

ARCHIWARE

Reference Manual 8.0



P5 Archive



P5 Backup



P5 Synchronize

Table of Contents

1 Preface.....	6
1.1 About This Manual.....	6
1.2 Who Should Read This Document?.....	6
1.3 Typographic Legend.....	7
2 Introduction.....	8
2.1 Backing Up.....	8
2.2 Synchronization.....	9
2.3 Archiving.....	10
2.4 Configuration.....	10
2.5 Program Description.....	13
2.6 Inside P5.....	13
2.6.1 Data Backup.....	13
2.6.2 Restoring Data.....	14
2.7 Software Licensing.....	14
2.8 Registration.....	15
2.9 Technical Support.....	15
2.10 Documentation and Help.....	15
3 Installation.....	16
3.1 Overview.....	16
3.2 System Requirements.....	17
3.3 System Resources.....	19
3.4 Installing on macOS.....	20
3.5 Installing on Windows.....	23
3.6 Linux /FreeBSD Installations.....	26
3.6.1 Installation Requirements.....	26
3.6.2 Creating an Installation Directory.....	27
3.6.3 Copying the Installation File.....	27
3.6.4 Unpacking the Program Directory Structure.....	27
3.6.5 Starting the server.....	27
3.7 Installing on Synology NAS.....	28
3.7.1 Log in to Your Synology NAS.....	28
3.7.2 Edit <i>Trust Level</i> /Settings.....	29
3.7.3 Upload the SPK Package File.....	29
3.7.4 Installation Process.....	30
3.7.5 Accessing the <i>Archiware P5 Suite</i>	32
3.8 Installing on QNAP NAS.....	33
3.9 Installing on Netgear ReadyNAS.....	34
3.10 Installing on iXsystems TrueNAS.....	35
3.11 Testing the Installation.....	37
3.12 Alternate Installations.....	37
3.13 Firewall Configuration Requirements.....	38
3.14 Upgrading to P5.....	39
3.14.1 Hints Regarding the Upgrade.....	39
3.14.2 Program Update.....	40
4 Signing In and Entering the License.....	41

- 4.1 Starting the Browser.....41
- 4.2 License Activation.....42
- 4.3 A General Introduction to the Server's User Interface.....43
 - 4.3.1 The Installation Overview.....46
- 4.4 Using the P5 Agent Setup GUI.....47
- 5 First Steps With P5.....48**
 - 5.1 First Steps to Set Up the Backup.....48
 - 5.2 First Steps to Set Up an Archive Plan.....49
 - 5.3 First Steps to Set Up a Synchronization Plan.....50
- 6 Configuring the P5 Server.....51**
 - 6.1 Licenses.....51
 - 6.1.1 Creating a License.....52
 - 6.1.2 Importing a License.....53
 - 6.2 Storage Manager.....54
 - 6.2.1 Tape Libraries / Media Changers.....55
 - 6.2.2 Single Tape Drives.....60
 - 6.2.3 Container Storage.....65
 - 6.2.4 Cloud Services.....68
 - 6.2.5 Creating a Virtual Disk Library.....71
 - 6.2.6 Creating a Removable Disk Storage.....73
 - 1.1.1 Storage Manager: Hamburger Menu.....76
 - 6.3 The Media Pool.....79
 - 6.3.1 Creating a Pool.....79
 - 6.3.2 Tape/Disk Pool : Additional Options.....81
 - 6.3.3 Tape or VTL Pool: Concurrency Options.....83
 - 6.3.4 Container Pool: Expert Settings.....85
 - 6.3.5 Pool Usage.....86
 - 6.3.6 Tape Based Cloud Pool: LTO Encryption.....87
 - 6.4 Volumes.....88
 - 6.4.1 Creating a Volume (Labeling).....88
 - 6.4.2 Modifying a Volume.....88
 - 6.4.3 Volume Manager: Gear Menu.....92
 - 6.5 Archive Plans.....97
 - 6.5.1 Creating an Archive Plan.....97
 - 6.5.2 Archive Overview.....111
 - 6.5.3 Archive to LTFS Tapes.....112
 - 6.6 Backup Plans.....113
 - 6.6.1 Creating a Backup Plan.....113
 - 6.6.2 Backup Overview.....121
 - 6.7 Host Sync Plans.....123
 - 6.7.1 Creating a Sync Plan.....123
 - 6.7.2 Synchronize Overview.....134
 - 6.7.3 Manual Synchronize.....136
 - 6.8 Cloud Sync Plans.....137
 - 6.8.1 Creating a Cloud Sync Plan.....137
 - 6.8.2 Cloud Sync Options.....139
 - 6.8.3 Cloud Sync Expert Settings.....140

6.9 Filters.....	142
6.9.1 Creating a Filter.....	142
6.10 Logs.....	145
6.10.1 The Job Monitor.....	145
6.10.2 The Log History.....	147
6.10.3 the Job Log.....	148
6.11 P5 Clients.....	149
6.11.1 Creating a Client.....	149
6.11.2 Client: Interactive Functions.....	153
6.11.3 Setting up Encryption for a Client.....	154
6.11.4 Setting up Reverse Connection for a Client.....	155
6.12 Index Databases.....	157
6.12.1 Creating a Database.....	157
6.12.2 Archive Database Options.....	159
6.12.3 Archive Index Database Backup and Cleanup.....	160
6.12.4 Archive Job Log Inventory.....	160
6.12.5 Database: Functions.....	161
6.13 Index Access Control (Login Areas).....	163
6.13.1 Creating a Log-In Area.....	163
6.14 User Preferences.....	165
6.15 Group Preferences.....	171
6.16 S3 Object Archive.....	173
6.16.1 Creating an Archive Bucket.....	173
6.16.2 Archive Bucket Options.....	175
6.16.3 S3 API Capabilities.....	175
6.17 LTO Encryption Key Management.....	176
6.17.1 Creating an Encryption Key.....	176
6.18 Advanced Server Settings.....	178
6.19 Lifecycle Management.....	179
6.19.1 Migration Plan.....	179
6.19.2 Verify Plan.....	186
6.19.3 Utilization.....	189
7 LTFS Exchange Functions.....	192
7.1 Mount / Unmount.....	193
7.2 Label.....	193
7.3 Read from LTFS.....	194
7.4 Write to LTFS.....	195
8 Restoring Data From Backup.....	196
8.1 Restoring Files.....	197
8.2 The Restore Selection.....	198
8.3 The Restore To Button.....	199
8.4 The Index Options.....	200
8.5 File Versions.....	201
8.6 The Search Button.....	202
9 Archiving and Restoring from Archive.....	205
9.1 Archiving Files.....	205
9.2 Restoring from Archives.....	207

- 9.2.1 The Restore-To Button.....209
- 9.2.2 The Restore Selection.....211
- 9.2.3 File Versions.....212
- 9.2.4 The Search Button.....212
- 9.2.5 Index Options.....216
- 9.2.6 Restore from a File List.....217
- 9.2.7 Editing the Archive Index.....218
- 10 Operating Procedures.....219**
 - 10.1 Labeling Media.....219
 - 10.2 Mounting Volumes.....220
 - 10.2.1 Mounting or Unmounting a Volume.....220
 - 10.3 Replacing Volumes in a Library.....221
 - 10.3.1 Jukebox Inventory.....221
 - 10.4 Handling Media Pools.....223
 - 10.5 Handling Network Shares on Windows.....224
 - 10.5.1 Mount a Network Share and Assign it to a Drive Letter.....224
 - 10.5.2 Mount a Network Share and Use the URL Directly.....224
- 11 Customizing.....225**
 - 11.1 Users and Groups.....225
 - 11.2 Index and File System Access.....225
 - 11.3 Customizing Reporting.....226
 - 11.4 Customized Views.....227
 - 11.5 The Dashboard.....229
 - 11.6 Setting up a HTTPS Certificate.....231
 - 11.7 Setting up Login with Two Factor Authentication.....232
- 12 Troubleshooting.....233**
 - 12.1 Downloading Support Data.....233
 - 12.2 Restoring the Configuration.....234
 - 12.2.1 Handling Encryption during Configuration Recovery.....236
 - 12.2.2 Disaster Recovery.....237
 - 12.3 Fixing Tape Drive Problems.....238
 - 12.3.1 Tape Access Problems.....238
 - 12.3.2 Tape Operation Problems.....240
 - 12.4 Network Connection Problems.....241
 - 12.4.1 The Network Connection Cannot be Established at All.....241
 - 12.4.2 Bad Connection / Connection Breakdown.....242
 - 12.4.3 Network Ports and Protocols used by P5.....242
- 13 Best Practice.....243**
 - 13.1 Prerequisite for Backup and Archive.....243
 - 13.2 Backup Scenarios.....244
 - 13.2.1 Standard Backup for Small Business.....244
 - 13.2.2 Progressive Backup for Huge Amounts of Data.....244
 - 13.2.3 Backup with Alternating Pools.....245
 - 13.2.4 Archiving Methods.....246
 - 13.2.5 Using Synchronize as Backup to Disk.....247
 - 13.2.6 Backup to Disk to Tape.....248
 - 13.2.7 Backup to Removable Disk Media.....248

1 Preface

1.1 About This Manual

This document describes the installation and configuration procedures for the Archiware P5 package. You will also find examples and case studies for installing, managing and employing P5 in a typical production environment.

1.2 Who Should Read This Document?

The information contained in this manual is intended for system administrators who will install and maintain P5 as well as for users who will employ P5 as a data management system.

Information for the system administrator:

- installing P5 on the server
- using the P5 data browser
- configuring a backup or an archive job
- customizing
- troubleshooting

Information is also included about the necessary periodic administrative tasks. These include:

- checking the daily backup
- managing volumes and devices

The user can read about how to:

- restore data
- archive data
- upload and download data
- and study some case scenarios

1.3 Typographic Legend

The following typography and symbols are used in this document for better clarity and simplicity:


Font or Symbol	Description and Example
bold	Emphasized text. Example: There must be a sufficient number of labeled volumes with enough free space for a backup to complete successfully.
<i>italic, fixed font</i>	Names of directories, files, machines, commands, glossary entries, as well as new terms defined in each chapter. Example: You will find errors logged in the P5 Server log file, which is located in the file <i>log/lexsrv.log</i> in the directory where you installed P5.
fixed font	Commands and text to be entered verbatim. Example: The command to start the P5 Server <code>start-server</code> .
<i>Italic, bold</i>	Information displayed on the display, window titles, dialog fields and other GUI elements. Example: You may find the names of the volumes required for the restore operation listed in the <i>Volume Listing</i> .
<i>italic, bold, fixed</i>	Shows the sequence of menu or application steps to be taken to initiate a command or a program. Example: To enter the backup administration area, click Login Area > Administration > Backup Management.
	Important information, notes or warnings about possible problem or error sources. Example: To ensure a sufficient level of data safety, it is reasonable to make a full backup at least once a week in addition to incremental backups.

Table of typographic conventions

2 Introduction

2.1 Backing Up

Data has to be backed up on regular basis to prevent or minimize the danger of data loss or corruption. When recovering the data, it is important to be able to restore the most recent state. All the files created on the diverse macOS and PC workstations in the network are saved on a central (for example Unix) file server. It must be possible to back up the data from such high-end servers in the shortest possible time.

There is a limited time window, generally during the night, within which such huge data volumes have to be backed up. This must be done automatically and without human intervention. P5 is equipped with many features which ensure a high standard of data safety in a diverse and complex production environment.

Depending on the configuration and demand, multiple media drives will be used simultaneously to perform a backup. This enables large amounts of data to be backed up safely and timely within the given time frame. For a completely autonomous backup, storage libraries (media changers or jukeboxes) are employed, which P5 controls and maintains. This way, neither a backup nor a restore operation requires manual media change.

Care has been taken to ensure that all the required operations in conjunction with the underlying hardware require a minimal amount of configuration, monitoring and maintenance. P5 logs all of its operational steps and activities and e-mails them on demand.

The backed up data is indexed into a file structure that can be browsed and queried, and by means of which the restore is later controlled. Versions of data files are automatically kept, so that older versions of files can be easily retrieved. P5 maintains the index files transparently and automatically. The index files are themselves automatically backed up and superfluous entries removed.

There are three different backup modes:

- **Full Backup**
All the files and folders contained within the given list of folders will be backed up. For each full backup, P5 creates a new file system snapshot in the index. Using the index, the user can locate his data and issue a restore.
- **Incremental Backup**
Only the files that changed since the last full or incremental backup will be saved. P5 will not create a new index, but will amend the existing index with the detected changes. Other backup systems use the modification date of the files to detect the changes and simply add the new data information into the index. P5, on the other hand, references its index and compares the state of the file system for each incremental backup to decide which data has changed and needs to be saved. This enables it to also detect moved files whose modification date did not change, and save them from the new place in the file system. Furthermore, intentionally deleted files will also be detected and marked in the index, thereby maintaining the exact picture of the file system at any particular point in time.

Additionally, the incremental backup takes care of continuously maintaining the backed up data, retaining the most recent state of the file system and cleaning up the index. It will prune the old data off old tapes and out of the index and recycle the tapes automatically, retaining only the data that falls within the configured retention window.

Nevertheless, a full snapshot of the current file system state is always maintained. This unique feature allows for running an everlasting incremental backup, thus eliminating the need for running a full backup – a feature essential for ensuring that large file systems are kept backed up at all times. Furthermore, it limits the backup to a controlled number of tapes and the index to a controlled size. Without it, it would be virtually impossible to make a full backup of such file systems within any reasonable time.

- **Synthetic Full Backup**

This feature enables creating a full backup from the most recent backed up state of the file system using previous backups, regardless on which media they may be, simply by referencing the index.

Synthetic Full Backup does not require data from the servers or the workstations. It employs the already backed up data from the existing backup media. Utilizing this data, the corresponding index and a new set of media, P5 creates a new full backup without penalizing the system's resources, imposing load on the network or impacting the day-to-day business on the machines. In this process, the P5 media management autonomously identifies and handles both the media to read from, as well as the new media to be written to, provided these are available within a library/media-changing device.

- **Progressive Backup**

- Traditionally, after some period of time running only Incremental backups, it becomes necessary to run a new full backup. This adds a new copy of all the data being saved to the storage pool, allowing older volumes within the pool to be recycled and reused to future backups. However, running regular full backups can be both time consuming and inconvenient. P5's Progressive backup handles the re-saving of data to the storage pool in smaller amounts, over several runs of incremental backups. This allows the task of renewing the data in the pool to be spread out over several days in smaller backup time windows.

2.2 Synchronization

Synchronization is a fast and effective form of backup. During synchronization, all the files of a directory or of a file system are copied (replicated) to another directory and kept synchronized thereafter. The source and the target directories can reside on different machines, one running a P5 server, the other a P5 client, or vice versa.

There are several advantages to using synchronization compared to a classical backup on tapes:

- a copy of the data is available directly on-line.
- you don't need to restore the data from tapes, as in a regular backup.
- the data is immediately available.

- the data can be made available network-wide on any remote server.

Furthermore, synchronizing can be employed as a work-flow element:

- to connect all the subsidiaries Internet-wide with the headquarters.
- to involve external service providers into the work flow of the company.

2.3 Archiving

P5 is also a production archiving and data-management tool. Thereby, data is copied (or moved) from disks or disk arrays to external media in order to better utilize the limited on-line disk space.

Projects are commonly organized and kept in folders or folder structures on the file server where, in the process of production, all the files relating to the project are stored. Upon project completion, you may archive the corresponding data and make the space available for new assignments. The archived data remains available to the production and can be re-integrated into new production on demand.

2.4 Configuration

The configuration entities in P5, such as devices (drives or changers), volumes, or clients are called resources. These are stored and kept in the resource (or configuration) database. The database has a tree-like structure that you can easily navigate.

Features are organized by functions. For instance, if you are looking for a drive, you will navigate into the devices area, where you can find all the installed drives. Likewise, all the labeled tapes (or disk containers) may be found in the volumes area. Each entry of a group has a unique ID, description and several additional parameters specific to that resource. These parameters are called the attributes. The resources may be linked or refer to each other. For instance, a backup plan refers to a media pool, or a volume is linked to its clone.

Using the database, P5 coordinates access to the various resources between its diverse modules, for example to restore different file off multiple volumes, where jobs were initiated from different users at the same time.

P5's browser-based user interface enables configuring of the resources. Most resources come with default or predefined values to simplify the setup and initial configuration. By changing these default values you can tune P5 to optimize the performance within a particular configuration and environment.

The table on next page lists all the resource groups and explains their purpose.

Resource Groups

Resource Group	Purpose and Use	Links
Client	Each server, which is to be backed up, must be configured as a client. The client resource carries all the necessary attributes to identify the server in the network, for example, the TCP/IP address.	Archive Plan, Backup Plan, Sync Plan
Archive Plan	The archive plan stores all the important settings for executing the plan. A user initiating an archive job can choose from a list of predefined and preconfigured plans, thus not having to get involved with all the technical details necessary for its execution.	Client, Filter
Backup Plan	Each backup plan defines one or more directories to back up. Additional attributes define the time the backup should be run and which pool to use to write the data to. Filters allow to exclude certain files from the backup, for example temporary files.	Client, Filter
Sync Plan	Each sync plan defines one or more sources and a target directory on a chosen P5 client. Additional options allow you to control, among other features, the synchronization method and the execution time. Filters allow to exclude certain files.	Client, Filter
Filter	Filters carry attributes for selecting files with desired properties, for example based on name or type, to be included or excluded in the course of executing a job.	Archive Plan, Backup Plan, Sync Plan
Volume	P5 stores data on volumes. Volumes carry a label. Each volume (labeled media, for example an LTO tape) has an entry in this resource group. Volumes are grouped and belong to pools. Volumes are labeled using the appropriate drive, meaning that a label is physically written in order to uniquely identify them.	Pool
Media Pool	Each volume (labeled media) belongs to a pool. Pools are an administrative and a control element. A job may be directed by means of a pool to use and reserve (limit the use of) dedicated resources.	Archive Plan, Backup Plan, Device, Volume
Cloud Service	A cloud service defines access to a cloud storage. It can be addressed directly by Archive Plans and can be referred to in Media pools to move or copy the data to a cloud storage.	Archive Plan, Media Pool, Volumes
Device	Entries in this resource group are associated with drives or other storage elements which P5 supports. The drives may be stand-alone or a part of a storage library. Attributes define the drive's properties such as the drive type and the operating system path to this device (its handle or name).	Jukebox, Pool

Jukebox	Each library (or changer) attached to the server and configured in P5 has an entry in this resource group. The attributes describe the hardware properties of the changer, for example the number of slots and their addresses, or the type and number of built-in drives.	Device
Job	All the tasks, such as labeling media or executing a backup, lie under the control of the job scheduler and are executed as jobs. A job exists as a resource only when it is scheduled for execution or being executed. Upon completion, the resource of the job is archived as a log entry.	
Database	P5 manages the configuration and the backup indexes as databases.	Log-in Area
Log-in Area	Log-in areas give the user structured access to the indexes and to the database.	Database
User Preferences	User configurations and preferences are managed through the user preferences resource. Using the attributes, users may be granted use of specific operations or access to diverse areas in P5.	
Group-Preferences	Group preferences are managed through the group preferences resource. Group members may be granted access to specific folders, also access to log-in areas and archive plans may be limited to specific groups.	Archive Plan, Log-in Area
Object Archive	Object Archives provide an S3 Object Storage service for Archiving Data to media managed by P5.	

Table of Resource Groups

2.5 Program Description

P5 is a pure server application. All the programs and services are executed on the server where it is installed. When running, P5 initiates one or multiple `nsd` processes which handle all the given tasks, for example executing a backup job or changing the attributes of a resource.

To check if the processes are running, you can issue the Unix `ps` command:

For macOS:

```
# ps -aux | grep nsd
root 3830    10 10:42:33 pts/12 0:00 /usr/local/aw/bin/nsd -t conf/lexsrv.80
```

For Linux:

```
# ps -ef | grep nsd
root 3830    10 10:42:33 pts/12 0:00 /usr/local/aw/bin/nsd -t conf/lexsrv.80
root 38333 8310 10:42:33 ?      0:09 /usr/local/aw/bin/nsd -t conf/lexsrv.80
```

The `nsd` application server is the P5 software platform and thus its kernel. It executes and controls all of its functions. P5 consists of a series of additional modules which are run and controlled by the application server.

