

- **Disks:** View information on available disks
 - **Directory:** Create and view information on *ext4* and *xfs* disks
 - **ZFS:** Create and view information on *ZFS* disks

4.4.4 Datastore



The Datastore section provides an interface for creating and managing datastores. It contains a subsection for each datastore on the system, in which you can use the top panel to view:

- **Content:** Information on the datastore's backup groups and their respective contents
- **Statistics:** Usage statistics for the datastore
- **Permissions:** View and manage permissions for the datastore

5.1 Disk Management

Device	Type	Usage	Size	GPT	Model	Serial	S.M.A.R.T.	Wearout
/dev/sda	Hard Disk	lvm	32.00 GiB	Yes	QEMU_HARDDISK	drive-scsi0	passed	N/A
/dev/sdb	Hard Disk	unused	64.00 GiB	No	QEMU_HARDDISK	drive-scsi3	passed	N/A
/dev/sdc	Hard Disk	unused	64.00 GiB	No	QEMU_HARDDISK	drive-scsi2	passed	N/A
/dev/sdd	Hard Disk	unused	64.00 GiB	No	QEMU_HARDDISK	drive-scsi1	passed	N/A

Proxmox Backup Server comes with a set of disk utilities, which are accessed using the `disk` subcommand. This subcommand allows you to initialize disks, create various filesystems, and get information about the disks.

To view the disks connected to the system, navigate to **Administration -> Disks** in the web interface or use the `list` subcommand of `disk`:

```
# proxmox-backup-manager disk list
```

name	used	gpt	disk-type	size	model	wearout	status
sda	lvm	1	hdd	34359738368	QEMU_HARDDISK	-	passed
sdb	unused	1	hdd	68719476736	QEMU_HARDDISK	-	passed

(continues on next page)

(continued from previous page)

sdc	unused	1	hdd	68719476736	QEMU_HARDDISK	-	passed
-----	--------	---	-----	-------------	---------------	---	--------

To initialize a disk with a new GPT, use the `initialize` subcommand:

```
# proxmox-backup-manager disk initialize sdX
```

You can create an ext4 or xfs filesystem on a disk using `fs create`, or by navigating to **Administration -> Disks -> Directory** in the web interface and creating one from there. The following command creates an ext4 filesystem and passes the `--add-datastore` parameter, in order to automatically create a datastore on the disk (in this case sdd). This will create a datastore at the location `/mnt/datastore/store1`:

```
# proxmox-backup-manager disk fs create store1 --disk sdd --filesystem ext4 --add-datastore_
→ true
```

You can also create a `zpool` with various raid levels from **Administration -> Disks -> Zpool** in the web interface, or by using `zpool create`. The command below creates a mirrored `zpool` using two disks (sdb & sdc) and mounts it on the root directory (default):

```
# proxmox-backup-manager disk zpool create zpool1 --devices sdb,sdc --raidlevel mirror
```

Note: You can also pass the `--add-datastore` parameter here, to automatically create a datastore from the disk.

You can use `disk fs list` and `disk zpool list` to keep track of your filesystems and zpools respectively.

Proxmox Backup Server uses the package `smartmontools`. This is a set of tools used to monitor

and control the S.M.A.R.T. system for local hard disks. If a disk supports S.M.A.R.T. capability, and you have this enabled, you can display S.M.A.R.T. attributes from the web interface or by using the command:

```
# proxmox-backup-manager disk smart-attributes sdX
```

Note: This functionality may also be accessed directly through the use of the `smartctl` command, which comes as part of the `smartmontools` package (see `man smartctl` for more details).

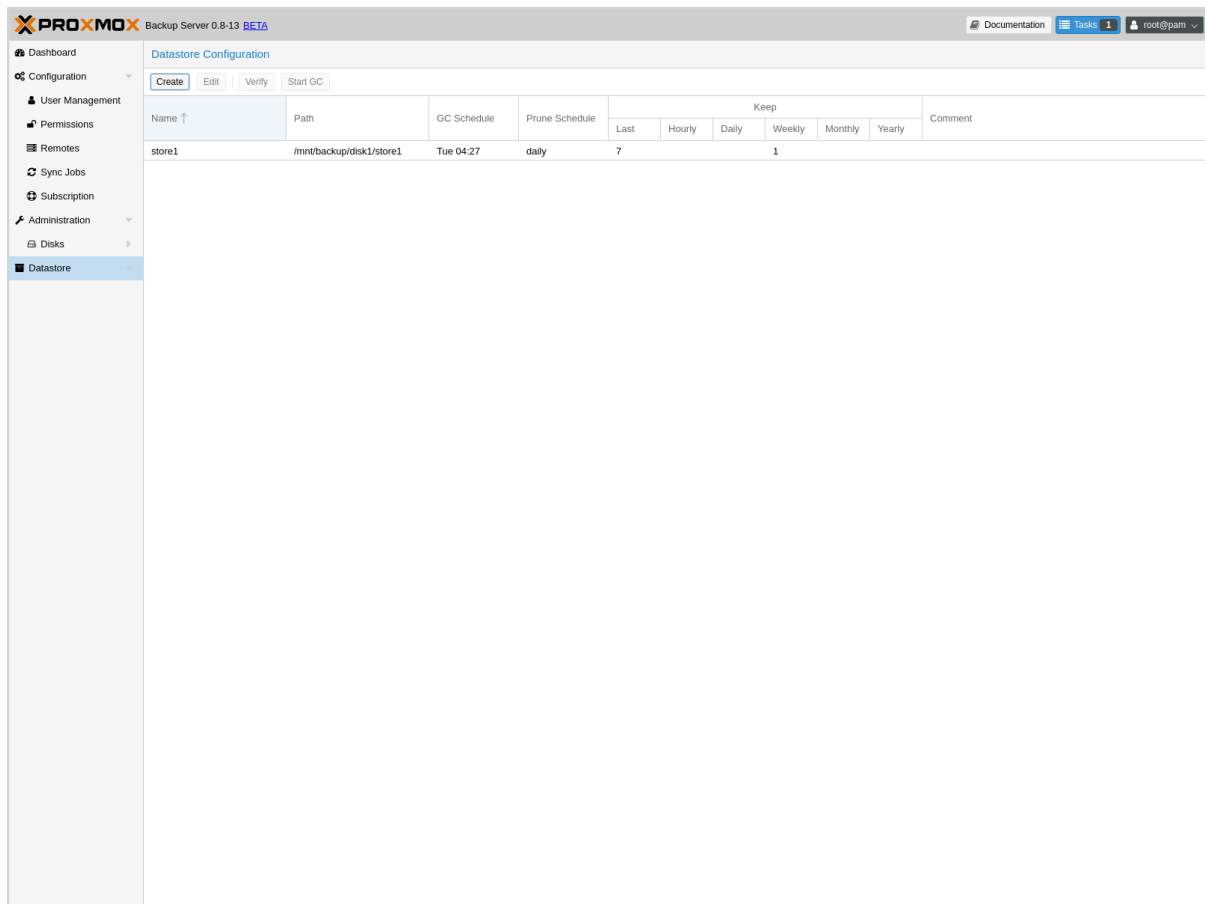
5.2 DataStore

A datastore refers to a location at which backups are stored. The current implementation uses a directory inside a standard Unix file system (ext4, xfs or zfs) to store the backup data.

Datastores are identified by a simple *ID*. You can configure this when setting up the datastore. The configuration information for datastores is stored in the file `/etc/proxmox-backup/datastore.cfg`.

Note: The *File Layout* requires the file system to support at least 65538 subdirectories per directory. That number comes from the 2^{16} pre-created chunk namespace directories, and the `.` and `..` default directory entries. This requirement excludes certain filesystems and filesystem configuration from being supported for a datastore. For example, ext3 as a whole or ext4 with the `dir_nlink` feature manually disabled.

5.2.1 Datastore Configuration



You can configure multiple datastores. Minimum one datastore needs to be configured. The datastore is identified by a simple *name* and points to a directory on the filesystem. Each datastore also has associated retention settings of how many backup snapshots for each interval of hourly, daily, weekly, monthly, yearly as well as a time-independent number of backups to keep in that store. [Pruning and Removing Backups](#) and [garbage collection](#) can also be configured to run periodically based on a configured schedule (see [Calendar Events](#)) per datastore.

Creating a Datastore

Add: Datastore

General
Prune Options

Name:
GC Schedule:

Backing Path:
Prune Schedule:

Comment:

Add

You can create a new datastore from the web GUI, by navigating to **Datastore** in the menu tree and clicking **Create**. Here:

- *Name* refers to the name of the datastore
- *Backing Path* is the path to the directory upon which you want to create the datastore
- *GC Schedule* refers to the time and intervals at which garbage collection runs

- *Prune Schedule* refers to the frequency at which pruning takes place
- *Prune Options* set the amount of backups which you would like to keep (see [Pruning and Removing Backups](#)).

Alternatively you can create a new datastore from the command line. The following command creates a new datastore called `store1` on `/backup/disk1/store1`

```
# proxmox-backup-manager datastore create store1 /backup/disk1/store1
```

Managing Datastores

To list existing datastores from the command line run:

```
# proxmox-backup-manager datastore list
```

name	path	comment
store1	/backup/disk1/store1	This is my default storage.

You can change the garbage collection and prune settings of a datastore, by editing the datastore from the GUI or by using the update subcommand. For example, the below command changes the garbage collection schedule using the update subcommand and prints the properties of the datastore with the show subcommand:

```
# proxmox-backup-manager datastore update store1 --gc-schedule 'Tue 04:27'
# proxmox-backup-manager datastore show store1
```

Name	Value
name	store1
path	/backup/disk1/store1
comment	This is my default storage.
gc-schedule	Tue 04:27
keep-last	7
prune-schedule	daily

Finally, it is possible to remove the datastore configuration:

```
# proxmox-backup-manager datastore remove store1
```

Note: The above command removes only the datastore configuration. It does not delete any data from the underlying directory.

File Layout

After creating a datastore, the following default layout will appear:

```
# ls -arilh /backup/disk1/store1
276493 -rw-r--r-- 1 backup backup      0 Jul  8 12:35 .lock
276490 drwxr-x--- 1 backup backup 1064960 Jul  8 12:35 .chunks
```

`.lock` is an empty file used for process locking.

The `.chunks` directory contains folders, starting from `0000` and taking hexadecimal values until `ffff`. These directories will store the chunked data after a backup operation has been executed.

```
# ls -arilh /backup/disk1/store1/.chunks
545824 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 ffff
545823 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 fffe
415621 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 fffd
415620 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 fffc
353187 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 fffb
344995 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 fffa
144079 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 fff9
144078 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 fff8
144077 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 fff7
...
403180 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 000c
403179 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 000b
403177 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 000a
402530 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 0009
402513 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 0008
402509 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 0007
276509 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 0006
276508 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 0005
276507 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 0004
276501 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 0003
276499 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 0002
276498 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 0001
276494 drwxr-x--- 2 backup backup 4.0K Jul  8 12:35 0000
276489 drwxr-xr-x 3 backup backup 4.0K Jul  8 12:35 ..
276490 drwxr-x--- 1 backup backup 1.1M Jul  8 12:35 .
```

NETWORK MANAGEMENT

Proxmox Backup Server provides both a web interface and a command line tool for network configuration. You can find the configuration options in the web interface under the **Network Interfaces** section of the **Configuration** menu tree item. The command line tool is accessed via the network subcommand. These interfaces allow you to carry out some basic network management tasks, such as adding, configuring, and removing network interfaces.

Note: Any changes made to the network configuration are not applied, until you click on **Apply Configuration** or enter the `network reload` command. This allows you to make many changes at once. It also allows you to ensure that your changes are correct before applying them, as making a mistake here can render the server inaccessible over the network.