With its integrated Web server, P5 provides the entire user interface through web browser, allowing access from any connected workstation and eliminating the need for installation of additional client applications.

2.6 Inside P5

P5 is a program for saving data. The data is contained in files and directories that are located in file systems on on-line storage such as disks or raid systems.

The data is saved to backup media via drives installed on the system. Diverse media types and corresponding drives are supported. To store large amounts of data, media libraries are also supported. In this case P5 controls the libraries (also called changers or jukeboxes) and automatically manages the media, allowing for large volumes of data to be processed without the need for human intervention.

2.6.1 Data Backup

The backup of data is managed by configuring backup plans. Backup plans are configured using the P5 data browser and carry attributes necessary for executing a backup. The most important parameters are:

- one or more clients where the data to be backed up is located.
- for each client, one or more paths in the file system to directories to be backed up.
- the periodic cycle to follow when performing the backup.
- the level of the backup.
- the time when the backup should be initiated.

The job scheduler starts a job automatically at a time configured for each of the backup plans. The job manager takes over the supervision of the job and takes into consideration all the configured attributes. It requests

volumes from the media manager for executing the backup. The media manager allocates one or more devices, depending on the available resources, or waits for devices to become available. For each device, it instantiates a job executor. At this moment the actual backup begins. If one of the volumes becomes filled, the media manager chooses the next available volume and places it in the drive. When the physical backup is completed, the job scheduler creates an index of the saved data. This index is used for the recovering of data.

2.6.2 Restoring Data

The user initiates a data restore. Using the P5 data browser, he/she marks a file or directory to recover. From the recover context menu, further options controlling the recovery are available:

- the user can choose a new client and/or path where to store the data,
- and he can define what to do if files with the same names are found on the file system.

After setting these options, P5 assembles a list of files and informs the user which volumes are necessary to complete its operation. If more than one volume is required to complete the operation, the job executor contracts the media manager to change the volume. The operator has to change the volumes manually upon request in case that no changer is available. The files and folders will be recovered identically to the original.

2.7 Software Licensing

Depending on the required performance and functionality, there are four basic license types available for P5:

- **P5 Backup**
The P5 Backup license enables you to backup and restore data from a single client (a server) using a single drive, based on time driven backup plans. For each additional machine (P5 client) you wish to back up, you will require an additional server or workstation agent license. For each additional drive, you will require a media drive license. If you wish to employ storage libraries, you will require a media changer license, which enables you to use any number of jukeboxes simultaneously.
- **P5 Backup2Go**
The P5 Backup2Go license enables you to backup and restore data from several workstations onto one Backup2Go server. An additional workstation agent license is required for each further workstation that is to be backup up to the Backup2Go server.
- **P5 Archive**
The P5 Archive license enables you to archive and restore data from a single client (a server) using a single drive, based on archive plans you define. For each additional machine (P5 client) from which you wish to archive files, you will require an additional server or workstation agent license. For each additional drive, you will require a media drive license. If you wish to employ storage libraries, you will require a media changer license, which enables you to use any number of jukeboxes simultaneously.
- **P5 Synchronize**
This license allows you to copy one or more directories from or to a client (a server) based on a time-driven synchronize plan defined on the server. An additional server or workstation agent license is required for each additional host (P5 client) from which or to which data is to be copied or moved.
- **P5 Trial**
In the trial version, you can test P5 for 30 days in its full, unrestricted functionality. Technical support is limited to the support offered through the [Internet](#), free of charge.

2.8 Registration

You can register your P5 software on-line on Archiware's web site. Upon correctly completing the form, you will receive an activation code. You have to create a license and enter this activation code to enable the functionality that it provides.

You may find further information regarding licensing in chapter 4.2, License Activation.

2.9 Technical Support

Contact information, software updates, technical documentation and information about the current support programs and options may be found on Archiware's web site.

2.10 Documentation and Help

P5 comes with a detailed documentation. This documentation is mainly intended for on-line use. This way, the user can have the documentation available while working with P5. It is, of course, possible to print out the documentation or selected parts of it.

This documentation can be found at the following locations:

- on the P5 server, after starting the software and logging in as a user. The 'P5' menu takes you to the documentation.
- on [Archiware's](#) web site.

3 Installation

This section describes how to install P5 on macOS, Windows and Linux.



These installation notes are intended for system administrators with a solid knowledge and understanding of the Unix commands and must be executed by the local system administrator.

3.1 Overview

Installing P5 takes several steps. It is recommended to read through this section before attempting to install the software.

1. Have your installation file downloaded from the Internet at hand. You will require at least one P5 license.
If you wish to test the software in its full functionality, obtain a trial license.
2. Before starting the installation, please make sure that the system requirements (as described below) are met.
3. Depending on the license, P5 will involve diverse system resources, as described in this manual, that you need to consider.
4. Follow the steps as described in the installation sections below.

3.2 System Requirements

P5 is available for the following platforms and operating systems:

- **macOS:** Intel x86 64-Bit: OSX 10.9 – 15.x
Apple M1: 11.x – 15.x
- **Windows:** Server 2016, Server 2019, Server 2022, Server 2025
Windows 10, 11
- **Linux:** Intel x86 64-Bit systems with glibc version 2.15 are supported, that includes:
OpenSuSE 12.2+ / SLES 12+
CentOS 7+ / RHEL 7+ / Fedora 19+
Ubuntu 12+
Debian 8+
- **Synology NAS:** DSM operating system 6.2, 7+, Intel x86 64-Bit CPU
- **QNAP NAS:** QNAP NAS Systems QTS 4.3.0+, 5.0 +
QTS hero
QTScloud
Intel x86 64-Bit CPU
- **ReadyNAS:** OS 6.6.0 +, Intel x86 64-Bit CPU
- **FreeBSD:** Version 13 – 14, Intel x86 64-Bit CPU
- **TrueNAS scale:** (Installation through TrueNAS Scale GUI)
- **TrueNAS on FreeBSD:** Based on FreeBSD 12 or 13
- **Synology 7 or Linux Docker:** Requires installed Docker (Linux) or Container Manager (Synology 7)
Intel x86 64 Bit CPU, Synology DSM7+ or Linux OS.



On Linux, P5 requires that the Unicode UTF-8 Language package is installed, also the locale en_US.UTF-8 must be installed.

P5 is released for the following file server products:

- Helios UB, UB+, UB2 and UB64
- Xinet Fullpress 14.0+
- Netatalk 2
- Windows Services for Macintosh

P5 supports the following browsers:

- Safari 13+
- FireFox 70+
- Chrome 80+

P5 supports the following virtualizations:

- Parallels
- VMWare (x86)
- Linux-Xen
- Hyper-V

P5 supports addressing the following cloud storage:

- Amazon S3
- Amazon Snowball
- Alibaba Cloud
- Backblaze B2
- Generic S3
- Google Cloud Storage
- Huawei Cloud
- Microsoft Azure
- Seagate Lyve
- Wasabi Cloud Storage
- Wasabi Ball
- Amazon Glacier (Legacy)

3.3 System Resources

Depending on the license, P5 will involve diverse system resources. The following section describes all the system resources that P5 requires. However, the requirements on the CPU as a resource are not mentioned here since those vary widely with the system configuration. Please note that the system load can be considerably high during a backup or synchronize run, since they involve a lot of I/O read and write operations that require system calls to be executed.

1. Disk Space

To ensure an optimal backup, P5 needs sufficient disk space available for its use.



It is recommended to configure the program and index areas on backed up file systems, for example a mirrored disk or a RAID system. This has the advantage that a crash of the disk does not influence the backup or, more importantly, the restore.

During installation, an application and an index directory are created. The application directory contains all the programs, libraries and protocols. At least 250 MB of disk space is recommended. The installation script suggests a directory with sufficient disk space.

The index directory contains the configuration files and the backup indexes. Each file or directory that is backed up has a corresponding entry in the index, allowing a search to be performed later. The amount of disk space depends on the number of files to be backed up and the length of time for which the file entries are to be kept in the backup index. On average, 1 KB for each file or a copy of it is required. This number can vary depending on the length of the file names, directory structure and depth.

2. Backup Devices

P5 requires one or more drives attached to the system it is installed on. Any SCSI II or SCSI III conforming drives common to the corresponding operating system are supported. LTO, AIT, DLT and DAT tape drives have been tested and qualified. As an alternative, P5 can save data into container files on the local hard disk.

3. Media Libraries

A media library or changer is required where larger amounts of data have to be backed up in order to enable fully automatic runs. A wide range of devices is supported.

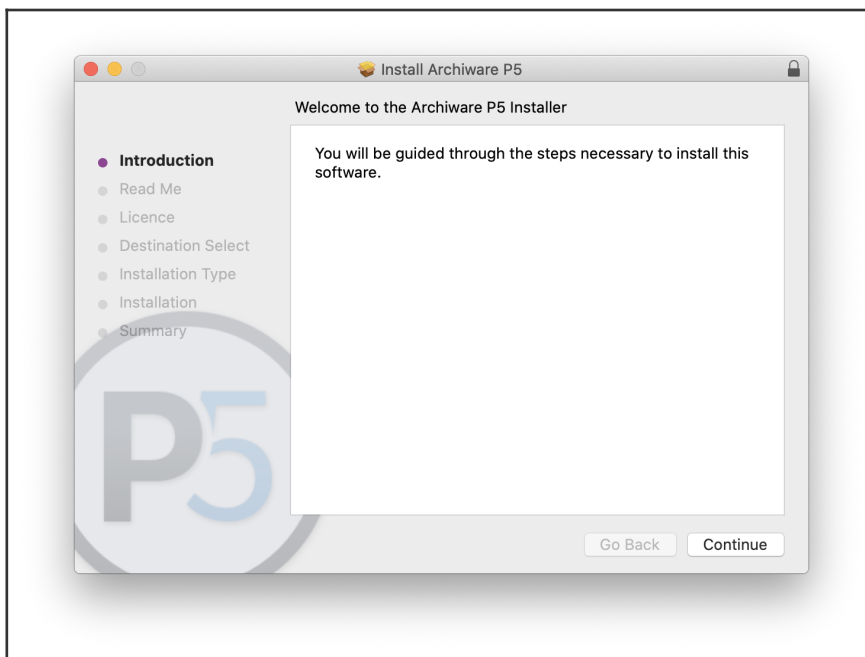
3.4 Installing on macOS

Archiware makes P5 available online. The installation file is packed and has the format 'awpstxyz.dmg'. The letters 'xyz' stand for the given P5 version. After mounting the disc image, the program 'awpstxyz.pkg' automatically starts the P5 installation.

In case you are upgrading from an earlier version of PresSTORE, please see section 3.14, Upgrading to P5.

The installation program automatically copies all the required modules into the P5 home directory: '/usr/local/aw'.

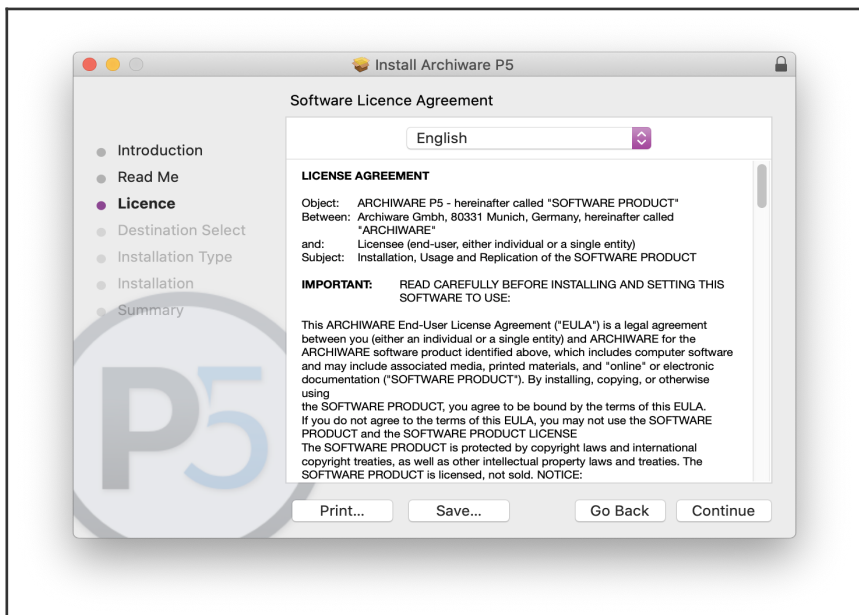
1. Begin installing on macOS. You will be asked for you user name and password. Please enter and confirm with 'OK'.



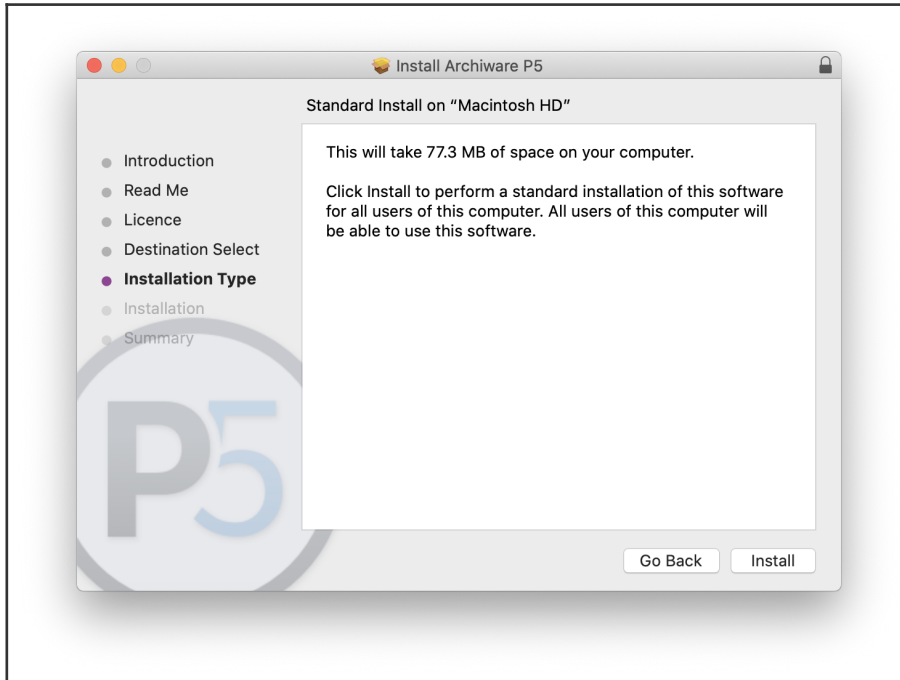
2. This window displays all the installation steps. Please click 'Continue' to proceed.



3. Please read through the License Agreement and familiarize yourself with the terms of use and of the application of P5. If you own a full or a time limited P5 license, click 'Continue' followed by 'Accept'.



4. Please choose the target volume where to install P5 and click **Continue**. Please note: A P5 server requires additional disk space for managing and storing the backup and archive indexes. See section 3.3, System Resources.



5. Following a successful P5 installation, you will continue to either configure the server using a browser or add the newly installed client to the client resources in the P5 server.



Starting with macOS version 10.14, additional steps must be taken to grant access to private files, i.e. Mail, Calendar, Photos ... as well network shares mounted by a user.

To do so:

- Navigate to System Preferences -> Security & Privacy -> Privacy -> Full Disk Access
- Open Apple System Preferences
- Open Security & Privacy
- Select the Privacy tab
- Click the lock icon in the lower-left corner
- In the left pane, select Full Disk Access
- In the right pane, click the "+" icon.
- In the Finder, select Go from the top menu bar, then Go to Folder..., type in `/usr/local/aw/bin` and hit Enter
- Locate the `nsd` binary and drag it into the Full Disk Access window.
- Close System Preferences
- Stop and restart P5.

3.5 Installing on Windows

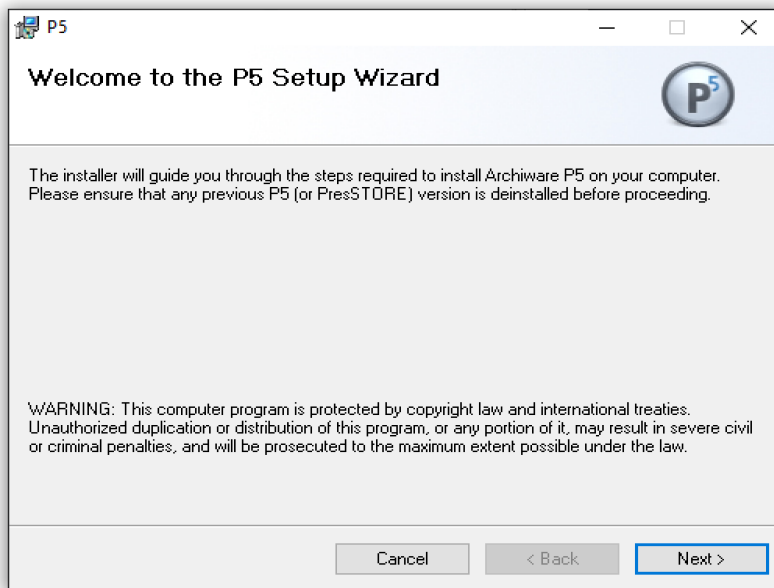
Archiware makes P5 available online. The installation file is packed and has the format 'awpstXYZwin.zip' . The letters 'xyz' stand for the given P5 version. When unpacking this file, you will find a P5-64.msi within. The installation is done by the Windows Installer by simply double-clicking on the msi installation file.

In case you are upgrading from an earlier version , please see section 3.14, Upgrading to P5.

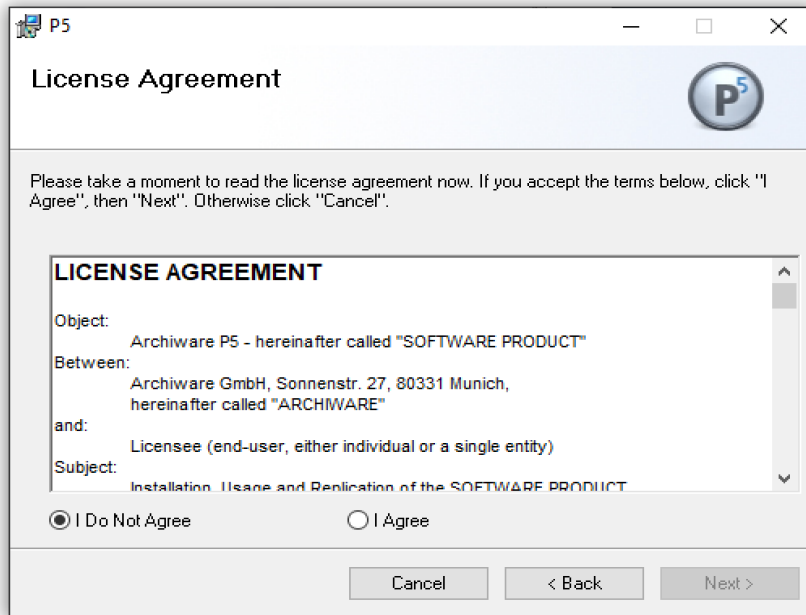


Administrator access rights are required to install P5.


1. After starting the installer, please click **Next** to continue.




2. Please read through the License Agreement and familiarize yourself with the terms of use and of the application of P5. Accept the License Agreement with **I agree** or exit the installation.



3. To accept the suggested installation path, click **Next**. To choose another path, click **Browse**.



4. If necessary, manually close the open command prompt window.
5. P5 is installed and started as a service. The default browser port is set to 8000. Subsequently, the P5 Server Manager is started.



A green icon indicates that P5 is running, a red icon that P5 could not be started. The log files stored in `log\lexsrv.log` and `log\install.log` in the P5 home folder contain additional information.

To continue with the configuration, right-click the P5 Server Manager icon in the task bar and click ***Start browser***. This step is not necessary when installing a client.

3.6 Linux /FreeBSD Installations

This section describes:

- installation requirements
- creating an installation directory
- copying the installation file
- unpacking the program directory structure
- starting the installation script

In case you are upgrading from an earlier version of P5, please see section 3.14, Upgrading to P5.

3.6.1 Installation Requirements

P5 is available online at <http://archiware.com/download>. The installation file is packed and has the format 'awpstxyz.tgz'. The letters 'xyz' stand for the given P5 version. If you downloaded P5 from the web, you can directly proceed to creating the installation directory.



You can install P5 only as the *root* user.

3.6.2 Creating an Installation Directory

Please create a directory for P5, taking into consideration the required disk space as described in the previous section.

Example: Let us assume that '/usr/local' is a directory on a mirrored partition with 5 GB of free space, thus satisfying the installation requirements.

```
# cd /usr/local
# mkdir aw
```

3.6.3 Copying the Installation File

Having created the P5 home directory, move the downloaded file there or download the file into this directory. Please replace the letters 'xyz' in the example below with the given PresSTORE version.

```
# cp ~/Download/awpstxyz.tgz /usr/local/aw
```

Correct the source directory according to your Download folder.

3.6.4 Unpacking the Program Directory Structure

You have to unpack the installation file before you can start the actual installation. First change to the P5 home directory.

```
# cd /usr/local/aw
```

Then unpack the installation file by using the following combined commands. Please replace the letters 'xyz' in the example below with the given PresSTORE version.

```
# gzip -d < awpstxyz.tgz | tar xf -
```

3.6.5 Starting the server

Now you can directly start the P5 software with.

```
# ./start-server
```

This script will check and proceed with all required installation steps.

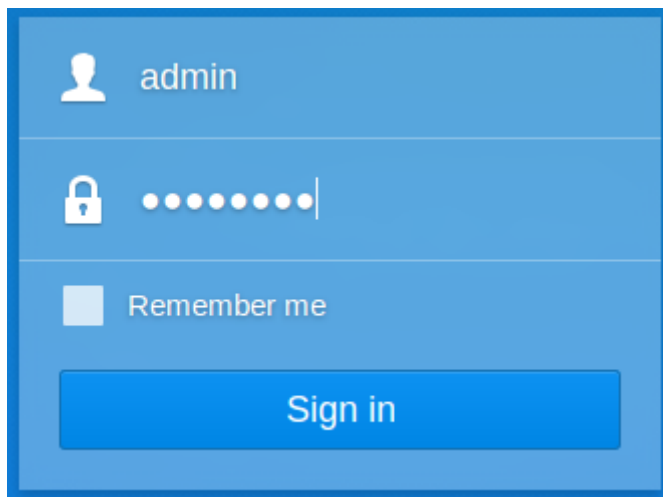
3.7 Installing on Synology NAS

Before the *Archiware P5 Suite* can be installed, it is necessary to obtain the corresponding installation package and upload it to the NAS. This can be done either manually or, if the NAS device has Internet connectivity, automatically through the Synology Package Center.


The *Archiware P5* installation package for Synology (SPK file) can be downloaded from the Archiware download area at <http://www.archiware.com/download>

3.7.1 Log in to Your Synology NAS

Open a web browser and navigate to your Synology NAS web interface. Log in as a user with administrator privileges.




After a successful login, open the *Package Center App* by clicking on the corresponding icon.

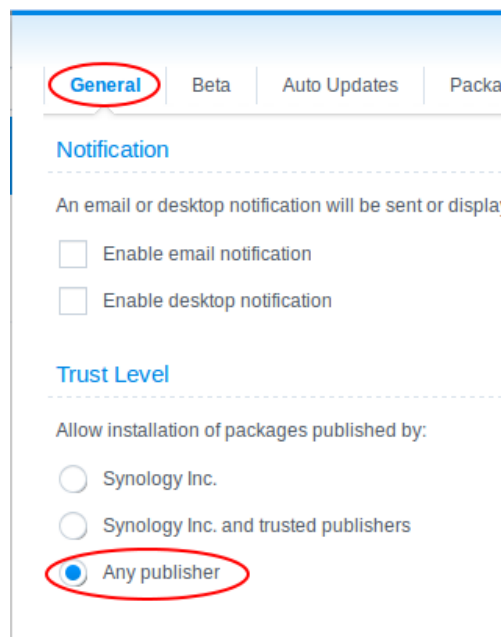


3.7.2 Edit Trust Level Settings

In order to allow the installation of packages from third parties, it is necessary to edit *Trust Level*/settings and allow the installation of packages published by any publisher.




- When the *Package Center* is opened, click on the *Settings* button.
- Within the *General* tab, select *Any publisher* in the *Trust Level* section, as shown below. Confirm your selection using the *OK* button.



3.7.3 Upload the SPK Package File

In the *Package Center* App, select *Manual Install*.



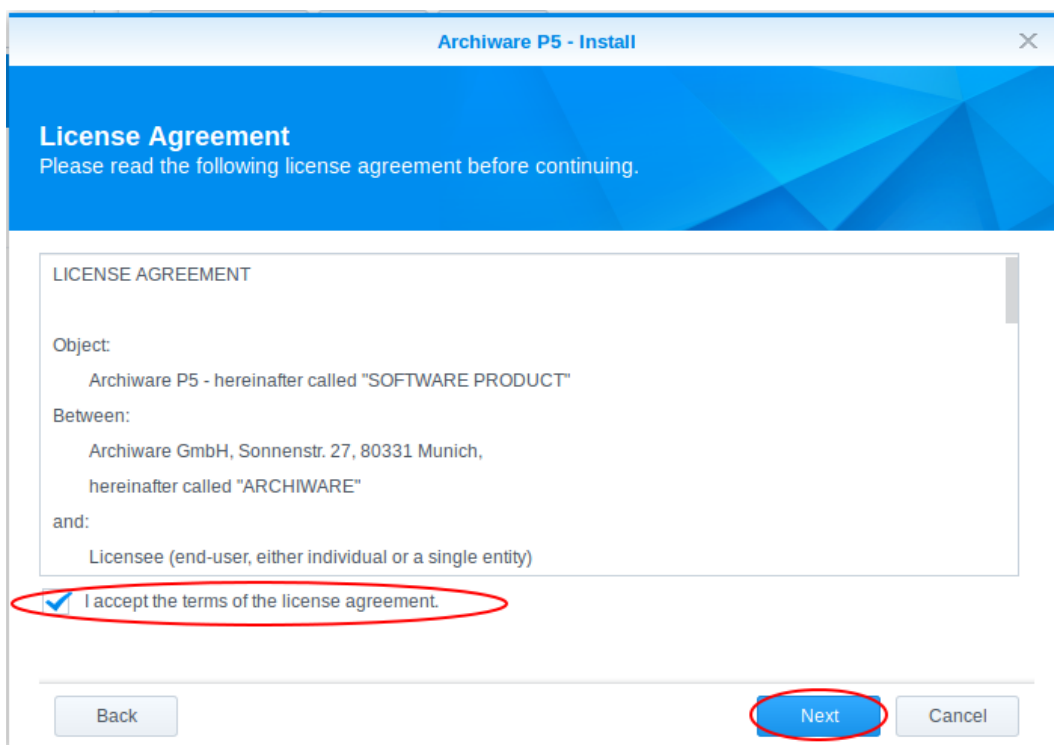
Within the newly opened window, click on the *Browse* button and select the downloaded SPK installation package file. Click *Next* to start uploading the installation package to the Synology NAS device. The installation procedure will start automatically when the installation package has been successfully uploaded to the NAS.

3.7.4 Installation Process

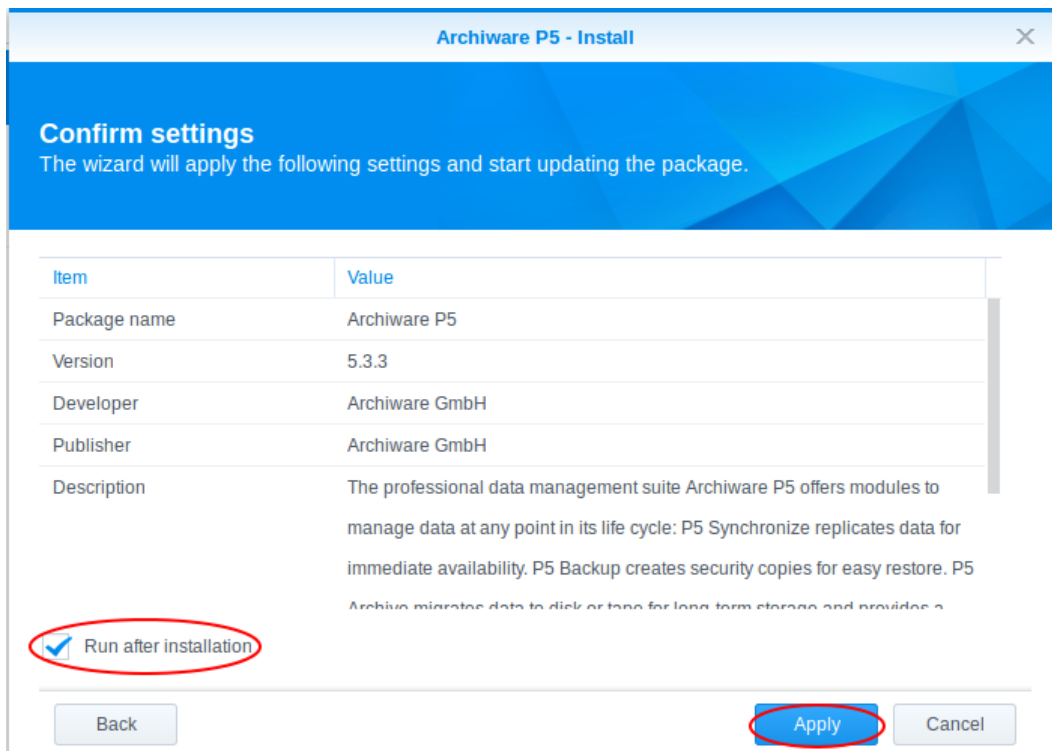
The *Archware P5 Suite* installation process is started after successfully performing all the steps outlined in *the above*

Read and Accept the License Agreement

The first step is to read and accept the *License Agreement*. Please read the agreement carefully in its entirety and accept it by ticking the checkbox *I accept the terms of the license agreement*. Confirm your selection by clicking the *Next* button.



Review Package Information and Configure Automatic Application Start




Please review the detailed package description for the selected *Archiware P5 Suite* installation. If you wish to have the *Archiware P5 Suite* run automatically after the installation, please tick the checkbox *Run after installation*. You can continue the installation process by clicking the *Apply* button.


Check That the Archiware P5 Suite is Correctly Installed and Running

After the installation process has finished, you can check the status of the *Archiware P5 Suite* by opening the *Package Center App* and selecting the *Installed* category. You will be presented with the list of all installed third party applications, including the *Archiware P5 Suite*, indicating their current status (*Running* or *Stopped*). If the installation was successful, the *Archiware P5 Suite* will be present in the list. If you have selected the *Run after installation* option, the status should be *Running*.

3.7.5 Accessing the Archiware P5 Suite

The *Archiware P5 Suite* uses a web interface accessible by any web browser. In order to open the interface, click on the *Main menu* icon  within the Synology web interface.

Within the menu, select and click on the *Archiware P5* icon.



Your web browser will open the *Archiware P5 Suite* web interface in a new browser tab.

Alternatively, you can access the *Archiware P5 Suite* web interface directly as described in section 4.1, *Starting the Browser*.



Known Limitation

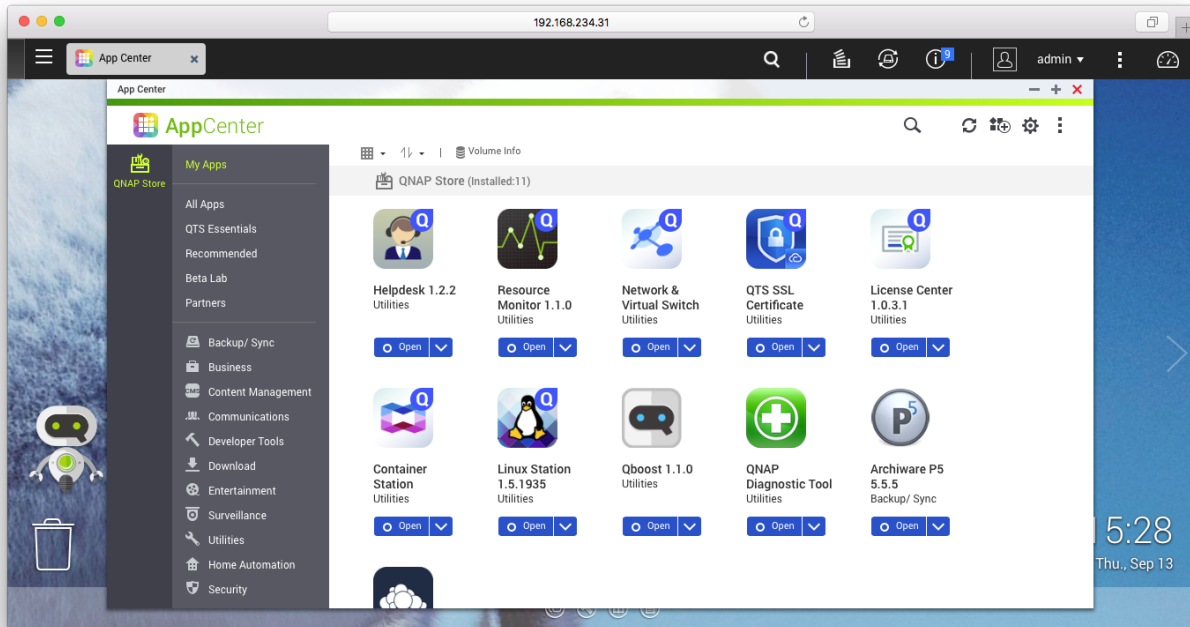
Due to the proprietary way in which the Synology NAS handles the management of Access Control Lists (ACLs) for the supported file systems, it is not possible to directly restore the files with their original ACLs. To work around this issue, please first restore the files to the desired location using the methods outlined in the *Archiware P5 Suite Reference Manual*. Once the files are restored, configure the desired file access permissions using the Synology web interface.

New Versions or P5 on Synology DSM 7.x

The installation on Synology DSM 7 requires that the application is approved by Synology. This process takes some time, so new P5 versions for Synology DSM 7 are released after a delay. During that time, the most recently approved version of P5 is available for download.

3.8 Installing on QNAP NAS

On Qnap NAS Systems, the *Archiware P5 Suite* can be installed through the AppCenter:



Login through a web browser on the QNAP Nas and open the AppCenter. the P5 application server can be directly installed from there. Simply click the P5 icon to install.

After the installation, a P5 icon is placed on the QNAP desktop, click the icon to access P5 through a web browser.

Alternatively, you can access the *Archiware P5 Suite* web interface directly as described in section 4.1, Starting the Browser.



In case the ip address of the QNAP NAS is not known, please visit the QNAP page <https://www.qnap.com> and download the utility Qfinder to locate the NAS in your LAN.

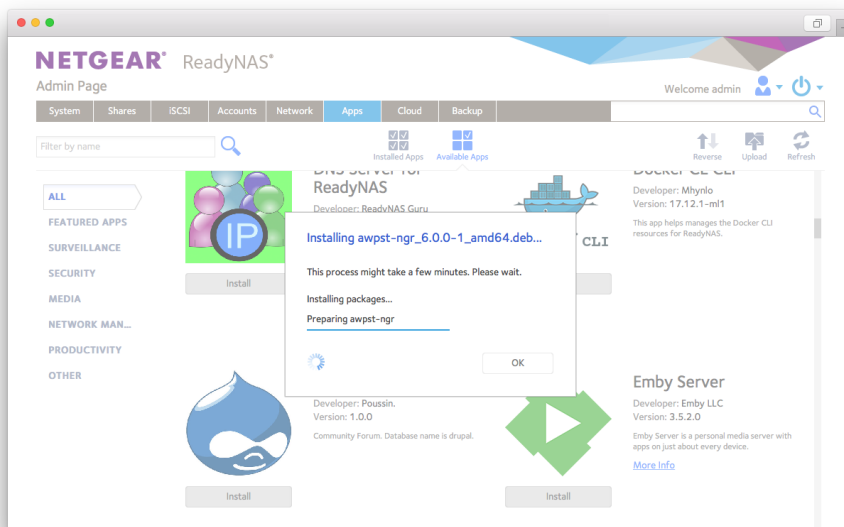
Note that new versions of P5 will not appear immediately in the AppCenter but need to be processed by Qnap first. The Qnap packages are also made available on the P5 download page <https://www.archiware.com/download-p5>. To install or update P5 on Qnap before the package appears in the AppCenter, download the package from the P5 download page to your local host, then select **Manual install** in the upper right corner of the AppCenter to upload and install the package that way.

3.9 Installing on Netgear ReadyNAS

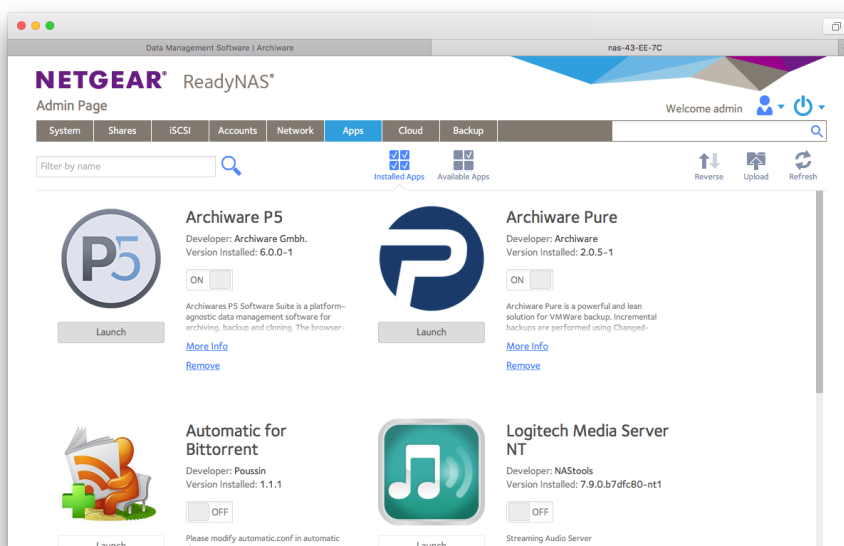
On ReadyNAS Systems, the *Archiware P5 Suite* can be installed through the Apps menu.

First, download the ReadyNas installer from <https://p5.archiware.com/download>

Log on to your ReadyNAS device web interface and open the Apps menu. Select Upload in the top right corner. Select Browse and navigate to the folder containing the downloaded P5 ReadyNAS package. Choose the .deb file and start the installation process by pressing Upload.



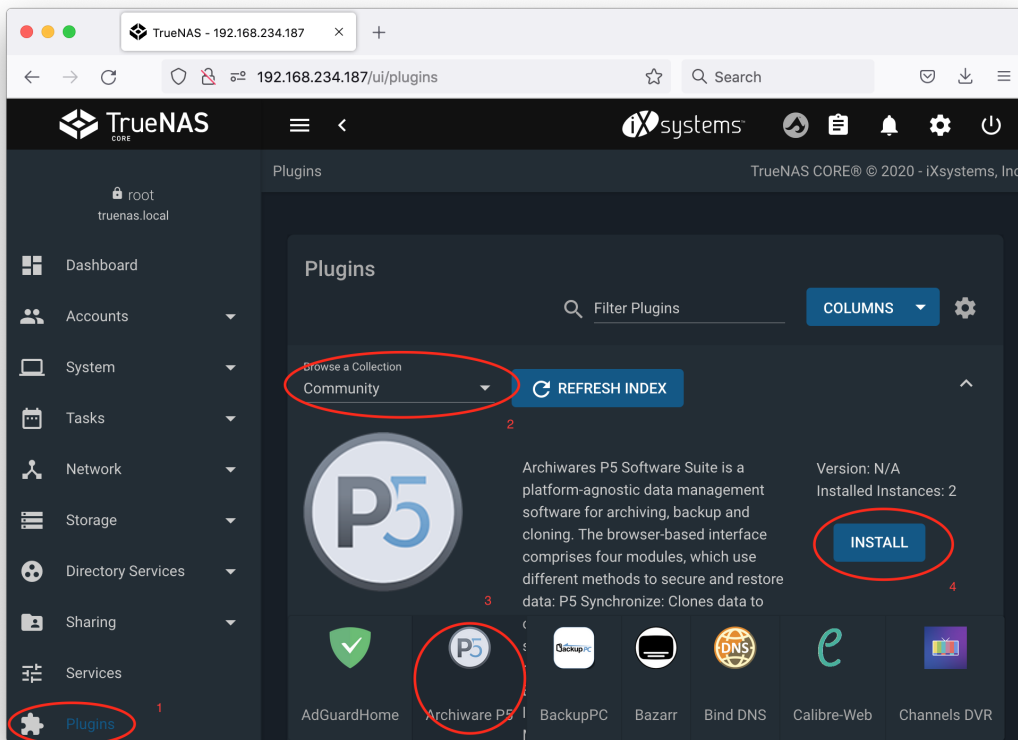
P5 is shown in the list of installed applications. From there it can be opened, started, stopped or uninstalled:




3.10 Installing on iXsystems TrueNAS

On TrueNAS Systems, the *Archiware P5 Suite* can be installed through the Plugins menu.