To get a list of available interfaces, use the following command:

```
# proxmox-backup-manager network list
```

name	type	autostart	method	address	gateway	ports/slaves
bond0	bond	1	static	x.x.x.x/x	x.x.x.x	ens18 ens19
ens18	eth	1	manual			
ens19	eth	1	manual			

To add a new network interface, use the `create` subcommand with the relevant parameters. For example, you may want to set up a bond, for the purpose of network redundancy. The following command shows a template for creating the bond shown in the list above:

```
# proxmox-backup-manager network create bond0 --type bond --bond_mode active-backup --slaves ens18,ens19 --autostart true --cidr x.x.x.x/x --gateway x.x.x.x
```

You can make changes to the configuration of a network interface with the `update` subcommand:

```
# proxmox-backup-manager network update bond0 --cidr y.y.y.y/y
```

You can also remove a network interface:

```
# proxmox-backup-manager network remove bond0
```

The pending changes for the network configuration file will appear at the bottom of the web interface. You can also view these changes, by using the command:

```
# proxmox-backup-manager network changes
```

If you would like to cancel all changes at this point, you can either click on the **Revert** button or use the following command:

```
# proxmox-backup-manager network revert
```

If you are happy with the changes and would like to write them into the configuration file, select **Apply Configuration**. The corresponding command is:

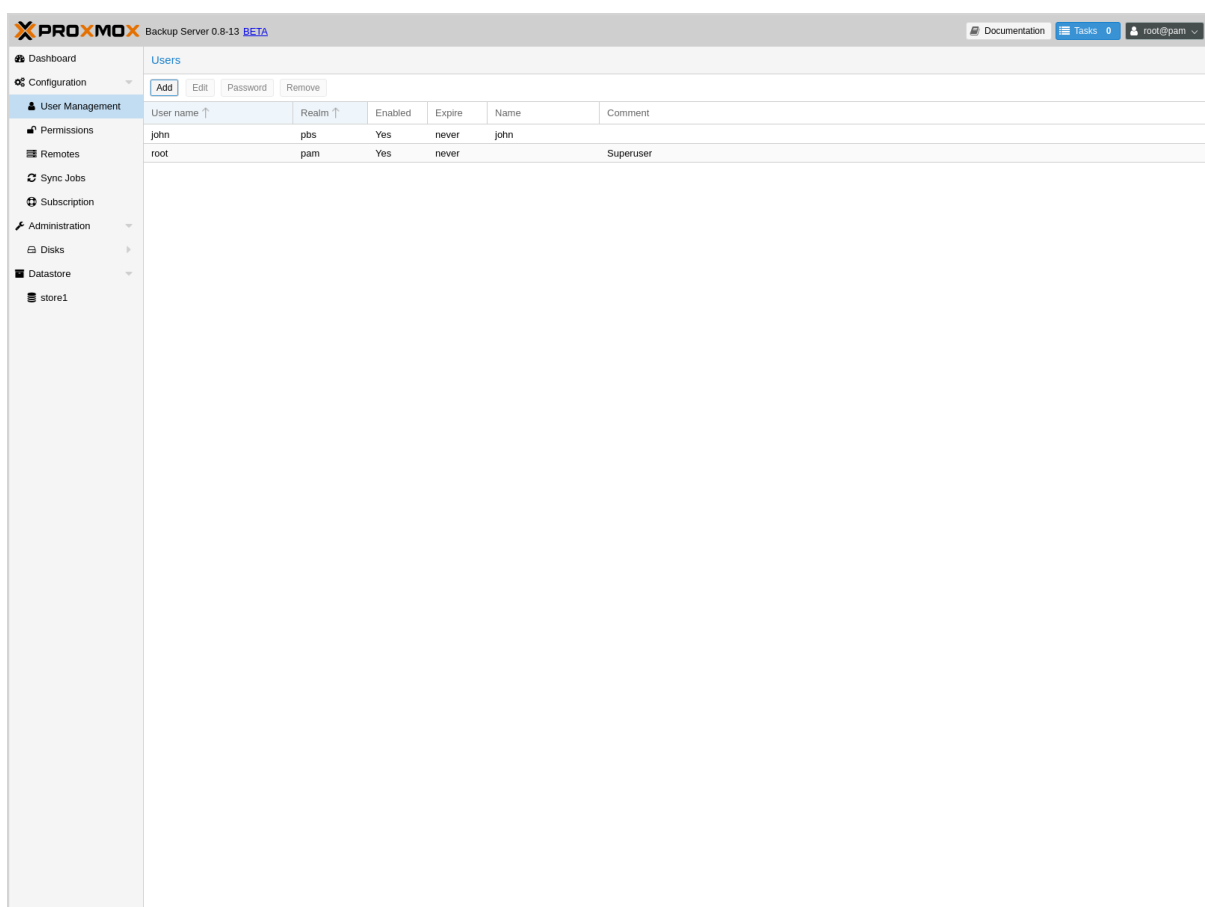
```
# proxmox-backup-manager network reload
```

Note: This command and corresponding GUI button rely on the `ifreload` command, from the package `ifupdown2`. This package is included within the Proxmox Backup Server installation, however, you may have to install it yourself, if you have installed Proxmox Backup Server on top of Debian or Proxmox VE.

You can also configure DNS settings, from the **DNS** section of **Configuration** or by using the `dns` subcommand of `proxmox-backup-manager`.

USER MANAGEMENT

7.1 User Configuration



Proxmox Backup Server supports several authentication realms, and you need to choose the realm when you add a new user. Possible realms are:

- pam** Linux PAM standard authentication. Use this if you want to authenticate as Linux system user (Users need to exist on the system).
- pbs** Proxmox Backup Server realm. This type stores hashed passwords in `/etc/proxmox-backup/shadow.json`.

After installation, there is a single user `root@pam`, which corresponds to the Unix superuser. User configuration information is stored in the file `/etc/proxmox-backup/user.cfg`. You can use the `proxmox-backup-manager` command line tool to list or manipulate users:

```
# proxmox-backup-manager user list
```

userid	enable	expire	firstname	lastname	email	comment
root@pam	1					Superuser

The superuser has full administration rights on everything, so you normally want to add other users with less privileges. You can create a new user with the `user create` subcommand or through the web interface, under **Configuration -> User Management**. The `create` subcommand lets you specify many options like `--email` or `--password`. You can update or change any user properties using the `update` subcommand later (**Edit** in the GUI):

```
# proxmox-backup-manager user create john@pbs --email john@example.com
# proxmox-backup-manager user update john@pbs --firstname John --lastname Smith
# proxmox-backup-manager user update john@pbs --comment "An example user."
```

The resulting user list looks like this:

```
# proxmox-backup-manager user list
```

userid	enable	expire	firstname	lastname	email	comment
john@pbs	1		John	Smith	john@example.com	An example user.
root@pam	1					Superuser

Newly created users do not have any permissions. Please read the next section to learn how to set access permissions.

If you want to disable a user account, you can do that by setting `--enable` to `0`

```
# proxmox-backup-manager user update john@pbs --enable 0
```

Or completely remove the user with:

```
# proxmox-backup-manager user remove john@pbs
```

7.2 Access Control

By default new users do not have any permission. Instead you need to specify what is allowed and what is not. You can do this by assigning roles to users on specific objects like datastores or remotes. The following roles exist:

NoAccess Disable Access - nothing is allowed.

Admin Can do anything.

Audit Can view things, but is not allowed to change settings.

DatastoreAdmin Can do anything on datastores.

DatastoreAudit Can view datastore settings and list content. But is not allowed to read the actual data.

DatastoreReader Can inspect datastore content and can do restores.

DatastoreBackup Can backup and restore owned backups.

DatastorePowerUser Can backup, restore, and prune owned backups.

RemoteAdmin Can do anything on remotes.

RemoteAudit Can view remote settings.

RemoteSyncOperator Is allowed to read data from a remote.

Access permission information is stored in `/etc/proxmox-backup/acl.cfg`. The file contains 5 fields, separated using a colon (":") as a delimiter. A typical entry takes the form:

```
acl:1:/datastore:john@pbs:DatastoreBackup
```

The data represented in each field is as follows:

1. acl identifier
2. A 1 or 0, representing whether propagation is enabled or disabled, respectively
3. The object on which the permission is set. This can be a specific object (single datastore, remote, etc.) or a top level object, which with propagation enabled, represents all children of the object also.
4. The user for which the permission is set
5. The role being set

You can manage datastore permissions from **Configuration -> Permissions** in the web interface. Likewise, you can use the `acl` subcommand to manage and monitor user permissions from the command line. For example, the command below will add the user `john@pbs` as a **DatastoreAdmin** for the datastore `store1`, located at `/backup/disk1/store1`:

```
# proxmox-backup-manager acl update /datastore/store1 DatastoreAdmin --userid john@pbs
```

You can monitor the roles of each user using the following command:

```
# proxmox-backup-manager acl list
```

ugid	path	propagate	roleid
john@pbs	/datastore/disk1	1	DatastoreAdmin

A single user can be assigned multiple permission sets for different datastores.

Note: Naming convention is important here. For datastores on the host, you must use the convention `/datastore/{storename}`. For example, to set permissions for a datastore mounted at `/mnt/backup/disk4/store2`, you would use `/datastore/store2` for the path. For remote stores, use the convention `/remote/{remote}/{storename}`, where `{remote}` signifies the

name of the remote (see *Remote* below) and {storename} is the name of the datastore on the remote.

MANAGING REMOTES

8.1 Remote

A remote refers to a separate Proxmox Backup Server installation and a user on that installation, from which you can *sync* datastores to a local datastore with a *Sync Job*. You can configure remotes in the web interface, under **Configuration -> Remotes**. Alternatively, you can use the `remote` subcommand. The configuration information for remotes is stored in the file `/etc/proxmox-backup/remote.cfg`.

To add a remote, you need its hostname or IP, a userid and password on the remote, and its certificate fingerprint. To get the fingerprint, use the `proxmox-backup-manager cert info` command on the remote, or navigate to **Dashboard** in the remote's web interface and select **Show Fingerprint**.

```
# proxmox-backup-manager cert info |grep Fingerprint
Fingerprint (sha256): 64:d3:ff:3a:50:38:53:5a:9b:f7:50:...:ab:fe
```

Using the information specified above, you can add a remote from the **Remotes** configuration panel, or by using the command:

```
# proxmox-backup-manager remote create pbs2 --host pbs2.mydomain.example --userid sync@pam --
password 'SECRET' --fingerprint 64:d3:ff:3a:50:38:53:5a:9b:f7:50:...:ab:fe
```

Use the `list`, `show`, `update`, `remove` subcommands of `proxmox-backup-manager remote` to manage your remotes:

```
# proxmox-backup-manager remote update pbs2 --host pbs2.example
# proxmox-backup-manager remote list
```

name	host	userid	fingerprint	comment
pbs2	pbs2.example	sync@pam	64:d3:ff:3a:50:38:53:5a:9b:f7:50:...:ab:fe	

```
# proxmox-backup-manager remote remove pbs2
```

8.2 Sync Jobs

The screenshot shows a web form titled "Add: SyncJob". It has the following fields:

- Sync Job ID:
- Source Remote:
- Source Datastore:
- Local Datastore:
- Schedule:
- Comment:
- Remove vanished: ☐
- An "Add" button is located at the bottom right.

Sync jobs are configured to pull the contents of a datastore on a **Remote** to a local datastore. You can manage sync jobs under **Configuration -> Sync Jobs** in the web interface, or using the `proxmox-backup-manager sync-job` command. The configuration information for sync jobs is stored at `/etc/proxmox-backup/sync.cfg`. To create a new sync job, click the add button in the GUI, or use the `create` subcommand. After creating a sync job, you can either start it manually on the GUI or provide it with a schedule (see [Calendar Events](#)) to run regularly.

```
# proxmox-backup-manager sync-job create pbs2-local --remote pbs2 --remote-store local --
→store local --schedule 'Wed 02:30'
# proxmox-backup-manager sync-job update pbs2-local --comment 'offsite'
# proxmox-backup-manager sync-job list
```

id	store	remote	remote-store	schedule	comment
pbs2-local	local	pbs2	local	Wed 02:30	offsite

```
# proxmox-backup-manager sync-job remove pbs2-local
```

MAINTENANCE TASKS

9.1 Garbage Collection

You can monitor and run *garbage collection* on the Proxmox Backup Server using the `garbage-collection` subcommand of `proxmox-backup-manager`. You can use the `start` subcommand to manually start garbage collection on an entire datastore and the `status` subcommand to see attributes relating to the *garbage collection*.

BACKUP CLIENT USAGE

The command line client is called **proxmox-backup-client**.

10.1 Repository Locations

The client uses the following notation to specify a datastore repository on the backup server.

`[[username@]server[:port]:]datastore`

The default value for username is `root@pam`. If no server is specified, the default is the local host (`localhost`).

You can specify a port if your backup server is only reachable on a different port (e.g. with NAT and port forwarding).

Note that if the server is an IPv6 address, you have to write it with square brackets (e.g. `[fe80::01]`).

You can pass the repository with the `--repository` command line option, or by setting the `PBS_REPOSITORY` environment variable.