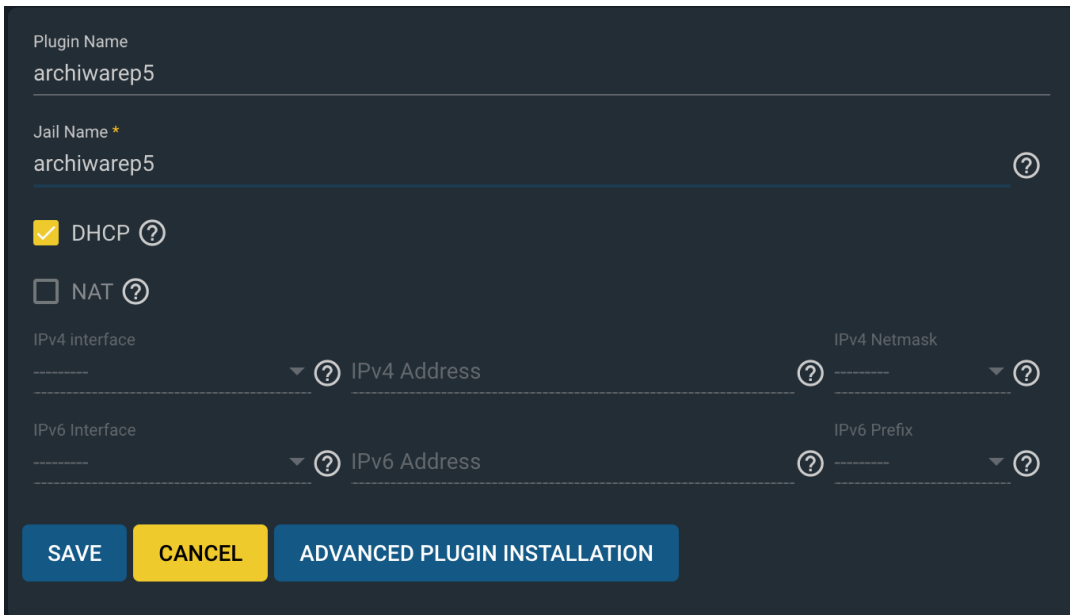
Log on to your ReadyNAS device web interface and open the Plugins menu (1). Select Community in the Browse a Collection menu on the top left (2). Select Archiware P5 in the list of the available plugins (3) and press Install (4) to start the installation process.



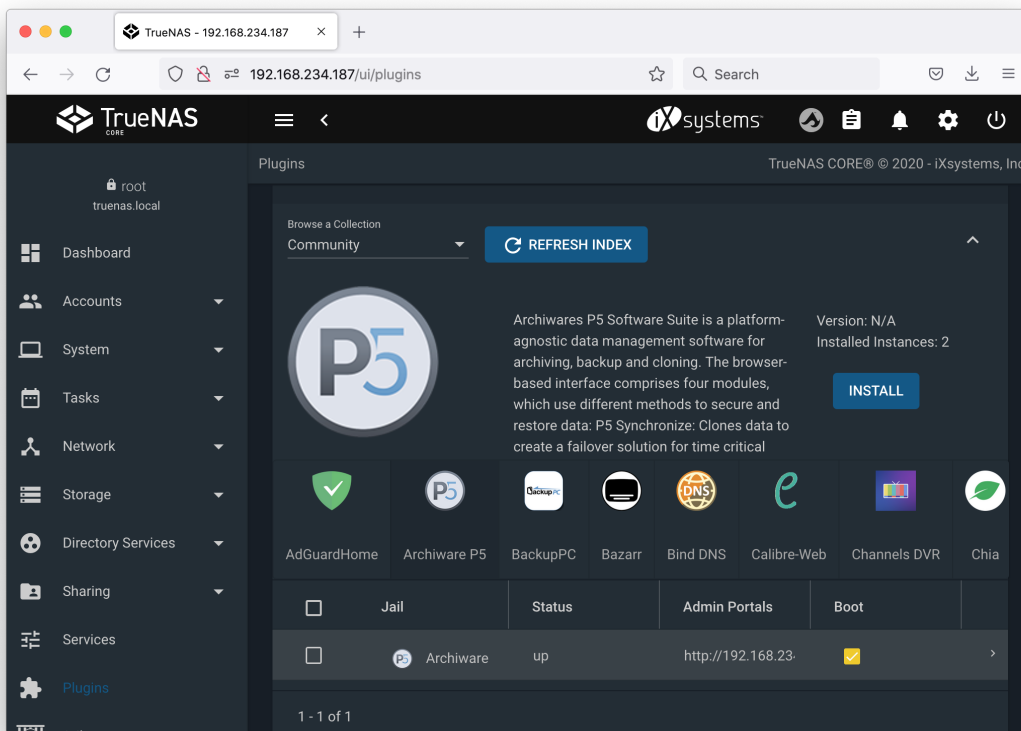
A new window is opened with a warning that this plugin is not produced nor supported by iXsystems, select Continue to proceed.



This leads to the next new window with installation options. For default installation just fill the desired name of a Jail for the new P5 plugin and select Save.



P5 is shown in the list of installed plugins. From there it can be opened, started, stopped or uninstalled:



3.11 Testing the Installation

After completing the installation, you can test the P5 application server. Use the command 'start-server' to start the server, if you have not already started it from the installation menu.

The program 'ping-server' in the P5 home directory attaches to the application server and returns connection information you will need to start a client session, as described in section 4.1, Starting the Browser.

```
# ./ping-server server
server pid: 442 (server is running)
access url: http://192.168.12.99:8000/login (use web-browser to connect)
```

Note: This command is not available on Windows installations.

3.12 Alternate Installations

For some operating systems, there are multiple installer downloads available:

- For macOS systems, there is an additional, generic .tgz installer that can be used to install on macOS by expanding the archive using the tar command in the Terminal.
- For Linux, there is a .deb for Debian – like package installers, a .rpm installer for systems using the RedHat package manager, and a .tgz install file that can be directly installed with tar. The generic .tgz install file does not use any system installer.

In all cases, the installed binaries are equal, i.e. the generic install file extracts the exactly same files as the package installers. The generic installer can be used in case of trouble with the package installer, for instance when downgrading to a previous system or in case the installer cannot resolve special dependencies.

3.13 Firewall Configuration Requirements

Background

P5 uses a proprietary protocol based on HTTP for communication between client and server. P5 therefore requires a transparent TCP transport without Layer-7 interference. Firewalls must perform stateful Layer-4 filtering only.

Firewall Configuration Requirement

For the dedicated communication channel between <IP_A> and <IP_B> on TCP port <PORT>, the firewall must be configured as follows:

- Allow TCP traffic based on source IP, destination IP, and port only.
- The service must be defined as generic TCP – not HTTP or any application-layer protocol.
- Only stateful Layer-4 inspection is permitted.
- Traffic must be forwarded transparently without protocol normalization.

Security

The security of this communication channel is ensured by two independent factors:

- Fixed IP addresses: Only the two known, pre-defined endpoints are permitted to communicate.
- Dedicated port: Communication is restricted to a single, specific TCP port.

3.14 Upgrading to P5

This section describes how to upgrade the P5 Server and Client.

3.14.1 Hints Regarding the Upgrade



- These installation notes are intended for system administrators with a solid knowledge and understanding of the Unix commands and must be executed by the *root* user.
- P5 versions prior to version 6.0 cannot be directly upgraded. Please either upgrade to version 6.0 before upgrading to P5 or remove all the files created with earlier versions (i.e. all the index files and configuration information) and install P5 anew.

Note: older versions may require intermediate installations so all updates are added to the configuration databases.

Versions prior to 3.2 must first be upgraded to 3.2.

Versions between 3.2 and 4.2 must first be upgraded to 4.2.

Versions between 4.2 and 4.4 must first be upgraded to 4.4.

Versions between 4.4 and 5.0 must first be upgraded to 5.1.

Version between 5.1 and 5.6 must first be upgraded to 6.1.

Version between 6.0 and 7.1 must first be upgraded to 7.4.

Version 7.4 can be updated to version 8

- *The latest release for these versions is available on the Archiware download page www.archiware.com/download-p5*
- A backup of the system is recommended before starting the installation. This backup should include both the P5 home directory and the existing backup and archive indexes.
- Before starting the upgrade, please check if the system requirements, as described in section 3.2 , System Requirements, are met.

3.14.2 Program Update

macOS:

Upgrading a macOS system requires the same procedure as a new installation.

See section 3.4, Installing on macOS. The installation procedure retains the existing configuration and index information.

Windows:

Upgrading a Windows system requires the same procedure as a new installation. See section 3.5, Installing on Windows .

After upgrading, P5 may automatically switch into a 30-day transitional period during which you are required to obtain and enter the activation keys for your new version. In this period, your product will remain active with your old license keys.

You have to enter the appropriate licenses for P5. For requesting new activation keys or to look up the existing license, please follow the steps in the online portal at www.ARCHIWARE.com/portal.

For macOS and Linux, an update through the browser is possible. Navigate to the P5 menu and pick Software update to proceed.

Systems with their own update management cannot update through the browser, these are Synology, QNAP, ReadyNAS and Windows.

4 Signing In and Entering the License

This section describes the basic elements of P5 and license activation.

4.1 Starting the Browser

P5 is a web-server application running in the background. The user interface runs in a web browser such as Firefox 3.5+ or Safari 5.1+ or Internet Explorer 10.0+.

To reach the interface sign in via your browser using the following URL:

[*http://localhost:8000/login*](http://localhost:8000/login)

Port 8000 is the default http port of P5 excepts on the following systems:

- On Synology NAS Systems, use the URL:
[*http://localhost:20000/login*](http://localhost:20000/login)
- On ReadyNAS Systems, use instead:
[*http://localhost:8500/login*](http://localhost:8500/login)
- On QNAP NAS Systems, use this URL:
[*http://localhost:8001/login*](http://localhost:8001/login)

In order to login via https protocol, the following URLs:

- On Synology NAS Systems, use the URL:
[*https://localhost:20443/login*](https://localhost:20443/login)
- On ReadyNAS Systems, use instead:
[*https://localhost:8943/login*](https://localhost:8943/login)
- On QNAP NAS Systems, use this URL:
[*https://localhost:8444/login*](https://localhost:8444/login)
- On all other systems, the https port is 8443, use this URL:
[*https://localhost:8443/login*](https://localhost:8443/login)

Starting with version 6, P5 does no longer support logging in with system accounts.



The default account to log in at the P5 web GUI is

Username: admin Password: admin

It is highly recommended to change the default password after login. P5 runs with administrator permissions, keeping the default password thus is a security risk.


Once you have signed into P5, the home screen is displayed.

4.2 License Activation

After signing in, you have to enter your licenses in order to be able to access all your licensed features. To create a license, please follow these steps:

1. Click home in the menu bar at the top.
2. Navigate to **My Installation > License Manager** in the main window.
3. Start entering a new license by clicking 'Add manually' at the bottom left.
4. In the license window, enter the data into the corresponding fields according to your Activation Key document. The version information in the document may differ from the installed version. Nevertheless, please enter the data exactly as shown in the document.

In case you do not own a P5 license, you may request a trial license from the [Archware portal](#) or through the `Get trial license` button.



The screenshot shows a window titled "192.168.234.197" with two tabs: "License Type 2" (selected) and "License Type 1". The version "V6.0.0" is displayed in the top right corner. The window is divided into two main sections: "License Data" and "System Identification".

License Data:

- Serial number:
- Status: (dropdown arrow)
- Product: (dropdown arrow)
- Expires: (calendar icon) Format: dd-mmm-yyyy

System Identification:

- Host ID:
- License key:

At the bottom right, there are two buttons: "Close" and "Apply".

5. Click **Apply** to confirm the input. A new license will be created. Provided the entry is correct, the license will be enabled.

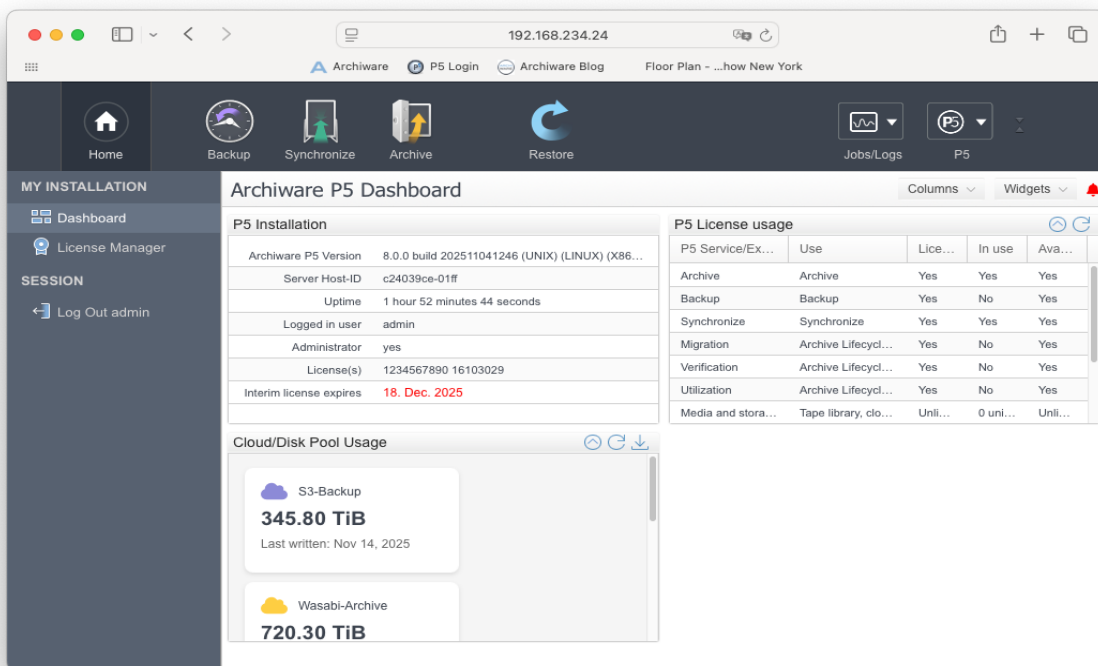
4.3 A General Introduction to the Server's User Interface

Having entered the activation keys, all the licensed modules can be accessed from the home screen of the P5 interface.

The P5 window consists of

- a top menu for selecting the desired module to control,
- a sidebar that depends on which module has been chosen, and
- a main window that depends on the current sidebar

The topmost or starting point of the P5 interface is the home screen:



The top menu holds the following items:



Home – brings you back to the home screen from anywhere in the interface.



Backup – leads to all the settings required to configure and run the backup.



Backup2Go – takes you to the screens required to configure Backup2Go.



Synchronize – allows you to set up data replication and transfers.



Archive – leads to all the settings required to configure and run archives.



Restore – allows to restore data from a backup or archive.




Jobs and Logs – allows access to logs of finished, planned and running jobs.



P5 – leads to help, support and further program information.

Some module icons in the top menu may be grayed out, depending on which modules are licensed. In trial mode, all the modules are accessible.

The left sidebar contains shortcuts to facilitate the navigation to the P5 configuration pages.



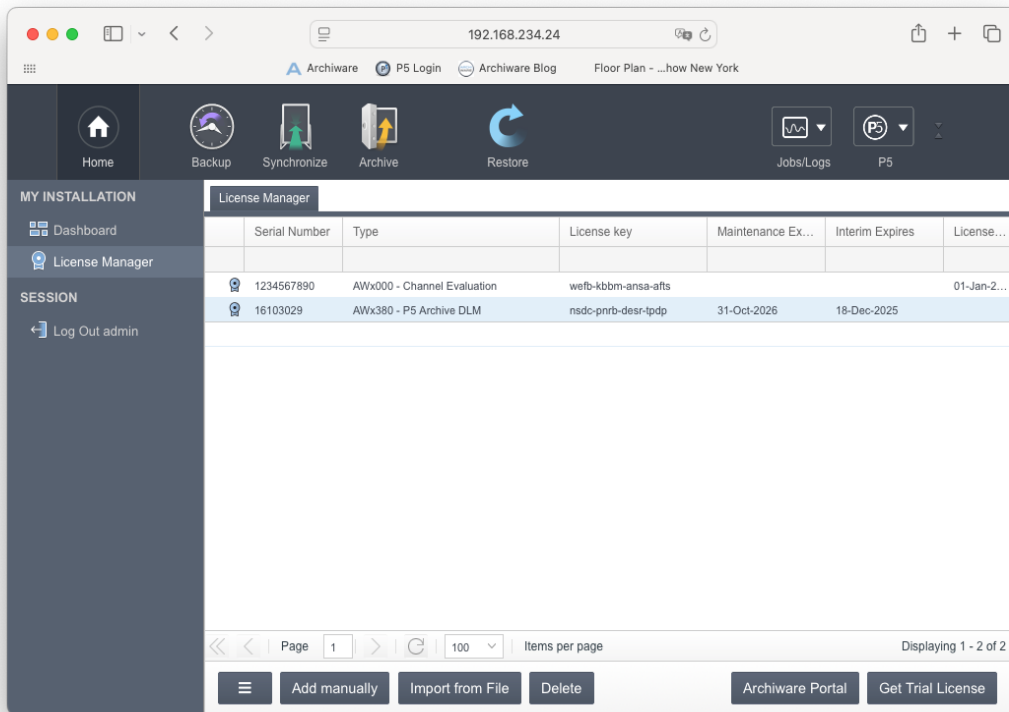
The content of the side bar depends on the selected module and takes you to the resources that must be configured for that module. For instance, you must set up a client if you wish to connect to another computer with Backup, Archive or Synchronize.

The **ADVANCED OPTIONS** section in the side bar takes you to the less frequently used configuration items like the index database or the log-in areas.

Each module stores its current selection within the side bar. This allows you to switch between modules without losing your current settings and to continue working where you left off in each particular module to proceed with its configuration.

4.3.1 The Installation Overview

Navigate to Home area and click the Overview tab get information about the installed version and license.

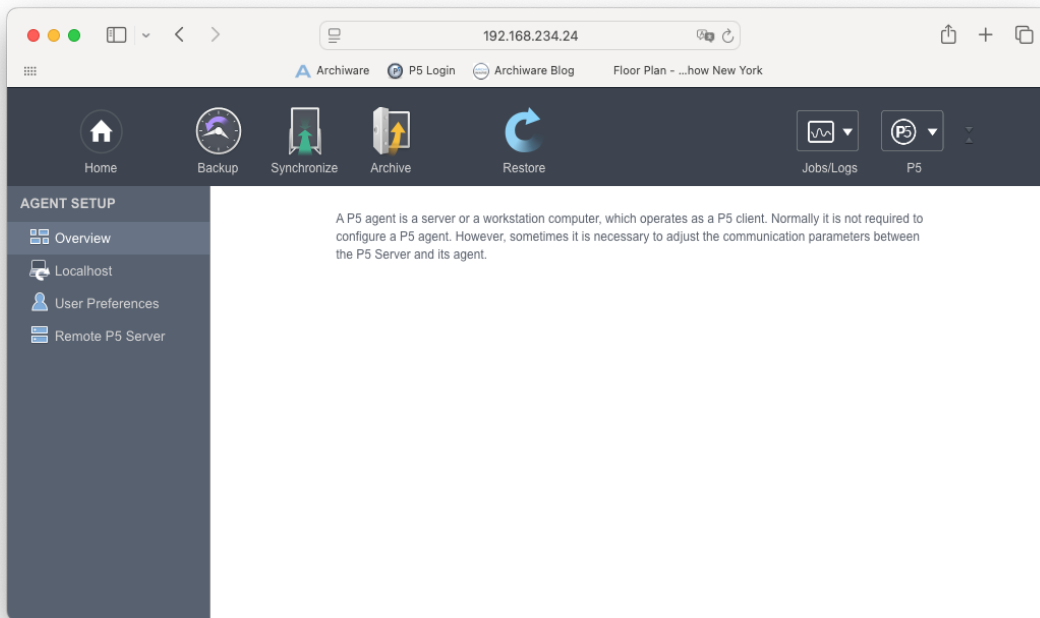


The display shows the installed software version and build, the host ID identifying the host where P5 is installed and the installed license.

Click License Usage to get detailed information about the used and free license components.

4.4 Using the P5 Agent Setup GUI

The P5 web GUI is also available on P5 client/agent or workstation installations. These installations run without license and the normal interface is locked. For client and workstation specific settings, a restricted interface is available in the P5 menu item Agent Setup:



This interface grants access to

- the Localhost client to configure settings such as encryption, see section 6.11, P5 Clients,
- to User Preferences to setup login permissions, see section 6.14, User Preferences, and
- on Backup2Go Workstations to setup the Backup2Go Server, see section Error: Reference source not found, Error: Reference source not found

Also, the P5 menu allows to change the password of the logged in user.



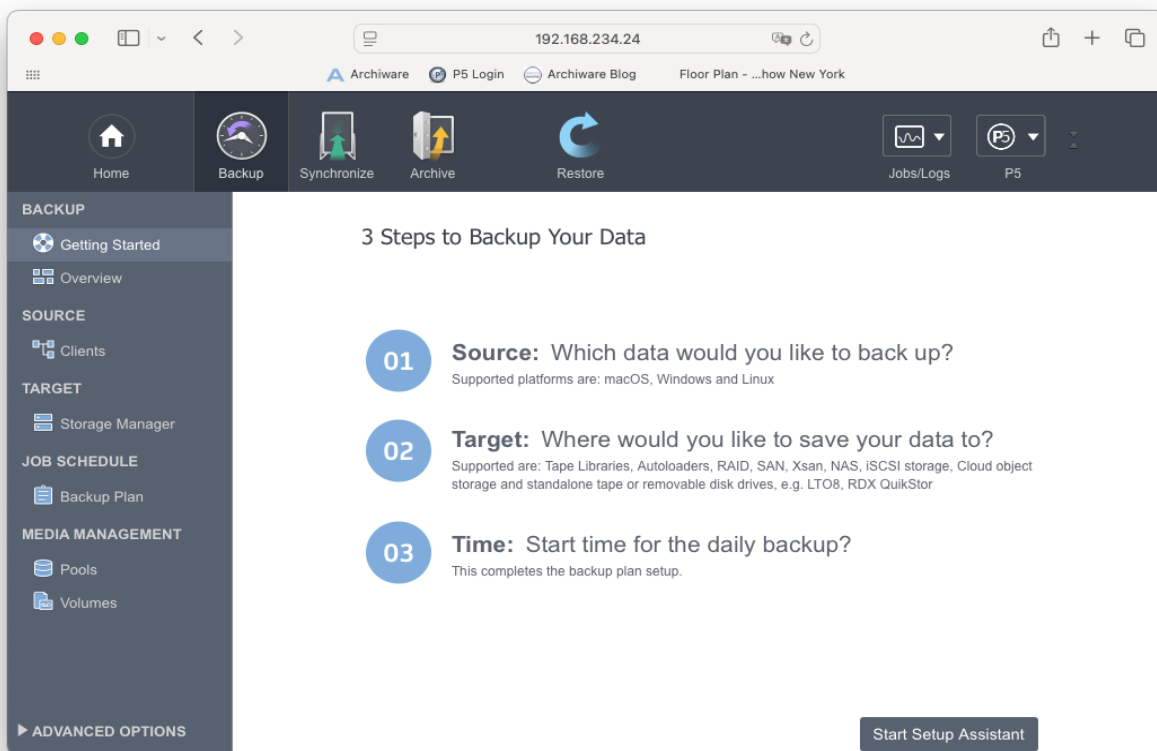
It is highly recommended to change the default password after login. P5 runs with administrator permissions, keeping the default password thus is a security risk.

5 First Steps With P5

This section shows how to set up a backup, archive or synchronization in a few steps. A detailed description of the setup can be found in sections 6 to 11 of the manual.

5.1 First Steps to Set Up the Backup

Select the Backup icon, the Getting Started item in the left side menu takes you to the Backup Assistant.



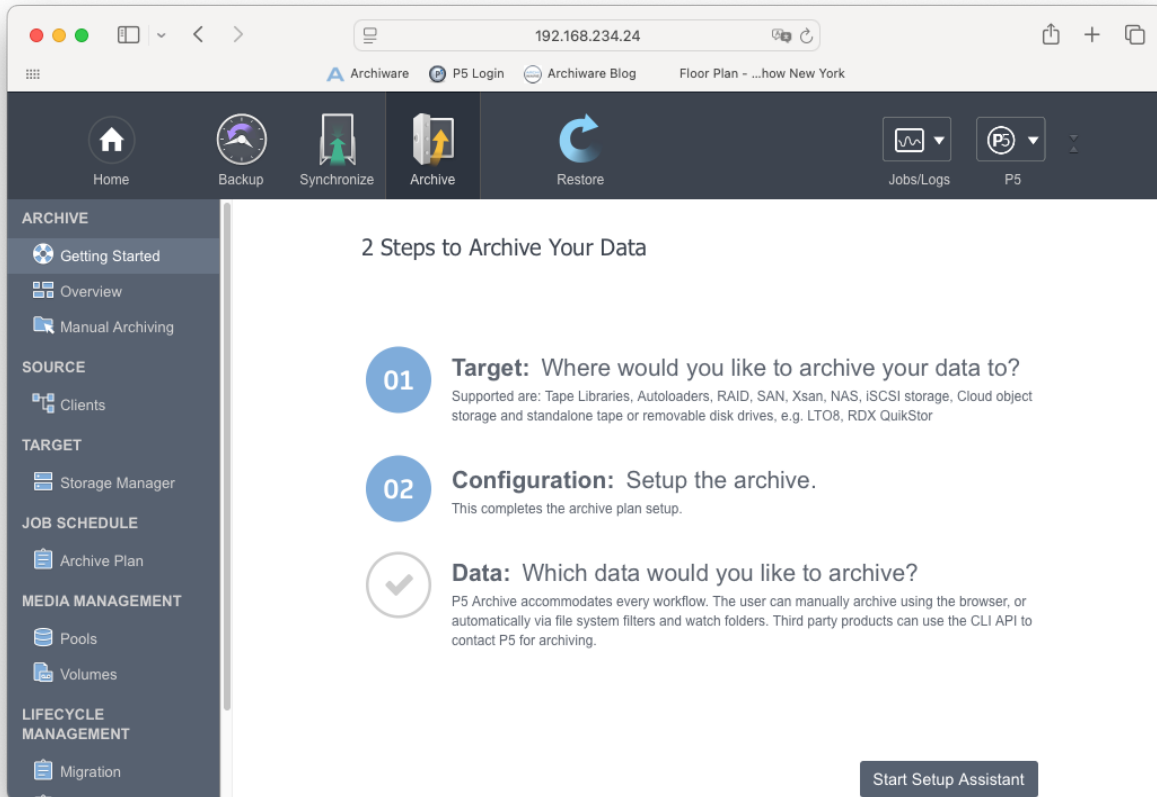
Simply follow the three steps to define which data shall be backed up, where and when.

The Backup Assistant will define a Backup Plan (see more in section 6.6, Backup Plans) that holds the info of what to save and the time schedule. Furthermore, a virtual or physical media changer or single drive is configured automatically and a media pool, either of Disk-Backup or Tape-Backup type, is created so that Volumes can be labeled in preparation for writing to them. See more about target storage in sections 6.2 to 6.3.

The first step also allows the user to define a client in order to save data from that client instead of the local machine. This configuration requires a P5 installation on that client computer. For further information about P5 clients please see section 6.11.1, Creating a Client.

5.2 First Steps to Set Up an Archive Plan

Select the Archive icon, the Getting Started item in the left side menu takes you to the Archive Assistant.



Simply follow the assistant's steps to define where and how to archive the data.

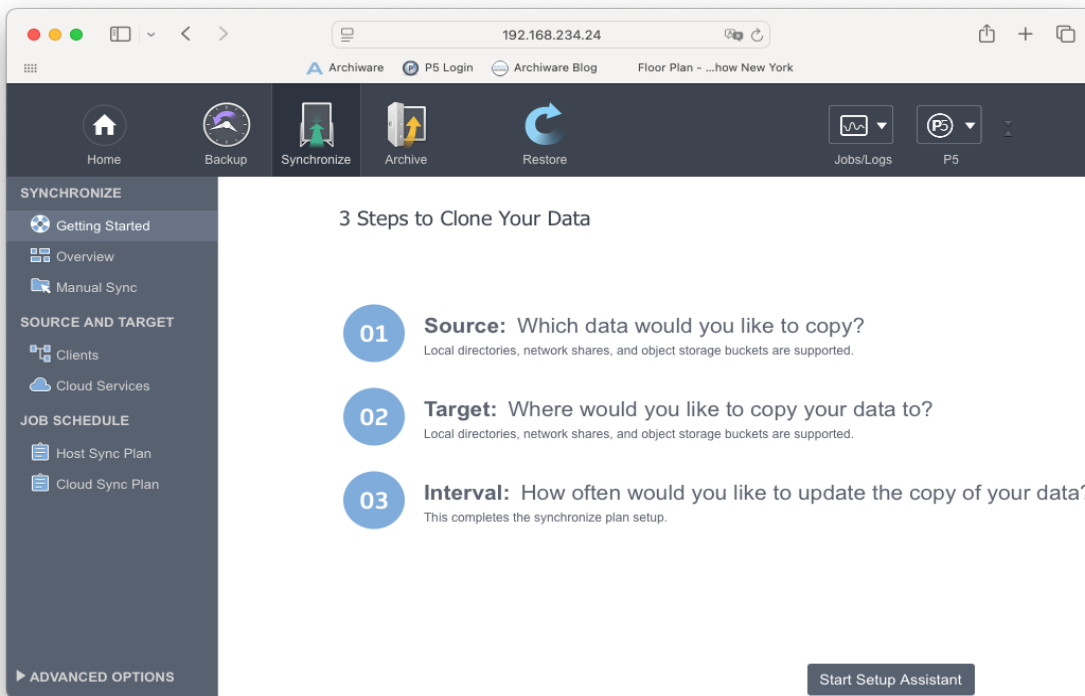
The Archive Assistant will define an Archive Plan that holds the configuration. Furthermore, a virtual or physical media changer or single drive is configured automatically and a media pool, either Disk-Archive or Tape-Archive, is created and Volumes can be labeled in preparation to write to them. See more about the target storage in sections 6.2 to 6.3.

The Assistant also allows the user to define a client in order to archive data from that client instead of the local machine. This configuration requires a P5 installation on that client computer. For further information about P5 clients please see section 6.11.1, Creating a Client.

As a last step, the assistant takes you to the file browser to select files for archiving. Automated archiving or preview generation can be set up in the Archive Plan, see more in section 6.5, Archive Plans.

5.3 First Steps to Set Up a Synchronization Plan

Select the Synchronize icon, the Getting Started item in the left side menu takes you to the Synchronize Assistant.



Simply follow the three steps to define which data shall be copied to where, and when. The Synchronize Assistant will define a Synchronize Plan that holds the configuration and the time schedule. For detailed settings in the synchronization Plan, please refer to section 6.7, Host Sync Plans.

The Assistant also allows the user to define a client in order to synchronize data to, or from, that client computer. This configuration requires a P5 installation on that client computer. For further information about P5 clients please see section 6.11.1, Creating a Client.

6 Configuring the P5 Server

Before P5 can begin saving data, it needs to be configured. This is done using the browser-based interface. The P5 user interface is run by an Internet browser. P5 stores its setup as resources within a configuration database that is managed by its application server. Among others, the application server is responsible for the control and execution of jobs, like Backup or Archive. The roles of the various resources are explained in section 2.4, Configuration.

To begin with, you will need a license, which enables the desired functionality. After obtaining and entering the license, just follow the `Getting Started` section of each of the P5 modules.

6.1 Licenses

Adding functions or upgrading a license is a simple task. It is not necessary to install any additional software. The license carries all the necessary information to enable the given functionality on the particular server identified by its host ID.

To activate a license, create a new resource in the licenses area of the Home screen.



WARNING: Before using a resource or running a job that is subject to licensing, P5 checks for a valid license. Resources that cannot be used due to insufficient licensing are automatically marked as 'Disabled'. After adding a license that adds new or additional resources, you will have to switch those resources to 'Enabled' manually in order for P5 to use them.

NOTE: P5 licenses are bound to an ID that clearly identifies the host P5 is installed on. In case of hardware changes, that HostID might change and the license will switch to an interim license. If so, you will need to login on the ARCHIWARE licensing portal at <http://www.archiware.com/portal> and request a new license key for the installation.

6.1.1 Creating a License

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Push Home and select License Manager from the side bar.
3. Choose New from the bottom bar in the main window.
4. Fill in the required fields.
5. To confirm your input, click Apply.

The screenshot shows a web browser window with the URL 192.168.234.197. The page has two tabs: 'License Type 2' (selected) and 'License Type 1'. The version number 'V6.0.0' is displayed in the top right corner. The 'License Data' section contains the following fields: 'Serial number' with the value '1234567890', 'Status' with a dropdown menu set to 'Enabled', 'Product' with a dropdown menu set to 'AWx000 - Channel Evaluation', and 'Expires' with the value '01-Jan-2020' and a calendar icon. Below this is the 'System Identification' section with 'Host ID' set to '299821d4-01ff' and 'License key' set to 'absb-pmtb-kaew-fnec'. At the bottom right of the form are 'Close' and 'Apply' buttons.

Serial number: Please enter the serial number from your license paper. Each serial number is unique and is validated.
The serial number is automatically filled in when selecting a trial license.

<i>Status:</i>	The status flag shows if this license and its resources are valid and/or enabled. It can take one of two values: Enabled: The license enables the corresponding resources. Disabled: The resources defined by this license cannot be used.
<i>Host ID:</i>	This value uniquely identifies the host where P5 is installed. You must know your serial number and the host ID when requesting an activation key. In case the host ID changes due to hardware modifications, please request a host key transfer from the Archiware license portal.
<i>Selection:</i>	This pull down field contains a list of all the available license groups. The serial number is automatically entered when a trial license is selected.
<i>Version:</i>	The version of the license – It does not have to match the installed P5 version. Enter the license version, not the P5 version into this field.
<i>Expires:</i>	Applies only to the trial license. Shows when the license expires.
<i>Maintenance Expires:</i>	Shows when the maintenance expires. This date has no influence on the installed P5, but updates released after that date will be unavailable.
<i>License Key:</i>	Archiware confirms your registration by sending your License Key and Registration paper.

6.1.2 Importing a License

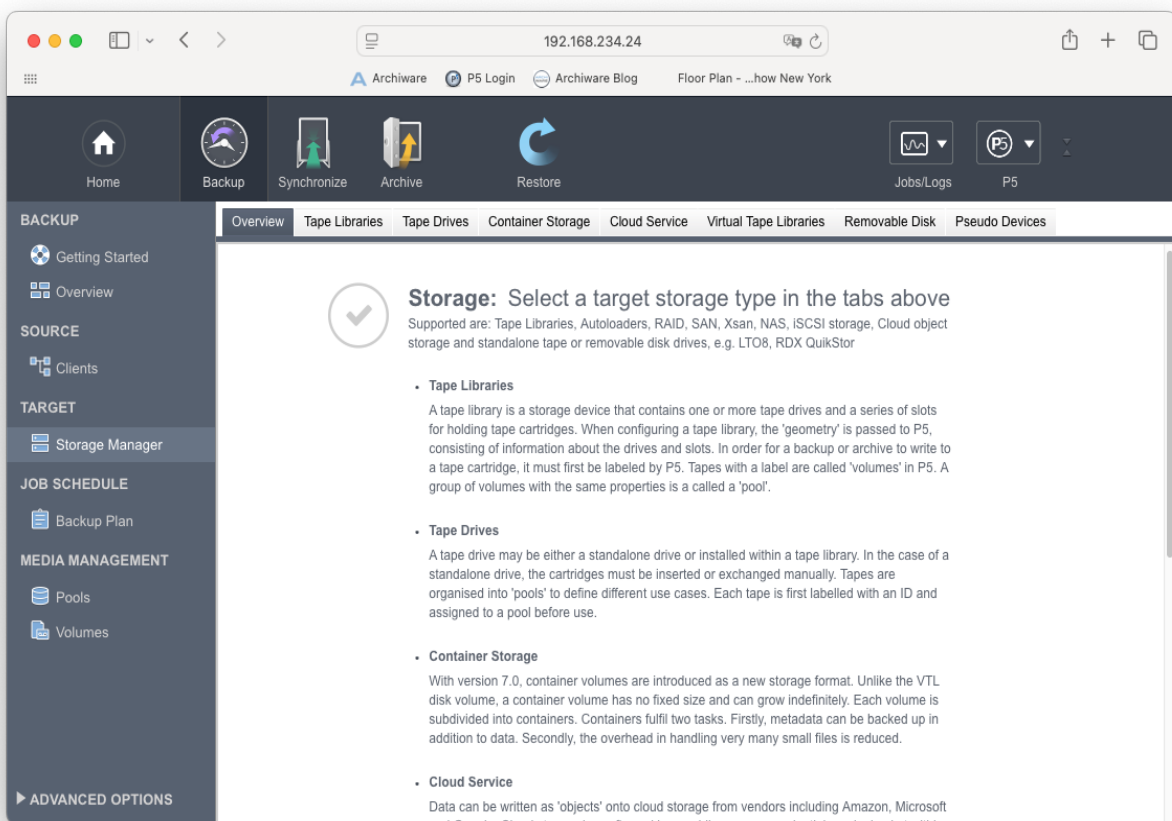
Starting with version 4.4, license files are sent by Mail when a license is activated in the ARCHIWARE licensing portal at <http://www.archiware.com/portal>.

In order to import that license into P5, open the licenses area of the Home screen and click the Import button. A dialog opens to select the license to import.

6.2 Storage Manager

The Storage Manager defines the targets where data can be written to with Archive or Backup. To enter the Storage Manager

1. Log into P5 using an Internet browser as an admin user. See section 4.1, Starting the Browser.
2. Select **Archive** or **Backup** from the top menu and then **Storage Manager**
3. Select the device type you want to create / configure. The **Overview** gives a brief description of the device types.



6.2.1 Tape Libraries / Media Changers

Tape libraries enable fully automated data save and restore operations without manual intervention. For large data volumes, a single tape is often not sufficient to store all the data. In such cases, media has to be replaced in the drive in the course of a backup in order to completely store all the data.

When several users require data to be restored at the same time, a media library is needed. It will insert the required media in the correct order into the drives, so that the data can be completely restored.

These tasks justify the use of a media library that in combination with P5 automates the storage processes.

The volumes, the devices and the changers are managed in the database. With a sufficient number of media in the changer, virtually no manual intervention is necessary, since P5 autonomously takes over all the operations. P5 automatically reuses the available volumes according to retention times defined by the backup plan.


Generally, P5 can use all SCSI libraries supported by the operating system.