Here some examples of valid repositories and the real values

Example	User	Host:Port	Datastore
mydatastore	root@pam	localhost:8007	mydatastore
myhostname:mydatastore	root@pam	myhostname:8007	mydatastore
user@pbs@myhostname:mydatastore	user@pbs	myhostname:8007	mydatastore
192.168.55.55:1234:mydatastore	root@pam	192.168.55.55:1234	mydatastore
[ff80::51]:mydatastore	root@pam	[ff80::51]:8007	mydatastore
[ff80::51]:1234:mydatastore	root@pam	[ff80::51]:1234	mydatastore

10.2 Environment Variables

PBS_REPOSITORY The default backup repository.

PBS_PASSWORD When set, this value is used for the password required for the backup server.

PBS_ENCRYPTION_PASSWORD When set, this value is used to access the secret encryption key (if protected by password).

PBS_FINGERPRINT When set, this value is used to verify the server certificate (only used if the system CA certificates cannot validate the certificate).

10.3 Output Format

Most commands support the `--output-format` parameter. It accepts the following values:

text Text format (default). Structured data is rendered as a table.

json JSON (single line).

json-pretty JSON (multiple lines, nicely formatted).

Please use the following environment variables to modify output behavior:

PROXMOX_OUTPUT_FORMAT Defines the default output format.

PROXMOX_OUTPUT_NO_BORDER If set (to any value), do not render table borders.

PROXMOX_OUTPUT_NO_HEADER If set (to any value), do not render table headers.

Note: The text format is designed to be human readable, and not meant to be parsed by automation tools. Please use the `json` format if you need to process the output.

10.4 Creating Backups

This section explains how to create a backup from within the machine. This can be a physical host, a virtual machine, or a container. Such backups may contain file and image archives. There are no restrictions in this case.

Note: If you want to backup virtual machines or containers on Proxmox VE, see [Proxmox VE Integration](#).

For the following example you need to have a backup server set up, working credentials and need to know the repository name. In the following examples we use `backup-server:store1`.

```
# proxmox-backup-client backup root.pxar:/ --repository backup-server:store1
Starting backup: host/elsa/2019-12-03T09:35:01Z
Client name: elsa
skip mount point: "/boot/efi"
skip mount point: "/dev"
skip mount point: "/run"
skip mount point: "/sys"
Uploaded 12129 chunks in 87 seconds (564 MB/s).
End Time: 2019-12-03T10:36:29+01:00
```

This will prompt you for a password and then uploads a file archive named `root.pxar` containing all the files in the `/` directory.

Caution: Please note that the `proxmox-backup-client` does not automatically include mount points. Instead, you will see a short `skip mount point` notice for each of them. The idea is to create a separate file archive for each mounted disk. You can explicitly include them using the `--include-dev` option (i.e. `--include-dev /boot/efi`). You can use this option multiple times for each mount point that should be included.

The `--repository` option can get quite long and is used by all commands. You can avoid having to enter this value by setting the environment variable `PBS_REPOSITORY`. Note that if you would like this to remain set over multiple sessions, you should instead add the below line to your `.bashrc` file.

```
# export PBS_REPOSITORY=backup-server:store1
```

After this you can execute all commands without specifying the `--repository` option.

One single backup is allowed to contain more than one archive. For example, if you want to backup two disks mounted at `/mnt/disk1` and `/mnt/disk2`:

```
# proxmox-backup-client backup disk1.pxar:/mnt/disk1 disk2.pxar:/mnt/disk2
```

This creates a backup of both disks.

The backup command takes a list of backup specifications, which include the archive name on the server, the type of the archive, and the archive source at the client. The format is:

```
<archive-name>.<type>:<source-path>
```

Common types are `.pxar` for file archives, and `.img` for block device images. To create a backup of a block device run the following command:

```
# proxmox-backup-client backup mydata.img:/dev/mylvm/mydata
```

10.4.1 Excluding files/folders from a backup

Sometimes it is desired to exclude certain files or folders from a backup archive. To tell the Proxmox Backup client when and how to ignore files and directories, place a text file called `.pxarexclude` in the filesystem hierarchy. Whenever the backup client encounters such a file in a directory, it interprets each line as glob match patterns for files and directories that are to be excluded from the backup.

The file must contain a single glob pattern per line. Empty lines are ignored. The same is true for lines starting with `#`, which indicates a comment. A `!` at the beginning of a line reverses the glob match pattern from an exclusion to an explicit inclusion. This makes it possible to exclude all entries in a directory except for a few single files/subdirectories. Lines ending in `/` match only on directories. The directory containing the `.pxarexclude` file is considered to be the root of the given patterns. It is only possible to match files in this directory and its subdirectories.

`\` is used to escape special glob characters. `?` matches any single character. `*` matches any character, including an empty string. `**` is used to match subdirectories. It can be used to, for example, exclude all files ending in `.tmp` within the directory or subdirectories with the following pattern `**/*.tmp`. `[...]` matches a single character from any of the provided characters within the brackets. `[!...]` does the complementary and matches any single character not contained within the brackets. It is also possible to specify ranges with two characters separated by `-`. For example, `[a-z]` matches any lowercase alphabetic character and `[0-9]` matches any one single digit.

The order of the glob match patterns defines whether a file is included or excluded, that is to say later entries override previous ones. This is also true for match patterns encountered deeper down the directory tree, which can override a previous exclusion. Be aware that excluded directories will **not** be read by the backup client. Thus, a `.pxarexclude` file in an excluded subdirectory will have no effect. `.pxarexclude` files are treated as regular files and will be included in the backup archive.

For example, consider the following directory structure:

```
# ls -aR folder
folder/:
.  ..  .pxarexclude  subfolder0  subfolder1

folder/subfolder0:
.  ..  file0  file1  file2  file3  .pxarexclude

folder/subfolder1:
.  ..  file0  file1  file2  file3
```

The different `.pxarexclude` files contain the following:

```
# cat folder/.pxarexclude
/subfolder0/file1
/subfolder1/*
!/subfolder1/file2
```

```
# cat folder/subfolder0/.pxarexclude
file3
```

This would exclude file1 and file3 in subfolder0 and all of subfolder1 except file2.

Restoring this backup will result in:

```
ls -aR restored
restored/:
.  ..  .pxarexclude  subfolder0  subfolder1

restored/subfolder0:
.  ..  file0  file2  .pxarexclude

restored/subfolder1:
.  ..  file2
```

10.5 Encryption

Proxmox Backup supports client-side encryption with AES-256 in [GCM](#) mode. To set this up, you first need to create an encryption key:

```
# proxmox-backup-client key create my-backup.key
Encryption Key Password: *****
```

The key is password protected by default. If you do not need this extra protection, you can also create it without a password:

```
# proxmox-backup-client key create /path/to/my-backup.key --kdf none
```

Having created this key, it is now possible to create an encrypted backup, by passing the `--keyfile` parameter, with the path to the key file.

```
# proxmox-backup-client backup etc.pxar:/etc --keyfile /path/to/my-backup.key
Password: *****
Encryption Key Password: *****
...
```

Note: If you do not specify the name of the backup key, the key will be created in the default location `~/.config/proxmox-backup/encryption-key.json`. `proxmox-backup-client` will also search this location by default, in case the `--keyfile` parameter is not specified.

You can avoid entering the passwords by setting the environment variables `PBS_PASSWORD` and `PBS_ENCRYPTION_PASSWORD`.

10.5.1 Using a master key to store and recover encryption keys

You can also use `proxmox-backup-client key` to create an RSA public/private key pair, which can be used to store an encrypted version of the symmetric backup encryption key alongside each backup and recover it later.

To set up a master key:

1. Create an encryption key for the backup:

```
# proxmox-backup-client key create
creating default key at: "~/.config/proxmox-backup/encryption-key.json"
Encryption Key Password: *****
...
```

The resulting file will be saved to `~/.config/proxmox-backup/encryption-key.json`.

2. Create an RSA public/private key pair:

```
# proxmox-backup-client key create-master-key
Master Key Password: *****
...
```

This will create two files in your current directory, `master-public.pem` and `master-private.pem`.

3. Import the newly created `master-public.pem` public certificate, so that `proxmox-backup-client` can find and use it upon backup.

```
# proxmox-backup-client key import-master-pubkey /path/to/master-public.pem
Imported public master key to "~/.config/proxmox-backup/master-public.pem"
```

4. With all these files in place, run a backup job:

```
# proxmox-backup-client backup etc.pxdar:/etc
```

The key will be stored in your backup, under the name `rsa-encrypted.key`.

Note: The `--keyfile` parameter can be excluded, if the encryption key is in the default path. If you specified another path upon creation, you must pass the `--keyfile` parameter.

5. To test that everything worked, you can restore the key from the backup:

```
# proxmox-backup-client restore /path/to/backup/ rsa-encrypted.key /path/to/target
```

Note: You should not need an encryption key to extract this file. However, if a key exists at the default location (`~/.config/proxmox-backup/encryption-key.json`) the program will prompt you for an encryption key password. Simply moving `encryption-key.json` out of this directory will fix this issue.

6. Then, use the previously generated master key to decrypt the file:

```
# openssl rsautl -decrypt -inkey master-private.pem -in rsa-encrypted.key -out /path/to/
target
Enter pass phrase for ./master-private.pem: *****
```

7. The target file will now contain the encryption key information in plain text. The success of this can be confirmed by passing the resulting json file, with the `--keyfile` parameter, when decrypting files from the backup.

Warning: Without their key, backed up files will be inaccessible. Thus, you should keep keys ordered and in a place that is separate from the contents being backed up. It can happen, for example, that you back up an entire system, using a key on that system. If the system then becomes inaccessible for any reason and needs to be restored, this will not be possible as the encryption key will be lost along with the broken system. In preparation for the worst case scenario, you should consider keeping a paper copy of this key locked away in a safe place.

10.6 Restoring Data

The regular creation of backups is a necessary step to avoiding data loss. More importantly, however, is the restoration. It is good practice to perform periodic recovery tests to ensure that you can access the data in case of problems.

First, you need to find the snapshot which you want to restore. The snapshot command provides a list of all the snapshots on the server:

```
# proxmox-backup-client snapshots
```

snapshot	size	files
host/elsa/2019-12-03T09:30:15Z	51788646825	root.pxar catalog.pcat1 index.json
host/elsa/2019-12-03T09:35:01Z	51790622048	root.pxar catalog.pcat1 index.json

```
...
```

You can inspect the catalog to find specific files.

```
# proxmox-backup-client catalog dump host/elsa/2019-12-03T09:35:01Z
```

```
...
d "./root.pxar.didx/etc/cifs-utils"
l "./root.pxar.didx/etc/cifs-utils/idmap-plugin"
d "./root.pxar.didx/etc/console-setup"
...
```

The restore command lets you restore a single archive from the backup.

```
# proxmox-backup-client restore host/elsa/2019-12-03T09:35:01Z root.pxar /target/path/
```

To get the contents of any archive, you can restore the `index.json` file in the repository to the target path `'.`. This will dump the contents to the standard output.

```
# proxmox-backup-client restore host/elsa/2019-12-03T09:35:01Z index.json -
```

10.6.1 Interactive Restores

If you only want to restore a few individual files, it is often easier to use the interactive recovery shell.

```
# proxmox-backup-client catalog shell host/elsa/2019-12-03T09:35:01Z root.pxar
Starting interactive shell
pxar:/ > ls
bin      boot      dev        etc        home       lib        lib32
...
```

The interactive recovery shell is a minimal command line interface that utilizes the metadata stored in the catalog to quickly list, navigate and search files in a file archive. To restore files, you can select them individually or match them with a glob pattern.

Using the catalog for navigation reduces the overhead considerably because only the catalog needs to be downloaded and, optionally, decrypted. The actual chunks are only accessed if the metadata in the catalog is not enough or for the actual restore.

Similar to common UNIX shells `cd` and `ls` are the commands used to change working directory and list directory contents in the archive. `pwd` shows the full path of the current working directory with respect to the archive root.

Being able to quickly search the contents of the archive is a commonly needed feature. That's where the catalog is most valuable. For example:

```
pxar:/ > find etc/**/*.txt --select
"/etc/X11/rgb.txt"
pxar:/ > list-selected
etc/**/*.txt
pxar:/ > restore-selected /target/path
...
```

This will find and print all files ending in `.txt` located in `etc/` or a subdirectory and add the corresponding pattern to the list for subsequent restores. `list-selected` shows these patterns and `restore-selected` finally restores all files in the archive matching the patterns to `/target/path` on the local host. This will scan the whole archive.