6.2.1.1 Creating a Tape Library

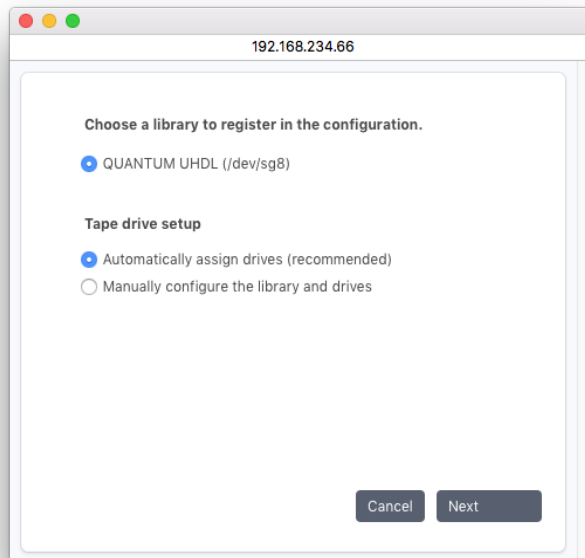
A wizard assists you in configuring one or more library devices. The drives built into the library are also configured in sequence and a device is created for each drive. Click `New Tape Library` to create and configure a tape Library.

The wizard walks you through the library configuration. The intermediate steps have to be confirmed by clicking `Next`.

The application server determines the number and type of the attached libraries.



In this dialog, P5 lists all the detected SCSI libraries. Mark the one you wish to create and click Next.



The wizard assigns a name for the library and suggests a description. You may modify the description and/or confirm by clicking Next.

The Jukebox Name is automatically assigned by the wizard and cannot be changed. The description can be changed later by editing the Jukebox definition.

When setting up the Jukebox manually, the application server determines the number of drives in the library and their types. Click Next to confirm.

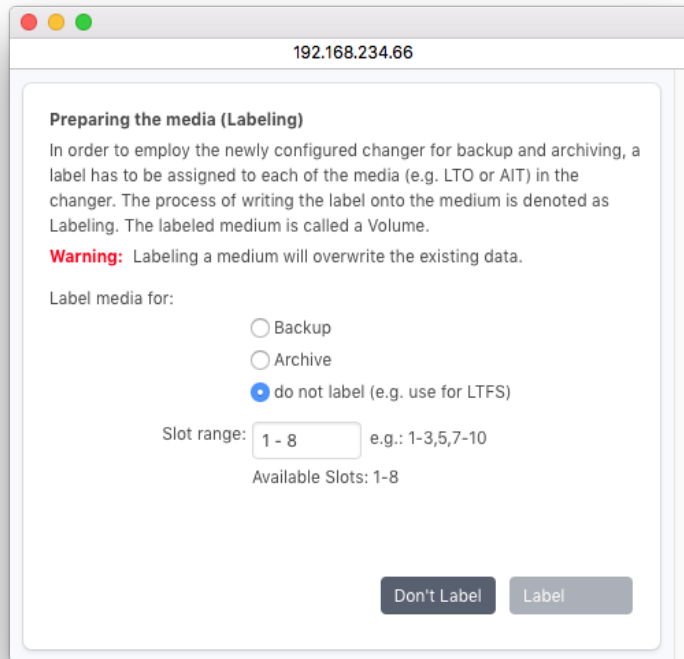
You will then be prompted to configure and create each of the drives. Complete or modify the information and click Install to confirm.

Note that during setup, the already configured drives and libraries must be connected to the system or must be removed from the configuration in P5.

- Name:** The drive's resource name. May be changed.
- Description:** Describes the drive. Changeable.
- OS path:** The operating system path (handle) to the driver's inode, e.g. /dev/rmt/2cn
- Type:** The largest media this drive supports. The pop-up shows a list of supported media. In case your media type is not in the list, select the largest supported media in the list.

The tape library installation will be saved by clicking Cancel, following the installation of the last drive.

Following the jukebox setup, the wizard offers the possibility to label tapes:



6.2.1.2 Modifying a Tape Library Setup

During the setup of the tape library in P5, the geometric addresses of all components is read from the physical library. This part of the setup cannot be edited manually.

In case the physical tape library setup is changed, please delete the library in the P5 Storage manager and create it new.

The screenshot shows a web-based configuration window for a tape library. The window title is '192.168.234.66'. It has two tabs: 'General Setup' (selected) and 'Additional Options'. The 'General Setup' tab contains the following sections:

- Library Information:** Library name: 'sg8', Description: 'QUANTUM UHDL', Status: 'Enabled' (dropdown menu).
- Setup:** OS path: '/dev/sg8', First slot: '1', Last slot: '8', Drive(s): 'sg7'.
- Library Geometry:** Vendor: 'QUANTUM', Type: 'UHDL', Slots: '8', Mailslots: '0', Drives: '1', Media type: 'TAPE'.
- Drive Cleaning:** Cleaning slot: '0', Maximum cycles: '0', Used cycles: '0'.
- Licensing:** Used license units: '8 Slot(s)', Free license units: 'Unlimited'.

At the bottom right, there are 'Close' and 'Apply' buttons.

Some values can be defined when editing the Library definition:

First and Last Slot:

These values define the slot range in the library that is used by P5. By default all slots are set, but in case for instance the license does not cover the full number of slots, it is possible to reduce the slot range and use a subset of the library.

To setup automatic drive cleaning in the library, define these values:

- Cleaning Slot:* The slot number where the cleaning cartridge can be found
- Maximum Cycles:* The number of cleaning procedures this cartridge can be used for, for instance 50.
- Used Cycles:* The number of cleaning procedures this cartridge has already been used for.

In case the tape drive signalizes that cleaning is required, P5 will automatically load the cleaning cartridge and perform the cleaning the next time a tape is to be loaded into the drive.



Please configure automatic cleaning either in the tape library or in P5, but not in both places. When automatic cleaning is defined in the tape library, the drive will not signal its cleaning requirement to P5 anymore.

6.2.2 Single Tape Drives

This section describes the configuration of standalone tape drives like for instance LTO tape drives. Tape drives can usually not be accessed directly. Before P5 can use a tape drive, it has to be configured. You can configure one or more devices that P5 can use to write data to. Each device carries several attributes, which describe the properties of the drive. Drives in changers are also represented as a device. Each such device is automatically generated when configuring a changer and must not be configured manually.


Generally, P5 can use all SCSI drives supported by the operating system.

You can configure a device either manually or with the help of the wizard. It is recommended to use the wizard since it automatically configures all the device attributes. It may however happen that P5 does not automatically recognize the desired device (see note). Then you will have to configure the device manually.

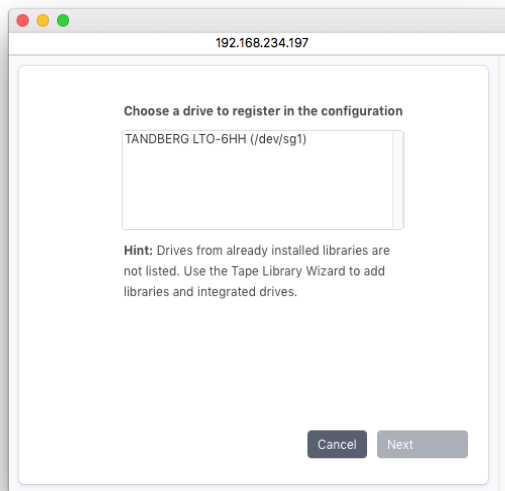
6.2.2.1 Creating a Device Using the Device Wizard

Click `Auto-detect standalone drives` from the bottom bar in the main window.

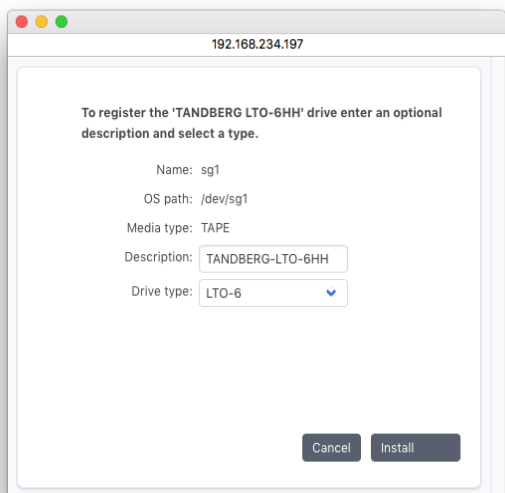
To start the automatic configuration, click `New Tape drive` and confirm by clicking `Next`.



A SCSI bus scan will be initiated and all detected sequential (tape) devices will be listed.



Following the scan, all the devices found will be listed. Devices that are already configured as part of media libraries, will be left out. Select a device and click Next to confirm.



In the *Description* field, you can additionally describe this device. Select the correct drive type from the pull-down menu and click Install to confirm your input.




Please do not configure drives built into a media changer as standalone tape drive as these will be hidden when the media changer / library is configured.

6.2.2.2 Creating a Device Manually

To start the configuration, select the Gear menu and click Add manually...

Enter the required information.

General Setup



The screenshot shows a web-based configuration window titled "192.168.234.197". It has two tabs: "General Setup" (selected) and "Additional Options". Under "General Setup", there is a table with columns "Drive", "Description", and "Status". The table contains one row: Drive "sg1", Description "TANDBERG-LTO-6HH", and Status "Enabled" (with a dropdown arrow). Below the table is a "Settings" section with three fields: "Media type: TAPE", "Drive type: LTO-6", and "OS path: /dev/sg1". At the bottom of the settings is an "Attention" section with "Cleaning required: No" (with a dropdown arrow). At the very bottom right are "Close" and "Apply" buttons.

Drive Name:

Enter the desired drive name. The following is checked:

1. Uniqueness: No other drive may carry the same name.
2. Validity: The name is at most 32 characters long and can be composed of ' [A-Z] [a-z] [0-9] - _ : . '. White space is not allowed.

Drive Status:

The status flag defines the availability of this drive. The possible states are:

- » Enabled: The drive is available for tasks as configured.
- » Disabled: The drive will not be used by the system.

Media type:

The media type of the device you are configuring. When defining a new drive, you may choose from a pop-up list of supported devices. Once the entry is stored, you cannot change this field any more. If you desire to do so, you will have to define a new device.

TAPE: Data is stored on tapes, e.g. AIT, DLT or LTO.

DISK: This type is present for compatibility reasons and should not be used. To setup disk storage, please refer to the virtual jukebox setup in the Storage Manager.

Drive type:

The type of the drive you are configuring. When defining a new drive, you may choose from a pop-up list of supported drive types. Once the entry is stored, you cannot change this field any more. If you desire to do so, you will have to define a new device.

OS path:

The operating system path to the device node on the file system, for example '/dev/rmt/2cn'.

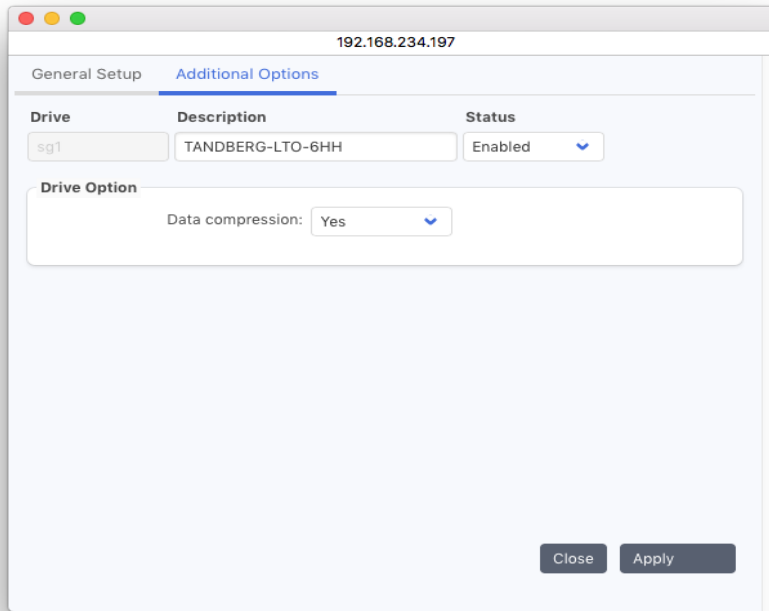
Cleaning required:

When a drive signals that it requires cleaning, P5 will automatically toggle this flag to 'Yes'. Once you have cleaned the drive, you will have to set this flag back to 'No' manually.

Note:

A drive that needs cleaning will not be used to store data.

Additional Options



Data Compression:

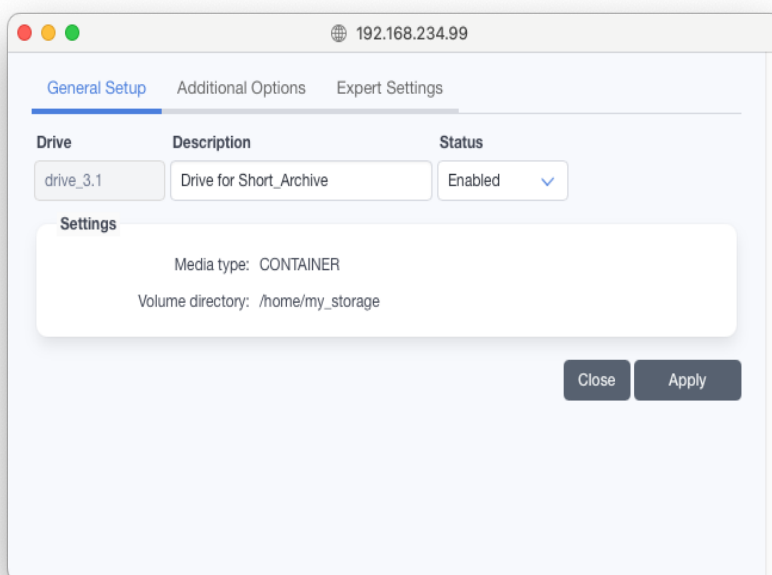
By default, P5 uses the hardware data compression function of the tape drive. In case of a virtual tape drive on disk, P5 uses a zip based compression algorithm when writing data. The compression can be turned off, for instance if the data is already in compressed format.

6.2.3 Container Storage

The Container Storage is a modern storage format for disk and cloud targets. The storage operates file wise for bigger files and also is capable to recycle data file wise. So the container storage can be used for incremental only backups, which in turn makes it an ideal storage format for cloud backups.

The container format is potentially unlimited in size, there is no splitting in multiple volumes as there is with tapes or virtual tape libraries.

When creating a container storage, P5 generates in addition a pseudo device as drive to write to and a media pool (see section 6.3, The Media Pool) to integrate the storage into the existing mechanisms in P5. A Container Storage corresponds to one P5 volume and one Media Pool.



Directory:

The container storage requires a directory where the data will be written to. In a disk backup, the single containers are placed in that folder, in a cloud backup the folder is used as staging area. Note that the container storage has no preset size limit, it will stop with an error in case no more space is available.

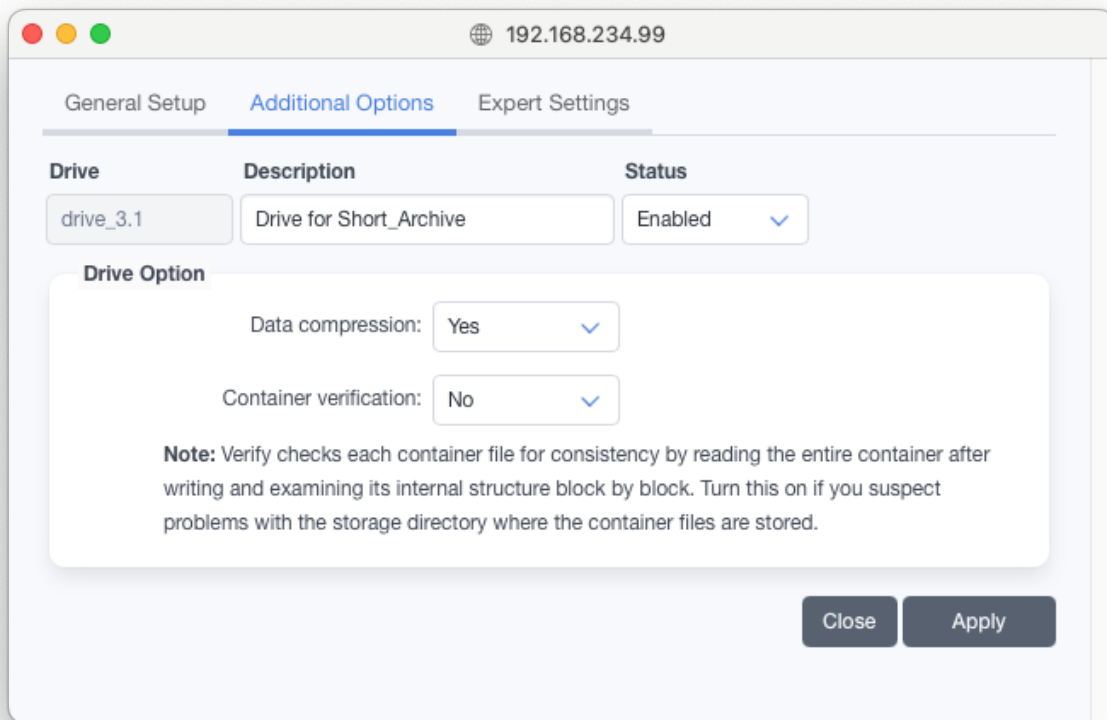


The container storage requires the target folder to be permanently available, it must not be used with removable hard drives.

A discussion of the new format can be found here:

<https://blog.archiware.com/blog/p5-version-7-new-container-format-concept-and-benefits/>

6.2.3.1 Container Storage Additional Options



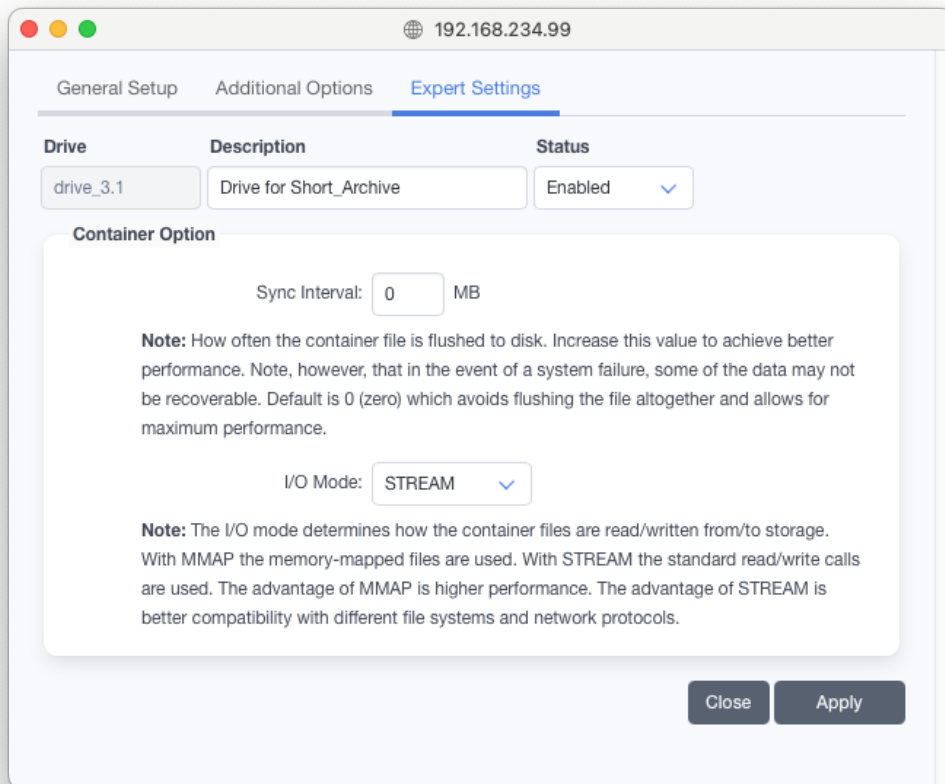
Data compression:

Set this option in order to compress the container files. This reduces the physical storage usage, depending on the compressability of the saved data.

Container verification:

The verification option checks the container file after writing. Turn this option on in case problems on the physical storage are expected. The verification will need to read the data after writing, that reduces the overall performance. In case the container storage is a cloud storage, this option causes additional data transfers when reading the container files.

6.2.3.2 Container Storage Expert Settings



Sync Interval:

While writing the container files to disk, hardware buffering applies so in case of a failure, the written file may be incomplete on disk. This parameter defines which size may be written without forcing a flush to disk. Flushing to disk causes a delay during write but ensures the written data so far is physically written. The default value of zero does not flush at all and is fastest. A nonzero value will enable intermediate flushing to disk after the given size. Bigger values are faster as less flushing appears.

I/O Mode: Stream

The I/O mode Stream is the normal disk access mode.

I/O Mode: MMAP

The MMAP or memory mapped mode allows faster access, but it does not work with all disk volumes. For MMAP mode, the volume must be a local volume that is directly accessed by the host.

6.2.4 Cloud Services

P5 supports archives and backups to be written or copied to or restored from a cloud service.

6.2.4.1 Defining a Cloud Service

Click **New Service** in the Cloud Service's bottom bar to define a Cloud Service.

The screenshot shows a configuration window for a cloud service. The window title is '192.168.234.24'. It is divided into two main sections: 'Object Storage' and 'Account Credentials'.
In the 'Object Storage' section:
- 'Cloud Service Provider' is a dropdown menu set to 'Amazon S3'.
- 'Description' is a text input field containing 'Cloud Service 1'.
In the 'Account Credentials' section:
- 'Access Key ID' is a text input field containing 'bla'.
- 'Secret Access Key' is a text input field containing '...'.
- 'Bucket Name' is a text input field containing 'blu'.
- 'Storage Class' is a dropdown menu set to 'S3 Standard'.
- 'Parallel uploads' is a text input field containing '1', with the label 'Number of files to upload concurrently' below it.
At the bottom of the window, there are four buttons: 'Compatibility ...', 'Test Connection', 'Close', and 'Apply'.

To use a cloud service as a target, assign it to a volume pool, see section 6.3, The Media Pool.

The values to setup a cloud storage depend on the storage provider and type, so this table just lists some common values:

- Access Key ID, Secret Access Key:* The account data that is used to login to the cloud service. Get these from your your Cloud Service Provider.
- Bucket Name:* The bucket in the cloud storage where the data shall be stored.
- Storage Class, Retrieval Option:* The storage class and retrieval option defines how the storage provider stored the data and how fast data can be retrieved. For instance a "Deep Archive" storage class usually means the data is usually not instantly accessible for restore but must first be brought back from an offline storage by the provider before it can be restored from the cloud storage.
- Parallel operations:* Define here how many parallel operations P5 may execute. Parallel operations may speed up up- and downloads, depending on the cloud storage. High values are however not recommended, 1.. 4 is a good starting point.
- Upload part size:* This size defines up to which size P5 can transfer data to the cloud in a single upload. Bigger data sizes are splitted and transferred in a multi part upload.
- Retain restore days:* When restoring from a deep archive, the files must be first be requested for download and are then made available on an online storage. This value defines how long to keep the files on the online storage. When restoring huge amounts of data (more than can be downloaded on one day), increasing this value speeds up the restore.

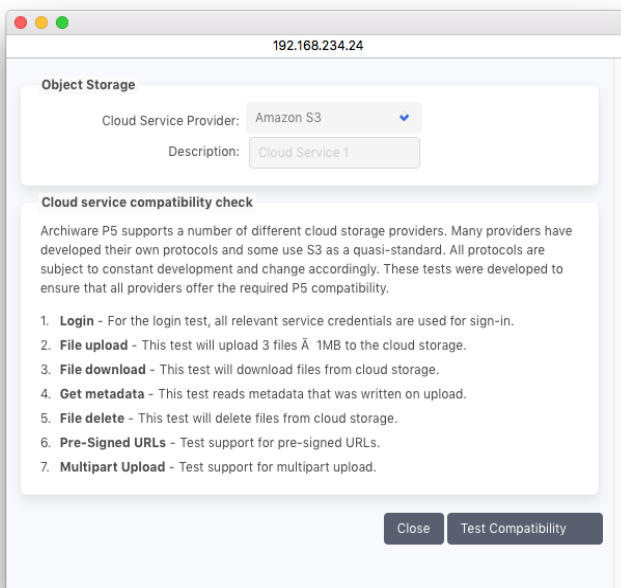
6.2.4.2 Supported Cloud Services

P5 supports addressing the following cloud storage:

- Amazon S3
- Amazon Snowball
- Alibaba Cloud
- Backblaze B2
- Generic S3
- Google Cloud Storage
- Huawei Cloud
- Microsoft Azure
- Seagate Lyve
- Wasabi Cloud Storage
- Wasabi Ball
- Amazon Glacier (Legacy)

6.2.4.3 Test the Cloud Storage

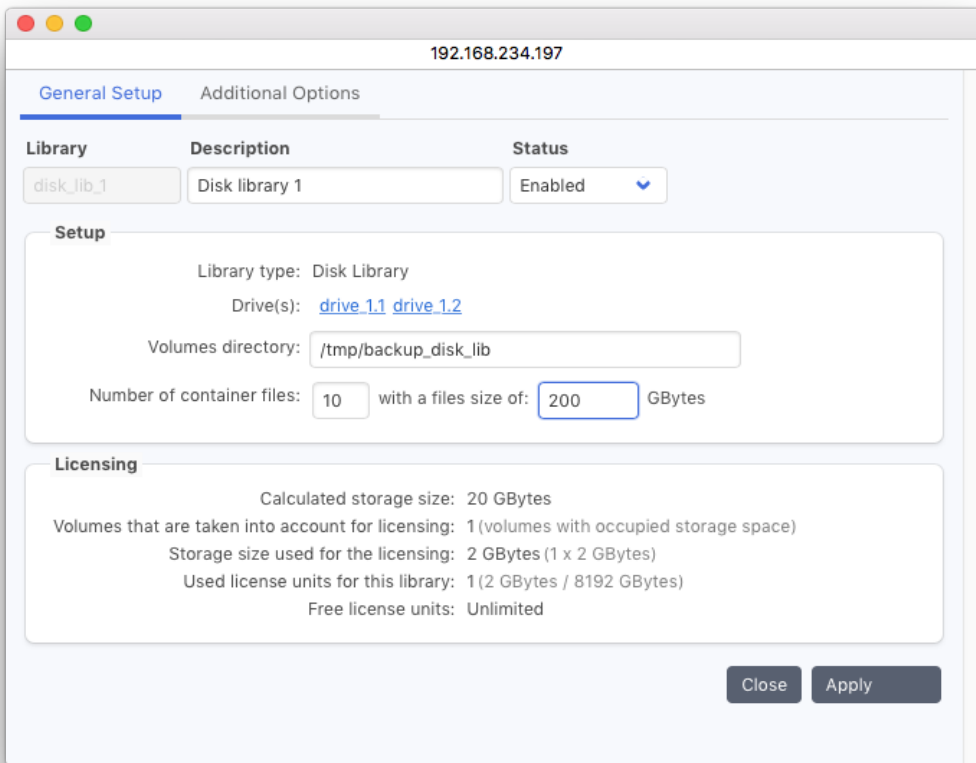
The connection test performs some basic tests regarding the accessibility of the cloud storage.



6.2.5 Creating a Virtual Disk Library

The virtual jukebox handles disk files as tapes of a physical library, thus allowing the management of disk and RAID systems, for example the automatic recycling and mounting of disk based volumes. A wizard aids you in creating and configuring one or more virtual jukeboxes. Two virtual (pseudo) drives associated with the changer are also automatically configured by the wizard.

The virtual tape library is a legacy format, for new projects we suggest creating container storage instead. In new installations, the virtual disk library is not visible. To enable it, please open the server settings and set `gui.Jukebox.mediaType.DISK.configurable.boolean` to 1. See section 6.18, Advanced Server Settings for details.



Fill in the fields and click Apply to start:

Volume directory: The directory that will be used for disk files. P5 will place the files in this directory.

P5 can use: Select here whether P5 should use the entire free space or limit the drive space.

Create (label) volumes: Choose whether the jukebox should be used for backup or archive purposes. P5 will then create a volume pool and prepare "virtual tapes" in that pool. You may select "do not label" to skip that step.



Each virtual jukebox is created with a default layout, depending on the available space on disk. By default, ten slots (for ten virtual tapes) with a minimum size of 2 GB are created. In case you want to change the library layout after a jukebox has been created, we suggest not to label the tapes, since the size of existing virtual tapes cannot be changed until they are relabeled.

In order to resize a virtual jukebox without losing data:

- Do not reduce the number of slots in the jukebox to less than the number of existing virtual tapes.
- Instead of increasing the size of the tapes, add further slots. Increasing the tape size will only increase the storage space after relabeling the tapes which would erase the data on that tape.

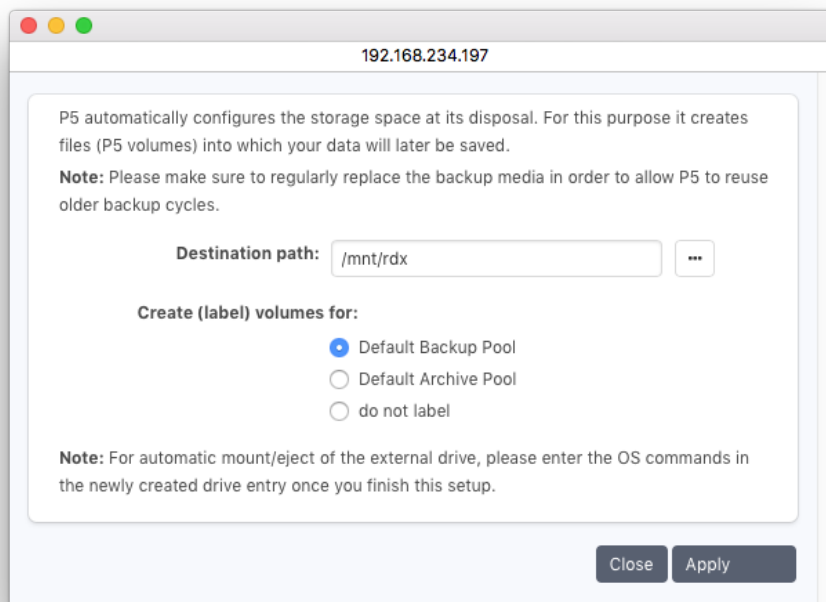
6.2.6 Creating a Removable Disk Storage

You can configure a removable disk with the wizard. It is recommended to use the wizard since it automatically configures all the device attributes. It may however happen that P5 does not automatically recognize the desired device. Then you will have to configure the device manually.

Disk drives can be configured in P5 as virtual tape library or as removable disks. Removable disks, like single tape drives, are written as one data storage. This method is intended for removable disks, not for bigger fixed disk arrays.

Removable Disk Storage is a legacy format and shall only be used in case of external hard disks that are planned to be exchanged, like for instance RDX disk cartridges. In new installations, the removable disk storage is not visible. To enable it, please open the server settings and set `gui.Device.deviceType.Removable.configurable.boolean` to 1. See section 6.18, Advanced Server Settings for details.

Click `New Removable Disk` to configure a removable disk storage.



Fill in the fields and Click `Apply` to create a removable disk storage.

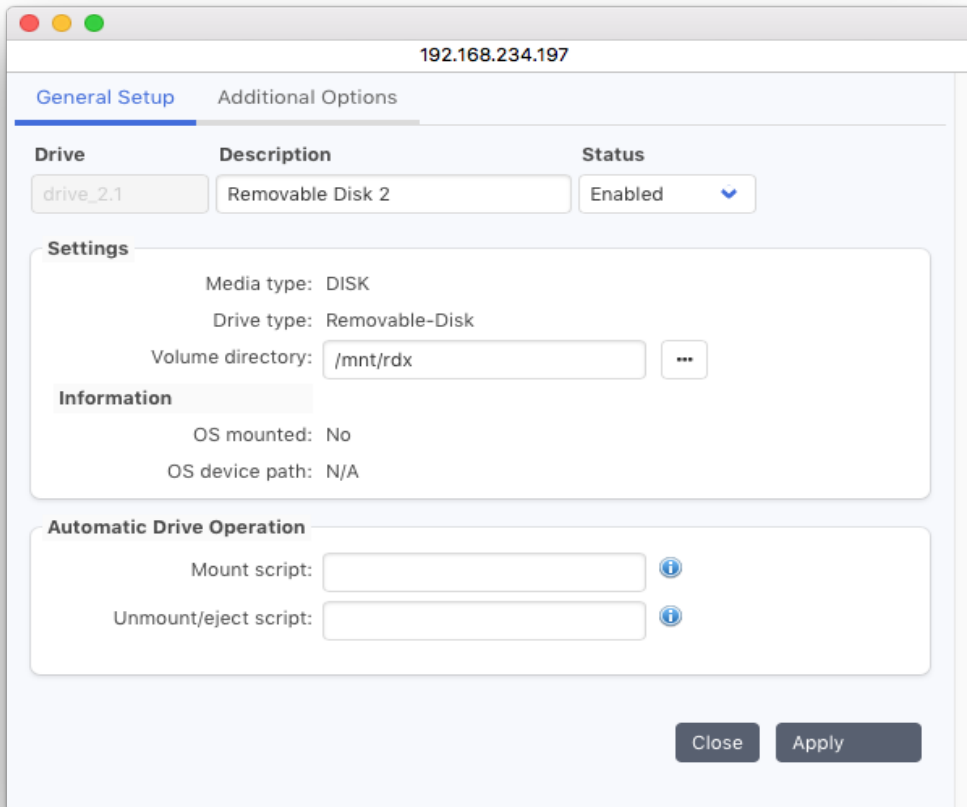
Destination Path: The directory that will be used as Removable Disk. P5 will place the files in this directory. Usually this is the path (on windows, the drive letter) where the removable disk is mounted.

Use media for: Choose whether the drive should be used for backup or archive purposes. P5 will then create a volume pool and prepare "virtual tapes" in that pool. You may select "do not label" to skip that step.



On OSX and Linux, removable disks are mounted by their volume name. In case multiple disks shall be used, format the disks with the same name. In this setup P5 is capable of handling disks with different sizes of up to 1 TB each.

Edit the created removable disk drive in order to change its configuration.



Volume directory: The directory that is used as Removable Disk. P5 will place the files in this directory. Usually this is the path (on Windows, the drive letter) where the removable disk is mounted.

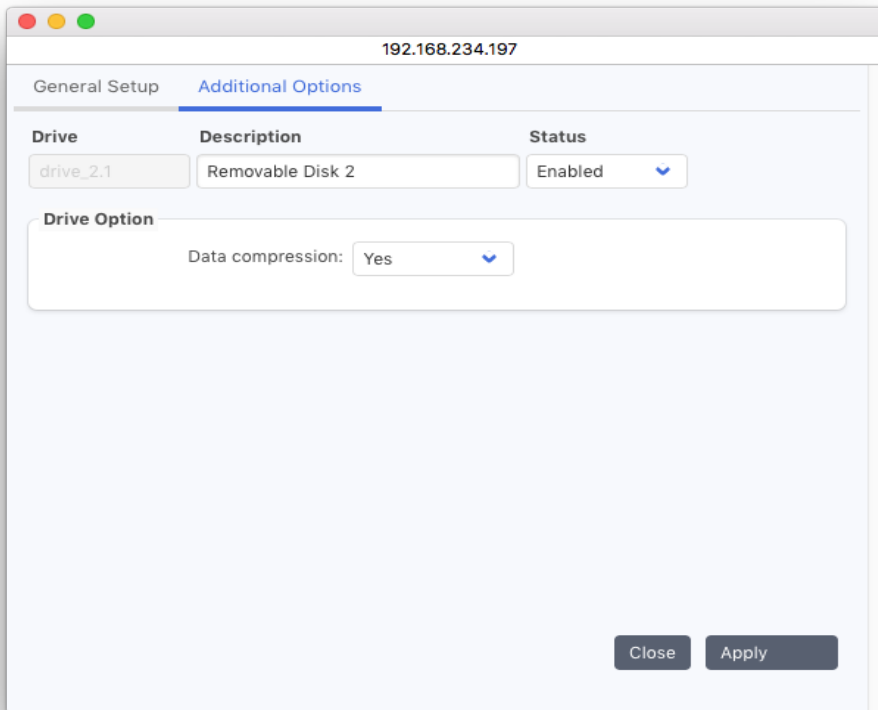
Mount script: An optional script can be defined to mount the external device before use.

Unmount script: The unmount script can be used to unmount the disk after use. Devices supporting a mechanical eject command (for instance RDX drives) can be ejected with this script.



Unfortunately, the device node used in mount commands mounted at the given path (on Windows, the drive letter) may change after rebooting the host. This is because the operating system enumerates the devices in the order they are found after booting.

A good practice to keep the assignment of device file and drive fixed is to connect external devices in the same order.



The Additional Options allow specifying whether or not compression shall be used when writing. By default, compression is turned on. P5 uses a ZIP compression mechanism on disk drives.

1.1.1 Storage Manager: Hamburger Menu

For each of the configured libraries, there are several options you can access from the Gear menu.

6.2.6.1 Label Volumes...

This function is available for discrete media like tapes or removable disks.

For a description read section 10.1, Labeling Media.

6.2.6.2 Take Inventory...

In This function is available for discrete media like tapes or removable disks.

In order for P5 to automatically manage media within a library, for example transport the required media into a drive, it has to take an inventory of the library. When running an inventory, P5 identifies the volumes and their positions within the library based either on the bar code on each of the media or on the label that was written on the media when labeling it. P5 keeps a copy of the inventory for each of the changers in its database. Any P5 job (for example Backup, Archive, Restore) that needs a volume within a library for its execution, uses the database information to automatically mount or unmount the volume.

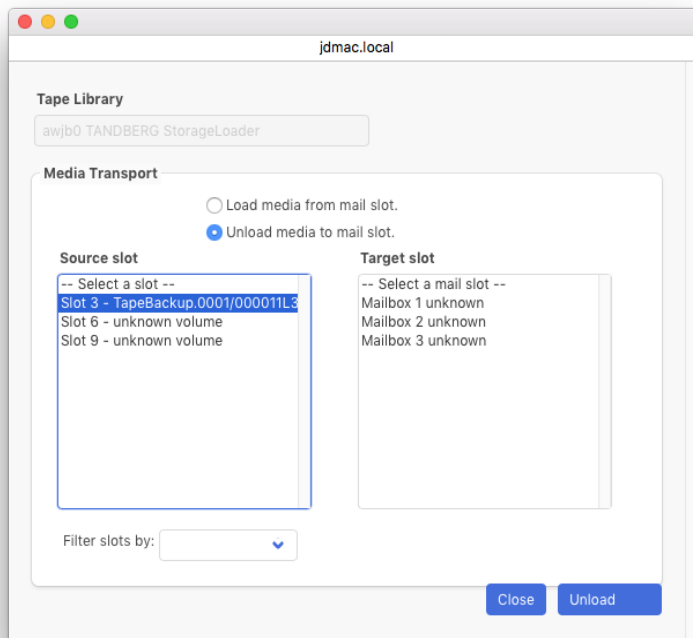
If you remove media from the library or replace it for off-site storage, or the exchange of cycles, you have to run an inventory for the library in order to make the changes known to P5. By default, all the slots of the jukebox are listed in the slot selection field of the inventory window. You can take only a partial inventory by choosing the desired slot range.

Users/operators often forget to run an inventory to update the P5 database after a manual exchange of the media. Any operations relying on a correct inventory would in such a case fail since the expected media would not be found in the given slots. P5 automates the process of updating the database inventory of the library to avoid premature termination of jobs. An inventory of the configured kind (bar code or mount) will be done the moment an unexpected condition has been detected in the library. P5 therefore attempts an automatic error recovery on behalf of the user.

1.1.1.1 Load/Unload Slot...

This function is available for tape libraries.

Media libraries or changers are often closed systems. To enable removing and replacing media, libraries are equipped with so called mail slots. They enable the user to replace or remove a media from the library without opening the library's door.



Load media from mail slot

The media from the mail slot will be transported into the chosen Target slot. If there is no bar code on the media or if the library is not equipped with a bar code reader, the check box *Inventory loaded media* should be selected to update the volume and library databases. When a bar code and a bar code reader are available, this operation will be done automatically.

Unload media to mail slot

The media chosen in the Source slot field will be transported into the mail slot for removal. The volume database will be updated accordingly.