With `restore /target/path` you can restore the sub-archive given by the current working directory to the local target path `/target/path` on your host. By additionally passing a glob pattern with `--pattern <glob>`, the restore is further limited to files matching the pattern. For example:

```
pxar:/ > cd /etc/
pxar:/etc/ > restore /target/ --pattern **/*.conf
...
```

The above will scan through all the directories below `/etc` and restore all files ending in `.conf`.

10.6.2 Mounting of Archives via FUSE

The [FUSE](#) implementation for the `pxar` archive allows you to mount a file archive as a read-only filesystem to a mountpoint on your host.

```
# proxmox-backup-client mount host/backup-client/2020-01-29T11:29:22Z root.pxar /mnt/
→ mountpoint
# ls /mnt/mountpoint
bin  dev  home  lib32  libx32  media  opt  root  sbin  sys  usr
boot  etc  lib  lib64  lost+found  mnt  proc  run  srv  tmp  var
```

This allows you to access the full contents of the archive in a seamless manner.

Note: As the FUSE connection needs to fetch and decrypt chunks from the backup server's datastore, this can cause some additional network and CPU load on your host, depending on the operations you perform on the mounted filesystem.

To unmount the filesystem use the `umount` command on the mountpoint:

```
# umount /mnt/mountpoint
```

10.7 Login and Logout

The client tool prompts you to enter the login password as soon as you want to access the backup server. The server checks your credentials and responds with a ticket that is valid for two hours. The client tool automatically stores that ticket and uses it for further requests to this server.

You can also manually trigger this login/logout using the `login` and `logout` commands:

```
# proxmox-backup-client login
Password: *****
```

To remove the ticket, issue a `logout`:

```
# proxmox-backup-client logout
```

10.8 Pruning and Removing Backups

You can manually delete a backup snapshot using the `forget` command:

```
# proxmox-backup-client forget <snapshot>
```

Caution: This command removes all archives in this backup snapshot. They will be inaccessible and unrecoverable.

Although manual removal is sometimes required, the `prune` command is normally used to systematically delete older backups. Prune lets you specify which backup snapshots you want to keep. The following retention options are available:

- keep-last <N>** Keep the last <N> backup snapshots.
- keep-hourly <N>** Keep backups for the last <N> hours. If there is more than one backup for a single hour, only the latest is kept.
- keep-daily <N>** Keep backups for the last <N> days. If there is more than one backup for a single day, only the latest is kept.
- keep-weekly <N>** Keep backups for the last <N> weeks. If there is more than one backup for a single week, only the latest is kept.

Note: Weeks start on Monday and end on Sunday. The software uses the [ISO week date](#) system and handles weeks at the end of the year correctly.

- keep-monthly <N>** Keep backups for the last <N> months. If there is more than one backup for a single month, only the latest is kept.
- keep-yearly <N>** Keep backups for the last <N> years. If there is more than one backup for a single year, only the latest is kept.

The retention options are processed in the order given above. Each option only covers backups within its time period. The next option does not take care of already covered backups. It will only consider older backups.

Unfinished and incomplete backups will be removed by the `prune` command unless they are newer than the last successful backup. In this case, the last failed backup is retained.

```
# proxmox-backup-client prune <group> --keep-daily 7 --keep-weekly 4 --keep-monthly 3
```

You can use the `--dry-run` option to test your settings. This only shows the list of existing snapshots and what actions prune would take.

```
# proxmox-backup-client prune host/elsa --dry-run --keep-daily 1 --keep-weekly 3
```

snapshot	keep
host/elsa/2019-12-04T13:20:37Z	1
host/elsa/2019-12-03T09:35:01Z	0
host/elsa/2019-11-22T11:54:47Z	1
host/elsa/2019-11-21T12:36:25Z	0
host/elsa/2019-11-10T10:42:20Z	1

Note: Neither the `prune` command nor the `forget` command free space in the chunk-store. The chunk-store still contains the data blocks. To free space you need to perform [Garbage Collection](#).

10.9 Garbage Collection

The `prune` command removes only the backup index files, not the data from the datastore. This task is left to the garbage collection command. It is recommended to carry out garbage collection on a regular basis.

The garbage collection works in two phases. In the first phase, all data blocks that are still in use are marked. In the second phase, unused data blocks are removed.

Note: This command needs to read all existing backup index files and touches the complete chunk-store. This can take a long time depending on the number of chunks and the speed of the underlying disks.

Note: The garbage collection will only remove chunks that haven't been used for at least one day (exactly 24h 5m). This grace period is necessary because chunks in use are marked by touching the chunk which updates the `atime` (access time) property. Filesystems are mounted with the `relatime` option by default. This results in a better performance by only updating the `atime` property if the last access has been at least 24 hours ago. The downside is, that touching a chunk within these 24 hours will not always update its `atime` property.

Chunks in the grace period will be logged at the end of the garbage collection task as *Pending removals*.

```
# proxmox-backup-client garbage-collect
starting garbage collection on store store2
Start GC phase1 (mark used chunks)
Start GC phase2 (sweep unused chunks)
percentage done: 1, chunk count: 219
percentage done: 2, chunk count: 453
...
percentage done: 99, chunk count: 21188
Removed bytes: 411368505
Removed chunks: 203
Original data bytes: 327160886391
Disk bytes: 52767414743 (16 %)
Disk chunks: 21221
Average chunk size: 2486565
TASK OK
```

10.10 Benchmarking

The backup client also comes with a benchmarking tool. This tool measures various metrics relating to compression and encryption speeds. You can run a benchmark using the `benchmark` subcommand of `proxmox-backup-client`:

```
# proxmox-backup-client benchmark
Uploaded 656 chunks in 5 seconds.
Time per request: 7659 microseconds.
TLS speed: 547.60 MB/s
SHA256 speed: 585.76 MB/s
```

(continues on next page)

(continued from previous page)

Compression speed: 1923.96 MB/s
Decompress speed: 7885.24 MB/s
AES256/GCM speed: 3974.03 MB/s

Name	Value
TLS (maximal backup upload speed)	547.60 MB/s (93%)
SHA256 checksum computation speed	585.76 MB/s (28%)
ZStd level 1 compression speed	1923.96 MB/s (89%)
ZStd level 1 decompression speed	7885.24 MB/s (98%)
AES256 GCM encryption speed	3974.03 MB/s (104%)

Note: The percentages given in the output table correspond to a comparison against a Ryzen 7 2700X. The TLS test connects to the local host, so there is no network involved.

You can also pass the `--output-format` parameter to output stats in json, rather than the default table format.

PROXMOX VE INTEGRATION

You need to define a new storage with type 'pbs' on your [Proxmox VE](#) node. The following example uses `store2` as storage name, and assumes the server address is `localhost`, and you want to connect as `user1@pbs`.

```
# pvesm add pbs store2 --server localhost --datastore store2
# pvesm set store2 --username user1@pbs --password <secret>
```

Note: If you would rather not pass your password as plain text, you can pass the `--password` parameter, without any arguments. This will cause the program to prompt you for a password upon entering the command.

If your backup server uses a self signed certificate, you need to add the certificate fingerprint to the configuration. You can get the fingerprint by running the following command on the backup server:

```
# proxmox-backup-manager cert info | grep Fingerprint
Fingerprint (sha256): 64:d3:ff:3a:50:38:53:5a:9b:f7:50:...:ab:fe
```

Please add that fingerprint to your configuration to establish a trust relationship:

```
# pvesm set store2 --fingerprint 64:d3:ff:3a:50:38:53:5a:9b:f7:50:...:ab:fe
```

After that you should be able to see storage status with:

```
# pvesm status --storage store2
```

Name	Type	Status	Total	Used	Available	%
store2	pbs	active	3905109820	1336687816	2568422004	34.23%

Having added the PBS datastore to [Proxmox VE](#), you can backup VMs and containers in the same way you would for any other storage device within the environment (see [PVE Admin Guide: Backup and Restore](#)).

PXAR COMMAND LINE TOOL

12.1 Description

pxar is a command line utility to create and manipulate archives in the *Proxmox File Archive Format* (.pxar). It is inspired by *casync file archive format*, which caters to a similar use-case. The .pxar format is adapted to fulfill the specific needs of the Proxmox Backup Server, for example, efficient storage of hardlinks. The format is designed to reduce storage space needed on the server by achieving a high level of deduplication.

12.2 Creating an Archive

Run the following command to create an archive of a folder named source:

```
# pxar create archive.pxar /path/to/source
```

This will create a new archive called archive.pxar with the contents of the source folder.

Note: pxar will not overwrite any existing archives. If an archive with the same name is already present in the target folder, the creation will fail.

By default, pxar will skip certain mountpoints and will not follow device boundaries. This design decision is based on the primary use case of creating archives for backups. It makes sense to not back up the contents of certain temporary or system specific files. To alter this behavior and follow device boundaries, use the `--all-file-systems` flag.

It is possible to exclude certain files and/or folders from the archive by passing the `--exclude` parameter with gitignore-style match patterns.

For example, you can exclude all files ending in .txt from the archive by running:

```
# pxar create archive.pxar /path/to/source --exclude '**/*.txt'
```

Be aware that the shell itself will try to expand all of the glob patterns before invoking pxar. In order to avoid this, all globs have to be quoted correctly.

It is possible to pass the `--exclude` parameter multiple times, in order to match more than one pattern. This allows you to use more complex file exclusion/inclusion behavior. However, it is recommended to use `.pxarexclude` files instead for such cases.

For example you might want to exclude all .txt files except for a specific one from the archive. This is achieved via the negated match pattern, prefixed by `!`. All the glob patterns are relative to the source directory.

```
# pxar create archive.pxar /path/to/source --exclude '**/*.txt' --exclude '!/folder/file.txt'
```

Note: The order of the glob match patterns matters as later ones override previous ones. Permutations of the same patterns lead to different results.

pxar will store the list of glob match patterns passed as parameters via the command line, in a file called `.pxarexclude-cli` at the root of the archive. If a file with this name is already present in the source folder during archive creation, this file is not included in the archive and the file containing the new patterns is added to the archive instead, the original file is not altered.

A more convenient and persistent way to exclude files from the archive is by placing the glob match patterns in `.pxarexclude` files. It is possible to create and place these files in any directory of the filesystem tree. These files must contain one pattern per line, again later patterns win over previous ones. The patterns control file exclusions of files present within the given directory or further below it in the tree. The behavior is the same as described in [Creating Backups](#).

12.3 Extracting an Archive

An existing archive, `archive.pxar`, is extracted to a target directory with the following command:

```
# pxar extract archive.pxar /path/to/target
```

If no target is provided, the content of the archive is extracted to the current working directory.

In order to restore only parts of an archive, single files, and/or folders, it is possible to pass the corresponding glob match patterns as additional parameters or to use the patterns stored in a file:

```
# pxar extract etc.pxar /restore/target/etc --pattern '**/*.conf'
```

The above example restores all `.conf` files encountered in any of the sub-folders in the archive `etc.pxar` to the target `/restore/target/etc`. A path to the file containing match patterns can be specified using the `--files-from` parameter.

12.4 List the Contents of an Archive

To display the files and directories contained in an archive `archive.pxar`, run the following command:

```
# pxar list archive.pxar
```

This displays the full path of each file or directory with respect to the archives root.

12.5 Mounting an Archive

pxar allows you to mount and inspect the contents of an archive via FUSE. In order to mount an archive named `archive.pxar` to the mountpoint `/mnt`, run the command:

```
# pxar mount archive.pxar /mnt
```

Once the archive is mounted, you can access its content under the given mountpoint.

```
# cd /mnt
# ls
bin  dev  home  lib32  libx32  media  opt  root  sbin  sys  usr
boot  etc  lib  lib64  lost+found  mnt  proc  run  srv  tmp  var
```

HOST SYSTEM ADMINISTRATION

[Proxmox Backup](#) is based on the famous [Debian](#) Linux distribution. That means that you have access to the whole world of Debian packages, and the base system is well documented. The [Debian Administrator's Handbook](#) is available online, and provides a comprehensive introduction to the Debian operating system.

A standard [Proxmox Backup](#) installation uses the default repositories from Debian, so you get bug fixes and security updates through that channel. In addition, we provide our own package repository to roll out all Proxmox related packages. This includes updates to some Debian packages when necessary.

We also deliver a specially optimized Linux kernel, where we enable all required virtualization and container features. That kernel includes drivers for [ZFS](#), and several hardware drivers. For example, we ship Intel network card drivers to support their newest hardware.

The following sections will concentrate on backup related topics. They either explain things which are different on [Proxmox Backup](#), or tasks which are commonly used on [Proxmox Backup](#). For other topics, please refer to the standard Debian documentation.

13.1 ZFS on Linux

ZFS is a combined file system and logical volume manager designed by Sun Microsystems. There is no need to manually compile ZFS modules - all packages are included.

By using ZFS, it's possible to achieve maximum enterprise features with low budget hardware, but also high performance systems by leveraging SSD caching or even SSD only setups. ZFS can replace cost intense hardware raid cards by moderate CPU and memory load combined with easy management.

General ZFS advantages

- Easy configuration and management with GUI and CLI.
- Reliable
- Protection against data corruption
- Data compression on file system level
- Snapshots
- Copy-on-write clone
- Various raid levels: RAID0, RAID1, RAID10, RAIDZ-1, RAIDZ-2 and RAIDZ-3
- Can use SSD for cache
- Self healing
- Continuous integrity checking
- Designed for high storage capacities

- Asynchronous replication over network
- Open Source
- Encryption

13.1.1 Hardware

ZFS depends heavily on memory, so you need at least 8GB to start. In practice, use as much you can get for your hardware/budget. To prevent data corruption, we recommend the use of high quality ECC RAM.

If you use a dedicated cache and/or log disk, you should use an enterprise class SSD (e.g. Intel SSD DC S3700 Series). This can increase the overall performance significantly.

IMPORTANT: Do not use ZFS on top of hardware controller which has its own cache management. ZFS needs to directly communicate with disks. An HBA adapter is the way to go, or something like LSI controller flashed in IT mode.

13.1.2 ZFS Administration

This section gives you some usage examples for common tasks. ZFS itself is really powerful and provides many options. The main commands to manage ZFS are *zfs* and *zpool*. Both commands come with great manual pages, which can be read with:

```
# man zpool
# man zfs
```

Create a new zpool

To create a new pool, at least one disk is needed. The *ashift* should have the same sector-size (2 power of *ashift*) or larger as the underlying disk.

```
# zpool create -f -o ashift=12 <pool> <device>
```

Create a new pool with RAID-0

Minimum 1 disk

```
# zpool create -f -o ashift=12 <pool> <device1> <device2>
```

Create a new pool with RAID-1

Minimum 2 disks

```
# zpool create -f -o ashift=12 <pool> mirror <device1> <device2>
```

Create a new pool with RAID-10

Minimum 4 disks

```
# zpool create -f -o ashift=12 <pool> mirror <device1> <device2> mirror <device3> <device4>
```


Create a new pool with RAIDZ-1

Minimum 3 disks

```
# zpool create -f -o ashift=12 <pool> raidz1 <device1> <device2> <device3>
```

Create a new pool with RAIDZ-2

Minimum 4 disks

```
# zpool create -f -o ashift=12 <pool> raidz2 <device1> <device2> <device3> <device4>
```

Create a new pool with cache (L2ARC)

It is possible to use a dedicated cache drive partition to increase the performance (use SSD).

As <device> it is possible to use more devices, like it's shown in "Create a new pool with RAID*".

```
# zpool create -f -o ashift=12 <pool> <device> cache <cache_device>
```

Create a new pool with log (ZIL)

It is possible to use a dedicated cache drive partition to increase the performance (SSD).

As <device> it is possible to use more devices, like it's shown in "Create a new pool with RAID*".

```
# zpool create -f -o ashift=12 <pool> <device> log <log_device>
```

Add cache and log to an existing pool

If you have a pool without cache and log. First partition the SSD in 2 partition with *parted* or *gdisk*

Important: Always use GPT partition tables.

The maximum size of a log device should be about half the size of physical memory, so this is usually quite small. The rest of the SSD can be used as cache.

```
# zpool add -f <pool> log <device-part1> cache <device-part2>
```

Changing a failed device

```
# zpool replace -f <pool> <old device> <new device>
```

Changing a failed bootable device

Depending on how Proxmox Backup was installed it is either using *grub* or *systemd-boot* as boot-loader.

The first steps of copying the partition table, reissuing GUIDs and replacing the ZFS partition are the same. To make the system bootable from the new disk, different steps are needed which depend on the bootloader in use.

```
# sgdisk <healthy bootable device> -R <new device>
# sgdisk -G <new device>
# zpool replace -f <pool> <old zfs partition> <new zfs partition>
```

Note: Use the `zpool status -v` command to monitor how far the resilvering process of the new disk has progressed.

With `systemd-boot`:

```
# pve-efiboot-tool format <new disk's ESP>
# pve-efiboot-tool init <new disk's ESP>
```

Note: *ESP* stands for EFI System Partition, which is setup as partition #2 on bootable disks setup by the {pve} installer since version 5.4. For details, see [xref:sysboot_systemd_boot_setup\[Setting up a new partition for use as synced ESP\]](#).

With `grub`:

Usually `grub.cfg` is located in `/boot/grub/grub.cfg`

```
# grub-install <new disk>
# grub-mkconfig -o /path/to/grub.cfg
```

Activate E-Mail Notification

ZFS comes with an event daemon, which monitors events generated by the ZFS kernel module. The daemon can also send emails on ZFS events like pool errors. Newer ZFS packages ship the daemon in a separate package, and you can install it using `apt-get`:

```
# apt-get install zfs-zed
```

To activate the daemon it is necessary to edit `/etc/zfs/zed.d/zed.rc` with your favorite editor, and uncomment the `ZED_EMAIL_ADDR` setting:

```
ZED_EMAIL_ADDR="root"
```

Please note Proxmox Backup forwards mails to `root` to the email address configured for the root user.

IMPORTANT: The only setting that is required is `ZED_EMAIL_ADDR`. All other settings are optional.

Limit ZFS Memory Usage

It is good to use at most 50 percent (which is the default) of the system memory for ZFS ARC to prevent performance shortage of the host. Use your preferred editor to change the configuration in `/etc/modprobe.d/zfs.conf` and insert:

```
options zfs zfs_arc_max=8589934592
```

This example setting limits the usage to 8GB.

Important: If your root file system is ZFS you must update your `initramfs` every time this value changes:

```
# update-initramfs -u
```

SWAP on ZFS

Swap-space created on a zvol may generate some troubles, like blocking the server or generating a high IO load, often seen when starting a Backup to an external Storage.

We strongly recommend to use enough memory, so that you normally do not run into low memory situations. Should you need or want to add swap, it is preferred to create a partition on a physical disk and use it as swap device. You can leave some space free for this purpose in the advanced options of the installer. Additionally, you can lower the *swappiness* value. A good value for servers is 10:

```
# sysctl -w vm.swappiness=10
```

To make the swappiness persistent, open `/etc/sysctl.conf` with an editor of your choice and add the following line:

```
vm.swappiness = 10
```

Table 1: Linux kernel *swappiness* parameter values :widths:auto

Value	Strategy
vm.swappiness = 0	The kernel will swap only to avoid an 'out of memory' condition
vm.swappiness = 1	Minimum amount of swapping without disabling it entirely.
vm.swappiness = 10	Sometimes recommended to improve performance when sufficient memory exists in a system.
vm.swappiness = 60	The default value.
vm.swappiness = 100	The kernel will swap aggressively.

ZFS Compression

To activate compression: .. code-block:: console

```
# zpool set compression=lz4 <pool>
```

We recommend using the *lz4* algorithm, since it adds very little CPU overhead. Other algorithms such as *lzjb* and *gzip-N* (where *N* is an integer 1-9 representing the compression ratio, 1 is fastest and 9 is best compression) are also available. Depending on the algorithm and how compressible the data is, having compression enabled can even increase I/O performance.

You can disable compression at any time with: .. code-block:: console

```
# zfs set compression=off <dataset>
```

Only new blocks will be affected by this change.

ZFS Special Device

Since version 0.8.0 ZFS supports *special* devices. A *special* device in a pool is used to store metadata, deduplication tables, and optionally small file blocks.

A *special* device can improve the speed of a pool consisting of slow spinning hard disks with a lot of metadata changes. For example workloads that involve creating, updating or deleting a large

number of files will benefit from the presence of a *special* device. ZFS datasets can also be configured to store whole small files on the *special* device which can further improve the performance. Use fast SSDs for the *special* device.

Important: The redundancy of the *special* device should match the one of the pool, since the *special* device is a point of failure for the whole pool.

Warning: Adding a *special* device to a pool cannot be undone!

Create a pool with *special* device and RAID-1:

```
# zpool create -f -o ashift=12 <pool> mirror <device1> <device2> special mirror <device3>
↪ <device4>
```

Adding a *special* device to an existing pool with RAID-1:

```
# zpool add <pool> special mirror <device1> <device2>
```

ZFS datasets expose the *special_small_blocks=<size>* property. *size* can be 0 to disable storing small file blocks on the *special* device or a power of two in the range between 512B to 128K. After setting the property new file blocks smaller than *size* will be allocated on the *special* device.

Important: If the value for *special_small_blocks* is greater than or equal to the *recordsize* (default 128K) of the dataset, *all* data will be written to the *special* device, so be careful!

Setting the *special_small_blocks* property on a pool will change the default value of that property for all child ZFS datasets (for example all containers in the pool will opt in for small file blocks).

Opt in for all file smaller than 4K-blocks pool-wide:

```
# zfs set special_small_blocks=4K <pool>
```

Opt in for small file blocks for a single dataset:

```
# zfs set special_small_blocks=4K <pool>/<filesystem>
```

Opt out from small file blocks for a single dataset:

```
# zfs set special_small_blocks=0 <pool>/<filesystem>
```

Troubleshooting

Corrupted cachefile

In case of a corrupted ZFS cachefile, some volumes may not be mounted during boot until mounted manually later.

For each pool, run:

```
# zpool set cachefile=/etc/zfs/zpool.cache POOLNAME
```

and afterwards update the *initramfs* by running:

```
# update-initramfs -u -k all
```

and finally reboot your node.

Sometimes the ZFS cachefile can get corrupted, and *zfs-import-cache.service* doesn't import the pools that aren't present in the cachefile.

Another workaround to this problem is enabling the *zfs-import-scan.service*, which searches and imports pools via device scanning (usually slower).

14.1 What distribution is Proxmox Backup Server (PBS) based on?

Proxmox Backup Server is based on [Debian GNU/Linux](#).

14.2 Which platforms are supported as a backup source (client)?

The client tool works on most modern Linux systems, meaning you are not limited to Debian-based backups.

14.3 Will Proxmox Backup Server run on a 32-bit processor?

Proxmox Backup Server only supports 64-bit CPUs (AMD or Intel). There are no future plans to support 32-bit processors.

14.4 How long will my Proxmox Backup Server version be supported?

Proxmox Backup Version	Debian Version	First Re-lease	Debian EOL	Proxmox Backup EOL
Proxmox Backup 1.x	Debian 10 (Buster)	tba	tba	tba

14.5 Can I copy or synchronize my datastore to another location?

Proxmox Backup Server allows you to copy or synchronize datastores to other locations, through the use of *Remotes* and *Sync Jobs*. *Remote* is the term given to a separate server, which has a datastore that can be synced to a local store. A *Sync Job* is the process which is used to pull the contents of a datastore from a *Remote* to a local datastore.

14.6 Can Proxmox Backup Server verify data integrity of a backup archive?

Proxmox Backup Server uses a built-in SHA-256 checksum algorithm, to ensure data integrity. Within each backup, a manifest file (index.json) is created, which contains a list of all the backup files, along with their sizes and checksums. This manifest file is used to verify the integrity of each backup.

14.7 When backing up to remote servers, do I have to trust the remote server?

Proxmox Backup Server supports client-side encryption, meaning your data is encrypted before it reaches the server. Thus, in the event that an attacker gains access to the server, they will not be able to read the data.

Note: Encryption is not enabled by default. To set up encryption, see the [Encryption](#) section of the Proxmox Backup Server Administration Guide.

14.8 Is the backup incremental/deduplicated?

With Proxmox Backup Server, backups are sent incremental and data is deduplicated on the server. This minimizes both the storage consumed and the network impact.

COMMAND SYNTAX

A.1 proxmox-backup-client

`proxmox-backup-client backup {<backupspec>} [OPTIONS]`

Create (host) backup.

<backupspec> **<array>** List of backup source specifications ([<label.ext>:<path>] ...)

Optional parameters:

- backup-id** **<string>** Backup ID.
- backup-time** **<integer>** **(1547797308 - N)** Backup time (Unix epoch.)
- backup-type** **<string>** Backup type.
- chunk-size** **<integer>** **(64 - 4096) (default=4096)** Chunk size in KB. Must be a power of 2.
- crypt-mode** **<string>** **(default=encrypt)** Defines whether data is encrypted (using an AEAD cipher), only signed, or neither.
- entries-max** **<integer>** **(default=1048576)** Max number of entries to hold in memory.
- exclude** **<array>** List of paths or patterns for matching files to exclude.
- include-dev** **<array>** Include mountpoints with same st_dev number (see `man fstat`) as specified files.
- keyfd** **<integer>** **(0 - N)** Pass an encryption key via an already opened file descriptor.
- keyfile** **<string>** Path to encryption key. All data will be encrypted using this key.
- repository** **<string>** Repository URL.
- skip-lost-and-found** **<boolean>** Skip lost+found directory.
- verbose** **<boolean>** Verbose output.

`proxmox-backup-client benchmark [OPTIONS]`

Run benchmark tests

Optional parameters:

- keyfile** **<string>** Path to encryption key. All data will be encrypted using this key.
- output-format** **<string>** Output format.
- repository** **<string>** Repository URL.
- verbose** **<boolean>** Verbose output.

`proxmox-backup-client catalog dump <snapshot> [OPTIONS]`

Dump catalog.

<snapshot> **<string>** Snapshot path.

Optional parameters:

--keyfd <integer> (0 - N) Pass an encryption key via an already opened file descriptor.

--keyfile <string> Path to encryption key.

--repository <string> Repository URL.

`proxmox-backup-client catalog shell <snapshot> <archive-name> [OPTIONS]`

Shell to interactively inspect and restore snapshots.

<snapshot> **<string>** Group/Snapshot path.

<archive-name> **<string>** Backup archive name.

Optional parameters:

--keyfd <integer> (0 - N) Pass an encryption key via an already opened file descriptor.

--keyfile <string> Path to encryption key.

--repository <string> Repository URL.

`proxmox-backup-client change-owner <group> <new-owner> [OPTIONS]`

Change owner of a backup group

<group> **<string>** Backup group.

<new-owner> **<string>** User ID

Optional parameters:

--repository <string> Repository URL.

`proxmox-backup-client files <snapshot> [OPTIONS]`

List snapshot files.

<snapshot> **<string>** Snapshot path.

Optional parameters:

--output-format <string> Output format.

--repository <string> Repository URL.

`proxmox-backup-client forget <snapshot> [OPTIONS]`

Forget (remove) backup snapshots.

<snapshot> **<string>** Snapshot path.

Optional parameters:

--repository <string> Repository URL.

`proxmox-backup-client garbage-collect [OPTIONS]`

Start garbage collection for a specific repository.

Optional parameters:

- output-format <string>** Output format.
 - repository <string>** Repository URL.
-

`proxmox-backup-client help [{<command>}] [OPTIONS]`

Get help about specified command (or sub-command).

<command> <array> Command. This may be a list in order to specify nested sub-commands.

Optional parameters:

- verbose <boolean>** Verbose help.

`proxmox-backup-client key change-passphrase [<path>] [OPTIONS]`

Change the encryption key's password.

<path> <string> Key file. Without this the default key's password will be changed.

Optional parameters:

- kdf <string> (default=scrypt)** Key derivation function for password protected encryption keys.
-

`proxmox-backup-client key create [<path>] [OPTIONS]`

Create a new encryption key.

<path> <string> Output file. Without this the key will become the new default encryption key.

Optional parameters:

- kdf <string> (default=scrypt)** Key derivation function for password protected encryption keys.
-

`proxmox-backup-client key create-master-key`

Create an RSA public/private key pair used to put an encrypted version of the symmetric backup encryption key onto the backup server along with each backup.

`proxmox-backup-client key import-master-pubkey <path>`

Import an RSA public key used to put an encrypted version of the symmetric backup encryption key onto the backup server along with each backup.

<path> <string> Path to the PEM formatted RSA public key.

`proxmox-backup-client key paperkey [<path>] [OPTIONS]`

Generate a printable, human readable text file containing the encryption key.

This also includes a scanable QR code for fast key restore.

<path> <string> Key file. Without this the default key's will be used.

Optional parameters:

--output-format <string> Paperkey output format

--subject <string> Include the specified subject as titel text.

proxmox-backup-client list [OPTIONS]

List backup groups.

Optional parameters:

--output-format <string> Output format.

--repository <string> Repository URL.

proxmox-backup-client login [OPTIONS]

Try to login. If successful, store ticket.

Optional parameters:

--repository <string> Repository URL.

proxmox-backup-client logout [OPTIONS]

Logout (delete stored ticket).

Optional parameters:

--repository <string> Repository URL.

proxmox-backup-client map <snapshot> <archive-name> [OPTIONS]

Map a drive image from a VM backup to a local loopback device. Use 'unmap' to undo. WARNING: Only do this with *trusted* backups!

<snapshot> <string> Group/Snapshot path.

<archive-name> <string> Backup archive name.

Optional parameters:

--keyfile <string> Path to encryption key.

--repository <string> Repository URL.

--verbose <boolean> (default=false) Verbose output and stay in foreground.

proxmox-backup-client mount <snapshot> <archive-name> <target> [OPTIONS]

Mount pxar archive.

<snapshot> <string> Group/Snapshot path.

<archive-name> <string> Backup archive name.

<target> <string> Target directory path.

Optional parameters:

--keyfile <string> Path to encryption key.

- repository <string>** Repository URL.
 - verbose <boolean> (default=false)** Verbose output and stay in foreground.
-

proxmox-backup-client prune <group> [OPTIONS]

Prune a backup repository.

<group> <string> Backup group.

Optional parameters:

- dry-run <boolean>** Just show what prune would do, but do not delete anything.
 - keep-daily <integer> (1 - N)** Number of daily backups to keep.
 - keep-hourly <integer> (1 - N)** Number of hourly backups to keep.
 - keep-last <integer> (1 - N)** Number of backups to keep.
 - keep-monthly <integer> (1 - N)** Number of monthly backups to keep.
 - keep-weekly <integer> (1 - N)** Number of weekly backups to keep.
 - keep-yearly <integer> (1 - N)** Number of yearly backups to keep.
 - output-format <string>** Output format.
 - quiet <boolean>** Minimal output - only show removals.
 - repository <string>** Repository URL.
-

proxmox-backup-client restore <snapshot> <archive-name> <target> [OPTIONS]

Restore backup repository.

<snapshot> <string> Group/Snapshot path.

<archive-name> <string> Backup archive name.

<target> <string> Target directory path. Use '-' to write to standard output.

We do not extract '.pxar' archives when writing to standard output.

Optional parameters:

- allow-existing-dirs <boolean>** Do not fail if directories already exists.
 - crypt-mode <string> (default=encrypt)** Defines whether data is encrypted (using an AEAD cipher), only signed, or neither.
 - keyfd <integer> (0 - N)** Pass an encryption key via an already opened file descriptor.
 - keyfile <string>** Path to encryption key. All data will be encrypted using this key.
 - repository <string>** Repository URL.
-

proxmox-backup-client snapshots [<group>] [OPTIONS]

List backup snapshots.

<group> <string> Backup group.

Optional parameters:

- output-format <string>** Output format.
- repository <string>** Repository URL.

`proxmox-backup-client status [OPTIONS]`

Get repository status.

Optional parameters:

--output-format <string> Output format.

--repository <string> Repository URL.

`proxmox-backup-client task list [OPTIONS]`

List running server tasks for this repo user

Optional parameters:

--all <boolean> Also list stopped tasks.

--limit <integer> (1 - 1000) (default=50) The maximal number of tasks to list.

--output-format <string> Output format.

--repository <string> Repository URL.

`proxmox-backup-client task log <upid> [OPTIONS]`

Display the task log.

<upid> <string> Unique Process/Task ID.

Optional parameters:

--repository <string> Repository URL.

`proxmox-backup-client task stop <upid> [OPTIONS]`

Try to stop a specific task.

<upid> <string> Unique Process/Task ID.

Optional parameters:

--repository <string> Repository URL.

`proxmox-backup-client unmap [<name>]`

Unmap a loop device mapped with 'map' and release all resources.

<name> <string> Archive name, path to loopdev (/dev/loopX) or loop device number.
Omit to list all current mappings and force cleaning up leftover instances.

`proxmox-backup-client upload-log <snapshot> <logfile> [OPTIONS]`

Upload backup log file.

<snapshot> <string> Group/Snapshot path.

<logfile> <string> The path to the log file you want to upload.

Optional parameters:

--crypt-mode <string> (default=encrypt) Defines whether data is encrypted (using an AEAD cipher), only signed, or neither.

- keyfd <integer> (0 - N)** Pass an encryption key via an already opened file descriptor.
 - keyfile <string>** Path to encryption key. All data will be encrypted using this key.
 - repository <string>** Repository URL.
-

`proxmox-backup-client version [OPTIONS]`

Show client and optional server version

Optional parameters:

- output-format <string>** Output format.
- repository <string>** Repository URL.

A.1.1 Catalog Shell Commands

Those command are available when you start an interactive restore shell:

`proxmox-backup-client shell <snapshot> <name.pxar>`

`cd [<path>]`

Change the current working directory to the new directory

<path> <string> target path.

`clear-selected`

Clear the list of files selected for restore.

`deselect <path>`

Deselect an entry for restore.

This will return an error if the entry was not found in the list of entries selected for restore.

<path> <string> path to entry to remove from list.

`exit`

Exit the shell

`find <pattern> [OPTIONS]`

Find entries in the catalog matching the given match pattern.

<pattern> <string> Match pattern for matching files in the catalog.

Optional parameters:

- select <boolean> (default=false)** Add matching filenames to list for restore.
-

`help [{<command>}] [OPTIONS]`

Get help about specified command (or sub-command).

<command> <array> Command. This may be a list in order to specify nested sub-commands.

Optional parameters:

--verbose <boolean> Verbose help.

list-selected [OPTIONS]

List entries currently selected for restore.

Optional parameters:

--patterns <boolean> (default=false) List match patterns instead of the matching files.

ls [<path>]

List the content of working directory or given path.

<path> <string> target path.

pwd

List the current working directory.

restore <target> [OPTIONS]

Restore the sub-archive given by the current working directory to target.

By further providing a pattern, the restore can be limited to a narrower subset of this sub-archive. If pattern is not present or empty, the full archive is restored to target.

<target> <string> target path for restore on local filesystem.

Optional parameters:

--pattern <string> match pattern to limit files for restore.

restore-selected <target>

Restore the selected entries to the given target path.

Target must not exist on the clients filesystem.

<target> <string> target path for restore on local filesystem.

select <path>

Select an entry for restore.

This will return an error if the entry is already present in the list or if an invalid path was provided.

<path> <string> target path.

stat <path>

Read the metadata for a given directory entry.

This is expensive because the data has to be read from the paxr archive, which means reading over the network.

<path> <string> target path.

A.2 proxmox-backup-manager

`proxmox-backup-manager acl list [OPTIONS]`

Access Control list.

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager acl update <path> <role> [OPTIONS]`

Update Access Control List (ACLs).

<path> <string> Access control path.

<role> <string> Role

Optional parameters:

--delete <boolean> Remove permissions (instead of adding it).

--digest <string> Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

--group <string> Group ID

--propagate <boolean> (default=true) Allow to propagate (inherit) permissions.

--userid <string> User ID

`proxmox-backup-manager cert info`

Display node certificate information.

`proxmox-backup-manager cert update [OPTIONS]`

Update node certificates and generate all needed files/directories.

Optional parameters:

--force <boolean> Force generation of new SSL certificate.

`proxmox-backup-manager datastore create <name> <path> [OPTIONS]`

Create new datastore config.

<name> <string> Datastore name.

<path> <string> Directory name

Optional parameters:

--comment <string> Comment (single line).

--gc-schedule <string> Run garbage collection job at specified schedule.

--keep-daily <integer> (1 - N) Number of daily backups to keep.

--keep-hourly <integer> (1 - N) Number of hourly backups to keep.

--keep-last <integer> (1 - N) Number of backups to keep.

--keep-monthly <integer> (1 - N) Number of monthly backups to keep.

--keep-weekly <integer> (1 - N) Number of weekly backups to keep.

--keep-yearly <integer> (1 - N) Number of yearly backups to keep.

--prune-schedule <string> Run prune job at specified schedule.

`proxmox-backup-manager datastore list [OPTIONS]`

Datastore list.

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager datastore remove <name> [OPTIONS]`

Remove a datastore configuration.

<name> <string> Datastore name.

Optional parameters:

--digest <string> Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

`proxmox-backup-manager datastore show <name> [OPTIONS]`

Show datastore configuration

<name> <string> Datastore name.

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager datastore update <name> [OPTIONS]`

Update datastore config.

<name> <string> Datastore name.

Optional parameters:

--comment <string> Comment (single line).

--delete <array> List of properties to delete.

--digest <string> Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

--gc-schedule <string> Run garbage collection job at specified schedule.

--keep-daily <integer> (1 - N) Number of daily backups to keep.

--keep-hourly <integer> (1 - N) Number of hourly backups to keep.

--keep-last <integer> (1 - N) Number of backups to keep.

--keep-monthly <integer> (1 - N) Number of monthly backups to keep.

--keep-weekly <integer> (1 - N) Number of weekly backups to keep.

--keep-yearly <integer> (1 - N) Number of yearly backups to keep.

--prune-schedule <string> Run prune job at specified schedule.

`proxmox-backup-manager disk fs create <name> --disk <string> [OPTIONS]`

Create a Filesystem on an unused disk. Will be mounted under '/mnt/datastore/<name>'.

<name> <string> Datastore name.

--disk <string> Block device name (/sys/block/<name>).

Optional parameters:

--add-datastore <boolean> Configure a datastore using the directory.

--filesystem <string>

`proxmox-backup-manager disk fs list [OPTIONS]`

List systemd datastore mount units.

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager disk initialize <disk> [OPTIONS]`

Initialize empty Disk with GPT

<disk> <string> Block device name (/sys/block/<name>).

Optional parameters:

--uuid <string> UUID for the GPT table.

`proxmox-backup-manager disk list [OPTIONS]`

Local disk list.

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager disk smart-attributes <disk> [OPTIONS]`

Show SMART attributes.

<disk> <string> Block device name (/sys/block/<name>).

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager disk zpool create <name> --devices <string>
--raidlevel <string> [OPTIONS]`

create a zfs pool

<name> <string> Datastore name.

--devices <string> A list of disk names, comma separated.

--raidlevel <string> The ZFS RAID level to use.

Optional parameters:

--add-datastore <boolean> Configure a datastore using the zpool.

--ashift <integer> (9 - 16) (default=12) Pool sector size exponent.

--compression <string> (default=0n) The ZFS compression algorithm to use.

`proxmox-backup-manager disk zpool list [OPTIONS]`

Local zfs pools.

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager dns get [OPTIONS]`

Read DNS settings

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager dns set [OPTIONS]`

Update DNS settings.

Optional parameters:

--delete <array> List of properties to delete.

--digest <string> Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

--dns1 <string> First name server IP address.

--dns2 <string> Second name server IP address.

--dns3 <string> Third name server IP address.

--search <string> Search domain for host-name lookup.

`proxmox-backup-manager garbage-collection start <store> [OPTIONS]`

Start garbage collection for a specific datastore.

<store> <string> Datastore name.

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager garbage-collection status <store> [OPTIONS]`

Show garbage collection status for a specific datastore.

<store> <string> Datastore name.

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager help [{<command>}] [OPTIONS]`

Get help about specified command (or sub-command).

<command> <array> Command. This may be a list in order to specify nested sub-commands.

Optional parameters:

--verbose <boolean> Verbose help.

proxmox-backup-manager network changes

Show pending configuration changes (diff)

proxmox-backup-manager network create <iface> [OPTIONS]

Create network interface configuration.

<iface> <string> Network interface name.

Optional parameters:

--autostart <boolean> Autostart interface.

--bond-primary <string> Network interface name.

--bond_mode <string> Linux Bond Mode

--bond_xmit_hash_policy <string> Bond Transmit Hash Policy for LACP (802.3ad)

--bridge_ports <string> A list of network devices, comma separated.

--bridge_vlan_aware <boolean> Enable bridge vlan support.

--cidr <string> IPv4 address with netmask (CIDR notation).

--cidr6 <string> IPv6 address with netmask (CIDR notation).

--comments <string> Comments (inet, may span multiple lines)

--comments6 <string> Comments (inet5, may span multiple lines)

--gateway <string> IPv4 address.

--gateway6 <string> IPv6 address.

--method <string> Interface configuration method

--method6 <string> Interface configuration method

--mtu <integer> (46 - 65535) (default=1500) Maximum Transmission Unit.

--slaves <string> A list of network devices, comma separated.

--type <string> Network interface type

proxmox-backup-manager network list [OPTIONS]

Network device list.

Optional parameters:

--output-format <string> Output format.

proxmox-backup-manager network reload

Reload network configuration (requires ifupdown2).

proxmox-backup-manager network remove <iface> [OPTIONS]

Remove network interface configuration.

<iface> <string> Network interface name.

Optional parameters:

--digest <string> Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

`proxmox-backup-manager network revert`

Revert network configuration (rm /etc/network/interfaces.new).

`proxmox-backup-manager network update <iface> [OPTIONS]`

Update network interface config.

<iface> <string> Network interface name.

Optional parameters:

--autostart <boolean> Autostart interface.

--bond-primary <string> Network interface name.

--bond_mode <string> Linux Bond Mode

--bond_xmit_hash_policy <string> Bond Transmit Hash Policy for LACP (802.3ad)

--bridge_ports <string> A list of network devices, comma separated.

--bridge_vlan_aware <boolean> Enable bridge vlan support.

--cidr <string> IPv4 address with netmask (CIDR notation).

--cidr6 <string> IPv6 address with netmask (CIDR notation).

--comments <string> Comments (inet, may span multiple lines)

--comments6 <string> Comments (inet5, may span multiple lines)

--delete <array> List of properties to delete.

--digest <string> Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

--gateway <string> IPv4 address.

--gateway6 <string> IPv6 address.

--method <string> Interface configuration method

--method6 <string> Interface configuration method

--mtu <integer> (46 - 65535) (default=1500) Maximum Transmission Unit.

--slaves <string> A list of network devices, comma separated.

--type <string> Network interface type

`proxmox-backup-manager pull <remote> <remote-store> <local-store> [OPTIONS]`

Sync datastore from another repository

<remote> <string> Remote ID.

<remote-store> <string> Datastore name.

<local-store> <string> Datastore name.

Optional parameters:

--output-format <string> Output format.

--remove-vanished <boolean> (default=true) Delete vanished backups. This remove the local copy if the remote backup was deleted.

`proxmox-backup-manager remote create <name> --host <string> --password <string> --userid <string> [OPTIONS]`

Create new remote.

<name> <string> Remote ID.

--host <string> DNS name or IP address.

--password <string> Password or auth token for remote host.

--userid <string> User ID

Optional parameters:

--comment <string> Comment (single line).

--fingerprint <string> X509 certificate fingerprint (sha256).

--port <integer> (default=8007) The (optional) port.

`proxmox-backup-manager remote list [OPTIONS]`

List configured remotes.

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager remote remove <name> [OPTIONS]`

Remove a remote from the configuration file.

<name> <string> Remote ID.

Optional parameters:

--digest <string> Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

`proxmox-backup-manager remote show <name> [OPTIONS]`

Show remote configuration

<name> <string> Remote ID.

Optional parameters:

--output-format <string> Output format.

`proxmox-backup-manager remote update <name> [OPTIONS]`

Update remote configuration.

<name> <string> Remote ID.

Optional parameters:

--comment <string> Comment (single line).

--delete <array> List of properties to delete.

--digest <string> Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

- fingerprint <string>** X509 certificate fingerprint (sha256).
- host <string>** DNS name or IP address.
- password <string>** Password or auth token for remote host.
- port <integer>** The (optional) port.
- userid <string>** User ID

```
proxmox-backup-manager sync-job create <id> --remote <string>  
--remote-store <string> --store <string> [OPTIONS]
```

Create a new sync job.

- <id> <string>** Job ID.
- remote <string>** Remote ID.
- remote-store <string>** Datastore name.
- store <string>** Datastore name.

Optional parameters:

- comment <string>** Comment (single line).
- remove-vanished <boolean> (default=true)** Delete vanished backups. This remove the local copy if the remote backup was deleted.
- schedule <string>** Run sync job at specified schedule.

```
proxmox-backup-manager sync-job list [OPTIONS]
```

Sync job list.

Optional parameters:

- output-format <string>** Output format.

```
proxmox-backup-manager sync-job remove <id> [OPTIONS]
```

Remove a sync job configuration

- <id> <string>** Job ID.

Optional parameters:

- digest <string>** Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

```
proxmox-backup-manager sync-job show <id> [OPTIONS]
```

Show sync job configuration

- <id> <string>** Job ID.

Optional parameters:

- output-format <string>** Output format.

```
proxmox-backup-manager sync-job update <id> [OPTIONS]
```

Update sync job config.

- <id> <string>** Job ID.

Optional parameters:

- comment <string>** Comment (single line).
- delete <array>** List of properties to delete.
- digest <string>** Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.
- remote <string>** Remote ID.
- remote-store <string>** Datastore name.
- remove-vanished <boolean> (default=true)** Delete vanished backups. This remove the local copy if the remote backup was deleted.
- schedule <string>** Run sync job at specified schedule.
- store <string>** Datastore name.

`proxmox-backup-manager task list [OPTIONS]`

List running server tasks.

Optional parameters:

- all <boolean>** Also list stopped tasks.
- limit <integer> (1 - 1000) (default=50)** The maximal number of tasks to list.
- output-format <string>** Output format.

`proxmox-backup-manager task log <upid>`

Display the task log.

<upid> <string> Unique Process/Task ID.

`proxmox-backup-manager task stop <upid>`

Try to stop a specific task.

<upid> <string> Unique Process/Task ID.

`proxmox-backup-manager user create <userid> [OPTIONS]`

Create new user.

<userid> <string> User ID

Optional parameters:

- comment <string>** Comment (single line).
- email <string>** E-Mail Address.
- enable <boolean> (default=true)** Enable the account (default). You can set this to '0' to disable the account.
- expire <integer> (0 - N) (default=0)** Account expiration date (seconds since epoch). '0' means no expiration date.
- firstname <string>** First name.
- lastname <string>** Last name.
- password <string>** User Password.

```
proxmox-backup-manager user list [OPTIONS]
```

List configured users.

Optional parameters:

--output-format <string> Output format.

```
proxmox-backup-manager user remove <userid> [OPTIONS]
```

Remove a user from the configuration file.

<userid> <string> User ID

Optional parameters:

--digest <string> Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

```
proxmox-backup-manager user update <userid> [OPTIONS]
```

Update user configuration.

<userid> <string> User ID

Optional parameters:

--comment <string> Comment (single line).

--digest <string> Prevent changes if current configuration file has different SHA256 digest. This can be used to prevent concurrent modifications.

--email <string> E-Mail Address.

--enable <boolean> (default=true) Enable the account (default). You can set this to '0' to disable the account.

--expire <integer> (0 - N) (default=0) Account expiration date (seconds since epoch). '0' means no expiration date.

--firstname <string> First name.

--lastname <string> Last name.

--password <string> User Password.

```
proxmox-backup-manager verify <store> [OPTIONS]
```

Verify backups

<store> <string> Datastore name.

Optional parameters:

--output-format <string> Output format.

A.3 pxar

`pxar create <archive> <source> [OPTIONS]`

Create a new .pxar archive.

<archive> <string> Archive name.

<source> <string> Source directory.

Optional parameters:

- all-file-systems <boolean> (default=false)** Include mounted sudirs.
- entries-max <integer> (0 - 9223372036854775807) (default=1048576)**
Max number of entries loaded at once into memory
- exclude <array>** List of paths or pattern matching files to exclude.
- no-acls <boolean> (default=false)** Ignore access control list entries.
- no-device-nodes <boolean> (default=false)** Ignore device nodes.
- no-fcaps <boolean> (default=false)** Ignore file capabilities.
- no-fifos <boolean> (default=false)** Ignore fifos.
- no-sockets <boolean> (default=false)** Ignore sockets.
- no-xattrs <boolean> (default=false)** Ignore extended file attributes.
- verbose <boolean> (default=false)** Verbose output.

`pxar extract <archive> [<target>] [OPTIONS]`

Extract an archive.

<archive> <string> Archive name.

<target> <string> Target directory

Optional parameters:

- allow-existing-dirs <boolean> (default=false)** Allows directories to already exist on restore.
- files-from <string>** File containing match pattern for files to restore.
- no-acls <boolean> (default=false)** Ignore access control list entries.
- no-device-nodes <boolean> (default=false)** Ignore device nodes.
- no-fcaps <boolean> (default=false)** Ignore file capabilities.
- no-fifos <boolean> (default=false)** Ignore fifos.
- no-sockets <boolean> (default=false)** Ignore sockets.
- no-xattrs <boolean> (default=false)** Ignore extended file attributes.
- pattern <array>** List of paths or pattern matching files to restore
- strict <boolean> (default=false)** Stop on errors. Otherwise most errors will simply warn.
- verbose <boolean> (default=false)** Verbose output.

`pxar help [{<command>}] [OPTIONS]`

Get help about specified command (or sub-command).

<command> <array> Command. This may be a list in order to specify nested sub-commands.

Optional parameters:

--verbose <boolean> Verbose help.

`pxar list <archive> [OPTIONS]`

List the contents of an archive.

<archive> <string> Archive name.

Optional parameters:

--verbose <boolean> (default=false) Verbose output.

`pxar mount <archive> <mountpoint> [OPTIONS]`

Mount the archive to the provided mountpoint via FUSE.

<archive> <string> Archive name.


<mountpoint> <string> Mountpoint for the file system.

Optional parameters:

--verbose <boolean> (default=false) Verbose output, running in the foreground (for debugging).

FILE FORMATS

B.1 Proxmox File Archive Format (.pxar)



BACKUP PROTOCOL

C.1 Backup Protocol API

GET .

Directory index.

Returns: <null>

UPLOAD blob

Upload binary blob file.

Required properties:

encoded-size <integer> (12 - 16777260) Encoded blob size.

file-name <string> Backup archive name.

Returns: <null>

UPLOAD dynamic_chunk

Upload a new chunk.

Required properties:

digest <string> Chunk digest (SHA256).

encoded-size <integer> (13 - 16777260) Encoded chunk size.

size <integer> (1 - 16777216) Chunk size.

wid <integer> (1 - 256) Dynamic writer ID.

Returns: <null>

POST dynamic_close

Close dynamic index writer.

Required properties:

chunk-count <integer> (1 - N) Chunk count. This is used to verify that the server got all chunks.

csum <string> Digest list checksum.

size <integer> (1 - N) File size. This is used to verify that the server got all data.

wid <integer> (1 - 256) Dynamic writer ID.

Returns: <null>

POST dynamic_index

Create dynamic chunk index file.

Required properties:

archive-name <string> Backup archive name.

Returns: <null>

PUT dynamic_index

Append chunk to dynamic index writer.

Required properties:

digest-list <array> Chunk digest list.

offset-list <array> Chunk offset list.

wid <integer> (1 - 256) Dynamic writer ID.

Returns: <null>

POST finish

Mark backup as finished.

Returns: <null>

UPLOAD fixed_chunk

Upload a new chunk.

Required properties:

digest <string> Chunk digest (SHA256).

encoded-size <integer> (13 - 16777260) Encoded chunk size.

size <integer> (1 - 16777216) Chunk size.

wid <integer> (1 - 256) Fixed writer ID.

Returns: <null>

POST fixed_close

Close fixed index writer.

Required properties:

chunk-count <integer> (0 - N) Chunk count. This is used to verify that the server got all chunks. Ignored for incremental backups.

csum <string> Digest list checksum.

size <integer> (0 - N) File size. This is used to verify that the server got all data. Ignored for incremental backups.

wid <integer> (1 - 256) Fixed writer ID.

Returns: <null>

POST fixed_index

Create fixed chunk index file.

Required properties:

archive-name <string> Backup archive name.

size <integer> (1 - N) File size.

Optional properties:

reuse-csum <string> If set, compare last backup's csum and reuse index for incremental backup if it matches.

Returns: <null>

PUT fixed_index

Append chunk to fixed index writer.

Required properties:

digest-list <array> Chunk digest list.

offset-list <array> Chunk offset list.

wid <integer> (1 - 256) Fixed writer ID.

Returns: <null>

DOWNLOAD previous

Download archive from previous backup.

Required properties:

archive-name <string> Backup archive name.

Returns: <null>

UPLOAD speedtest

Test upload speed.

Returns: <null>

C.2 Reader Protocol API

GET .

Directory index.

Returns: <null>

UPLOAD blob

Upload binary blob file.

Required properties:

encoded-size <integer> (12 - 16777260) Encoded blob size.

file-name <string> Backup archive name.

Returns: <null>

UPLOAD dynamic_chunk

Upload a new chunk.

Required properties:

digest <string> Chunk digest (SHA256).

encoded-size <integer> (13 - 16777260) Encoded chunk size.

size <integer> (1 - 16777216) Chunk size.

wid <integer> (1 - 256) Dynamic writer ID.

Returns: <null>

POST dynamic_close

Close dynamic index writer.

Required properties:

chunk-count <integer> (1 - N) Chunk count. This is used to verify that the server got all chunks.

csum <string> Digest list checksum.

size <integer> (1 - N) File size. This is used to verify that the server got all data.

wid <integer> (1 - 256) Dynamic writer ID.

Returns: <null>

POST dynamic_index

Create dynamic chunk index file.

Required properties:

archive-name <string> Backup archive name.

Returns: <null>

PUT dynamic_index

Append chunk to dynamic index writer.

Required properties:

digest-list <array> Chunk digest list.

offset-list <array> Chunk offset list.

wid <integer> (1 - 256) Dynamic writer ID.

Returns: <null>

POST finish

Mark backup as finished.

Returns: <null>

UPLOAD fixed_chunk

Upload a new chunk.

Required properties:

- digest** <string> Chunk digest (SHA256).
- encoded-size** <integer> (13 - 16777260) Encoded chunk size.
- size** <integer> (1 - 16777216) Chunk size.
- wid** <integer> (1 - 256) Fixed writer ID.

Returns: <null>

POST fixed_close

Close fixed index writer.

Required properties:

- chunk-count** <integer> (0 - N) Chunk count. This is used to verify that the server got all chunks. Ignored for incremental backups.
- csum** <string> Digest list checksum.
- size** <integer> (0 - N) File size. This is used to verify that the server got all data. Ignored for incremental backups.
- wid** <integer> (1 - 256) Fixed writer ID.

Returns: <null>

POST fixed_index

Create fixed chunk index file.

Required properties:

- archive-name** <string> Backup archive name.
- size** <integer> (1 - N) File size.

Optional properties:

- reuse-csum** <string> If set, compare last backup's csum and reuse index for incremental backup if it matches.

Returns: <null>

PUT fixed_index

Append chunk to fixed index writer.

Required properties:

- digest-list** <array> Chunk digest list.
- offset-list** <array> Chunk offset list.
- wid** <integer> (1 - 256) Fixed writer ID.

Returns: **<null>**

DOWNLOAD previous

Download archive from previous backup.

Required properties:

archive-name <string> Backup archive name.

Returns: **<null>**

UPLOAD speedtest

Test upload speed.

Returns: **<null>**

CALENDAR EVENTS

D.1 Introduction and Format

Certain tasks, for example pruning and garbage collection, need to be performed on a regular basis. Proxmox Backup Server uses a format inspired by the systemd Time and Date Specification (see [systemd.time manpage](#)) called *calendar events* for its schedules.

Calendar events are expressions to specify one or more points in time. They are mostly compatible with systemd's calendar events.

The general format is as follows:

Listing 1: Calendar event

```
[WEEKDAY] [[YEARS-]MONTHS-DAYS] [HOURS:MINUTES[:SECONDS]]
```

Note that there either has to be at least a weekday, date or time part. If the weekday or date part is omitted, all (week)days are included. If the time part is omitted, the time 00:00:00 is implied. (e.g. '2020-01-01' refers to '2020-01-01 00:00:00')

Weekdays are specified with the abbreviated English version: *mon, tue, wed, thu, fri, sat, sun*.

Each field can contain multiple values in the following formats:

- comma-separated: e.g., 01,02,03
- as a range: e.g., 01..10
- as a repetition: e.g., 05/10 (means starting at 5 every 10)
- and a combination of the above: e.g., 01,05..10,12/02
- or a * for every possible value: e.g., *:00

There are some special values that have specific meaning:

Value	Syntax
<i>minutely</i>	*-*-* *:00
<i>hourly</i>	*-*-* *:00:00
<i>daily</i>	*-*-* 00:00:00
<i>weekly</i>	mon *-*-* 00:00:00
<i>monthly</i>	*-*01 00:00:00
<i>yearly or annually</i>	*-01-01 00:00:00
<i>quarterly</i>	*-01,04,07,10-01 00:00:00
<i>semiannually or semi-annually</i>	*-01,07-01 00:00:00

Here is a table with some useful examples:

Example	Alternative	Explanation
<i>mon,tue,wed,thu,fri</i>	<i>mon..fri</i>	Every working day at 00:00
<i>sat,sun</i>	<i>sat..sun</i>	Only on weekends at 00:00
<i>mon,wed,fri</i>	–	Monday, Wednesday, Friday at 00:00
<i>12:05</i>	–	Every day at 12:05 PM
<i>*:00/5</i>	<i>0/1:0/5</i>	Every five minutes
<i>mon..wed *:30/10</i>	<i>mon,tue,wed *:30/10</i>	Monday, Tuesday, Wednesday 30, 40 and 50 minutes after every full hour
<i>mon..fri 8..17,22:0/15</i>	–	Every working day every 15 minutes between 8 AM and 6 PM and between 10 PM and 11 PM
<i>fri 12..13:5/20</i>	<i>fri 12,13:5/20</i>	Friday at 12:05, 12:25, 12:45, 13:05, 13:25 and 13:45
<i>12,14,16,18,20,22:0/2:5</i>	<i>12,14,16,18,20,22:0/2:5</i>	Every day starting at 12:05 until 22:05, every 2 hours
<i>*.*</i>	<i>0/1:0/1</i>	Every minute (minimum interval)
<i>*-05</i>	–	On the 5th day of every Month
<i>Sat *-1..7 15:00</i>	–	First Saturday each Month at 15:00
<i>2015-10-21</i>	–	21st October 2015 at 00:00

D.2 Differences to systemd

Not all features of systemd calendar events are implemented:

- no Unix timestamps (e.g. @12345): instead use date and time to specify a specific point in time
- no timezone: all schedules use the set timezone on the server
- no sub-second resolution
- no reverse day syntax (e.g. 2020-03~01)
- no repetition of ranges (e.g. 1..10/2)

D.3 Notes on scheduling

In [Proxmox Backup](#) scheduling for most tasks is done in the *proxmox-backup-proxy*. This daemon checks all job schedules if they are due every minute. This means that even if *calendar events* can contain seconds, it will only be checked once a minute.

Also, all schedules will be checked against the timezone set in the [Proxmox Backup](#) server.

GLOSSARY

Virtual machine A virtual machine is a program that can execute an entire operating system inside an emulated hardware environment.

Container A container is an isolated user space. Programs run directly on the host's kernel, but with limited access to the host resources.

Datastore A place to store backups. A directory which contains the backup data. The current implementation is file-system based.

Rust Rust is a new, fast and memory-efficient system programming language. It has no runtime or garbage collector. Rust's rich type system and ownership model guarantee memory-safety and thread-safety. I can eliminate many classes of bugs at compile-time.

Sphinx Is a tool that makes it easy to create intelligent and beautiful documentation. It was originally created for the documentation of the Python programming language. It has excellent facilities for the documentation of software projects in a range of languages.

reStructuredText Is an easy-to-read, what-you-see-is-what-you-get plaintext markup syntax and parser system.

FUSE Filesystem in Userspace (**FUSE**) defines an interface which makes it possible to implement a filesystem in userspace as opposed to implementing it in the kernel. The fuse kernel driver handles filesystem requests and sends them to a userspace application.

Remote A remote Proxmox Backup Server installation and credentials for a user on it. You can pull datastores from a remote to a local datastore in order to have redundant backups.

GNU FREE DOCUMENTATION LICENSE

Version 1.3, 3 November 2008

Copyright (C) 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc. <<https://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”.) To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document’s license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.

- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled “Endorsements”. Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled “Endorsements” or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version’s license notice. These titles must be distinct from any other section titles.

You may add a section Entitled “Endorsements”, provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgements”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements”.

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the

copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <https://www.gnu.org/licenses/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

"Massive Multiauthor Collaboration Site" (or "MMC Site") means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those

works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

INDEX

C

Container, [89](#)

D

Datastore, [89](#)

F

FUSE, [89](#)

R

Remote, [89](#)

reStructuredText, [89](#)

Rust, [89](#)

S

Sphinx, [89](#)

V

Virtual machine, [89](#)