6.2.6.3 Import Volume...

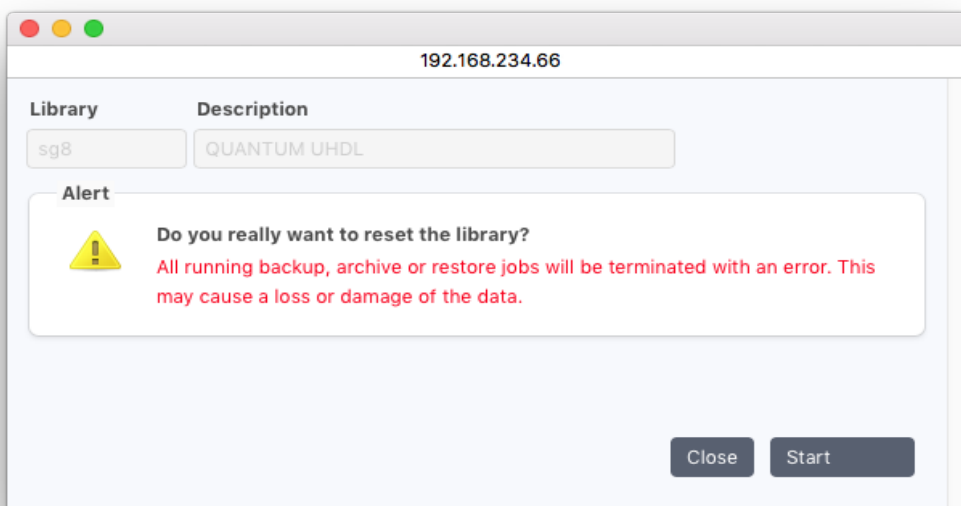
The import function is available for tape drives and removable disk drives.

This option imports into your database a P5 volume that has been written using some other P5 installation, e.g. from a subsidiary. A separate index will be created for imported **archive** volumes from other hosts. Imported **backup** volumes are added as separate sections in backup indexes. In both cases, the volumes are added to a separate media pool.

6.2.6.4 Reset Jukebox...

This option is available for tape libraries. It resets the selected library.

Before resetting a library, make sure that no jobs are using this library or its drives, since these jobs will be interrupted by the reset, which can lead to a loss of data.



6.3 The Media Pool

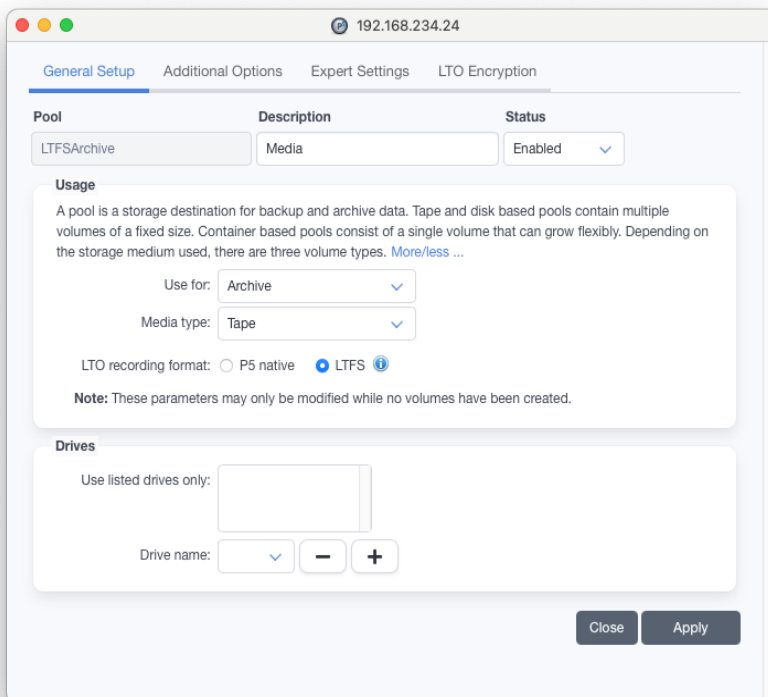
The media pool can be viewed as a large storage where the capacity is limited by the number of volumes. If the capacity needs to be increased, volumes can be added by labeling media. Container pools are potentially unlimited in size, depending only on the underlying physical storage.

A further advantage of pools is their controlling and channeling function. It enables you to use only chosen drives by defining them in the pool or to limit the used slots by having media of the corresponding pool occupy the given slots.

When storing data, the P5 media manager will give the job executor only valid volumes according to all the constraints defined by the pool. This way, data can be backed up only on volumes in defined slots and drives of a specific changer.

6.3.1 Creating a Pool

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the Archive or Backup module from the top menu and select Pools.
3. Select New to create a new pool.
4. Fill in the required fields of the Media Pool Manager's dialog window:



Pool Description:

Enter the desired pool name. The following is checked:

1. Uniqueness: No other pool may carry the same name.
2. Validity: The name is at most 32 characters long and can be composed of [A-Z][a-z][0-9]-_: . White space is not allowed.

Pool Status:

The status flag defines the availability of this pool. The possible states are:

- Enabled: The pool is available for use by the system.
Disabled: The pool will not be used by the system.

Use for:

This parameter restricts the pool to be used either for backup or for archive, which prevents unintentional mixing of data. The usage may be for:

- Backup: The pool will be used to backup data.
Archive: The pool will be used to archive data.

Media type:

Determines which media type comprises this pool. The possible types are:

- Tape: Backup and archive to tapes.
Disk: Backup and archive to removable disks and virtual tape libraries.
Container: Backup and archive to disk and cloud storage.
 Container storage should be created through the storage manager in order to create a pool, drive and a container storage together in one step.

Tape / Disk Pool - Use listed drives only:

This list directs P5 to explicitly use only the given drives for backup and restore. Depending on the drive parallelism, this can either restrict the number of used drives, when less drives are given than the allowed drive parallelism, or be used as a choice list, when the number of listed devices exceeds the drive parallelism limit. If this list is empty, any available device of the given media type can be used.

Tape / Disk Pool - Drive name:

The pull-down menu lists all the configured devices.

- + The selected device will be added to the list.
- The selected device will be removed from the list.

Tape / Disk Pool - Clone Data:

Check this box to create an exact copy of the media during the writing process. Cloning requires an additional drive as both clone tapes must be written at the same time and both media contain the same data. It is highly recommended to use the same device types when cloning because the bigger media gets closed when the size limit of the smaller media is reached. Media can be recovered from their exact clone using `Regenerate Volume` from the Gear Menu in the `Volume Manager`. Cloning is useful for Archive media pools to keep a backup copy of archived data.

Container Pool – Assign to Cloud Service:

Check this box to save the data of this media pool in the cloud. Data written to this media pool is then stored on a cloud service. See section 6.2.3, Container Storage for how to define access to a cloud server.

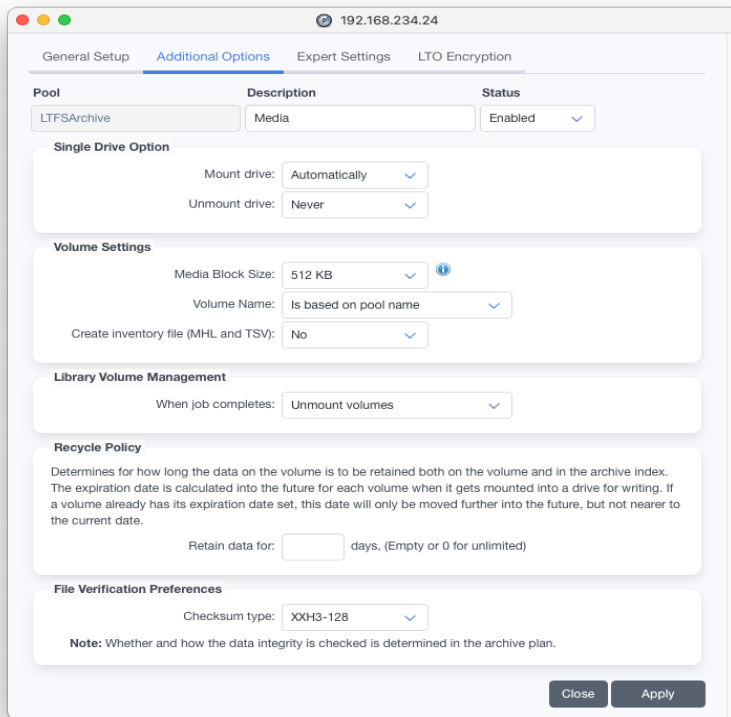
Container Pool – Cloud Service:

Select here the cloud service to connect to. See section 6.2.3, Container Storage for how to define access to a cloud service.

Container Pool – Volume Directory:

Here, the local directory is set where the data shall be stored. When the pool is bound to a cloud storage without local storage, the directory is used as staging area. Do not change the directory in case the storage is in use already, excepts when the storage has been moved on disk.

6.3.2 Tape/Disk Pool : Additional Options



Single Drive Options:

Mount Drive:

Before P5 can use a volume from a given pool in a single drive (not part of a library), the volume has to be mounted into a drive. Mounting means reserving the given drive for use by a process. A mounted drive is locked against use by other tasks. This is how a drive is handled in a multi-user environment. Unmounting releases the drive.

Manually: Mounting a volume in a single drive (device) must be done manually.
Automatically: Each of the drives in the list of 'Use Listed Drives Only' will be automatically mounted before a backup or archive using this pool is done.

Unmount drive:

After completing a job, the tape can be unmounted or ejected. This flag affects the behavior of single tape devices, library devices are not affected.

Never: The tape is never ejected.
When full: The tape is ejected if full after a job completion
Always: The tape is always ejected.

Media Block Size:

The media block size defines the block size P5 uses to write to tape. The default size of 64 KB is optimal for most hardware setups. In some cases, namely when using SAS connected drives, a bigger block size may increase the read and write speed.

Volume Name:

During labeling, the volumes are given a unique name identifying the physical tape. By default, the pool name and an incremented number is used to name the volume. As an alternative, this menu allows to select the barcode label as volume name instead. This may make it simpler to identify the tape by its barcode sticker in bigger tape sets.

Library Volume Management - When job completes:

Select whether a volume should be unmounted and transported back to its slot or kept in the drive after using it.

Time of labeling:

This parameter controls when a recyclable volume will be relabeled and made available for new data.
When volumes are needed:

Volumes whose expiration dates have been reached will be relabeled only when there are no more appendable volumes that can be written. Data of old backup cycles will be retained as long as possible despite the recyclable status.

When expiration date is reached:

Recyclable volumes will be relabeled and become available for writing immediately when a need for new volumes arises, even though further appendable volumes are available. Data of old backup cycles are therefore erased as soon as possible.

To confirm your input, click Apply. In case all the entered data is correct, the pool will be updated.

File Verification Preferences - Checksum type:

This setting allows to select the checksum type that is calculated while archiving. The type XXH3-128 is the fastest, this may make a difference while archiving data from a host with a slow CPU.

Note that on a P5 format pool, the checksum is applied on tape blocks containing the file with all its attributes, so that checksum is different to the checksum created for a file on disk.

In LTFS format pools, the checksum is the same as the checksum for the same file on disk and is stored as extended attribute on the file during restore.

6.3.3 Tape or VTL Pool: Concurrency Options

This window allows to configure the parallelisms in P5. P5 attempts to channel multiple data sources to each tape drive in order to achieve a high performance. This can be configured separately for each pool. The following window allows to define how jobs and tape drives may be used in parallel.



Tasks per stream:

The number of tasks whose data will be multiplexed into a single data stream. A task represents the data received from a single client. The default value of zero stands for no limit.

Streams per pool:

This value defines how many streams in total are allowed for this pool. A stream is a logical unit that receives data from tasks, multiplexes the data and distributes it among the available tapes.

Drives per stream:

This value defines the number of drives that will be assigned to each of the data streams.

Note: the data within one stream is spread over the given drives. Therefore, the data of a single client will be written across multiple tapes (and possibly mixed with the data of other clients). However, the data of a single file will always be located on a single media. The optimal setting of these three parameters depends on the applied hardware and network topology. There is no "universal or fastest" setting. The default combination 0-1-1 has proven to be a good setup for most cases.

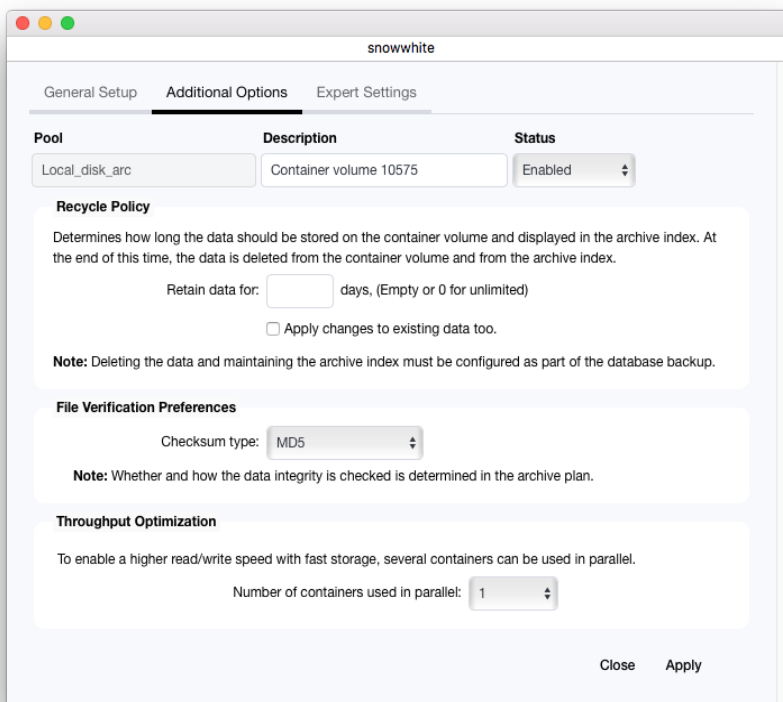
LTFS Settings – Import mode:

Define here whether during import of LTFS tapes only the last recent state of a tape or whether also older versions (of otherwise overwritten files) shall be imported.

LTFS Settings – Disable media spanning:

Select Yes here when file spanning over multiple tapes shall be disabled. A file bigger than the available space on the tape will then be written to a new tape. Note that when media spanning is disabled, some space on the tapes will be wasted.

Container Pool: Options



Recycle Policy – Retain Data for:

It is possible to change the retain time for archived data, the default is 0 for unlimited. This allows removing old data from an archive container storage, cthe check box allows to apply the changed retain time also to existing data.

File Verification Preferences – Checksum Type:

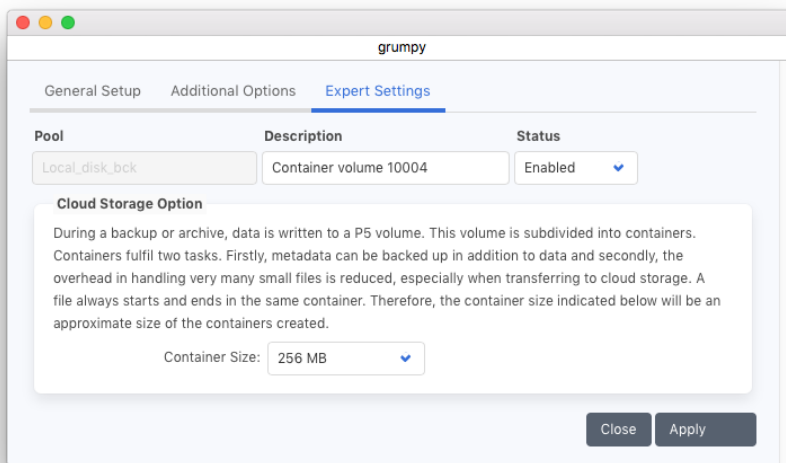
Define here which checksum tape shall be used during file verification. XXH3 is calculated fastest.

Throughput Optimization – Number of containers used in parallel:

it is possible to write to the container storage with multiple threads in parallel. This may make sense in case the container storage resides on a system that is optimized for parallel operations like for instance RAID disks.

Note that this setting may optimize the container storage only, for a higher total throughput the complete process during a job must be regarded.

6.3.4 Container Pool: Expert Settings



Container Size:

In a container storage, files are stored in containers. These contain besides the actual files also its attributes. Define here the size of data containers. Files smaller than the container size are combined to one container during backup or Archive. Bigger files will however be stored in bigger containers.

6.3.5 Pool Usage

Backup or Archive to Cloud

When assigning a media pool to a cloud service, a container storage (see section 6.2.3, Container Storage) as disk storage must be defined that holds virtual volumes as local copy for the data.

When the check mark *Keep local copy as clone* is not checked, the virtual jukebox is used as a staging area and will require approximately 0.1% of the defined space on disk.

During operation, the staging area may require more disk space, up to the size of the data that is to be saved or restored.

In case a cleanup of local data is necessary and does not happen automatically, check in the Volumes section for volumes that are marked as out of sync and run the Re-Sync function from the context menu.

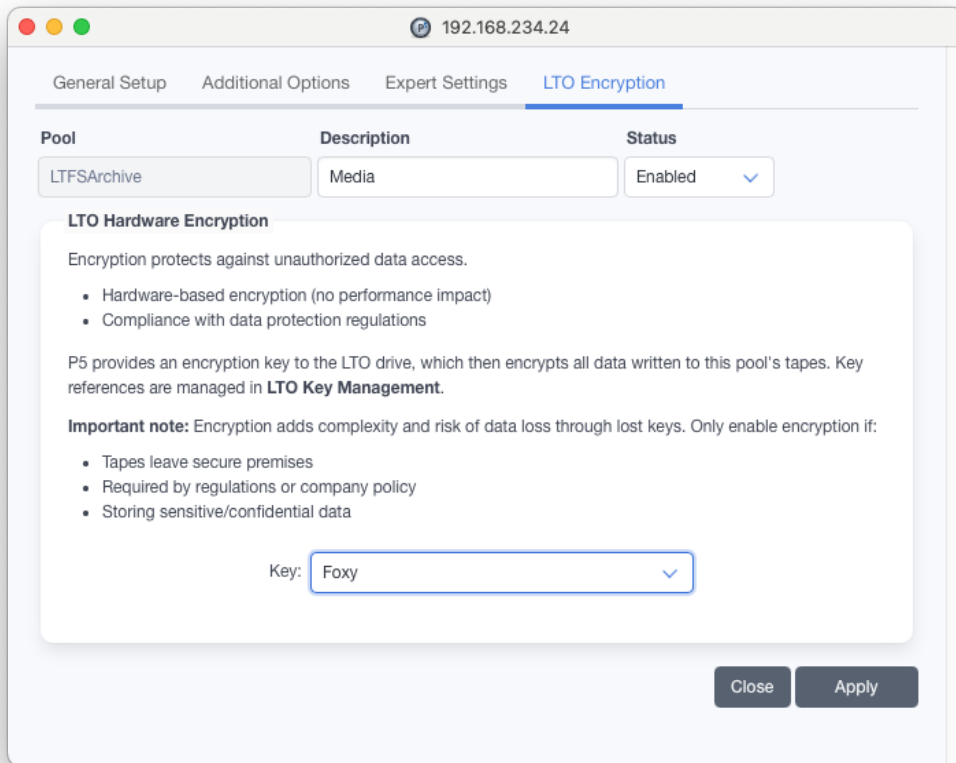
Archive to LTFS Tapes

As of version 6.1, it is possible to select LTFS (Linear Tape File System) tapes for archiving. To archive to LTFS tapes, a pool of type LTFS must be created and tapes must be labeled into it. These tapes are, except for the LTFS format, handled the same as archive tapes in the regular P5 format. However the tape cloning function is not supported for tapes written in LTFS format. Note that LTFS tapes must be labeled before the job is started, it is not possible later on to interrupt the job in order to label the tape. For more info regarding LTFS archives, see section 6.5.3, Archive to LTFS Tapes.

6.3.6 Tape Based Cloud Pool: LTO Encryption

Starting with version 8, P5 supports in addition to its client-related storage encryption also the hardware encryption implemented in LTO drives.

For the hardware encryption key management, see section 6.17, LTO Encryption Key Management.



Select her the key that shall be used for hardware encryption.
 P5 will then use this key for all read- and write operations in this media pool.



PLEASE STORE THIS KEY AND KEY NAME SECURELY OUTSIDE OF P5.
 In case the key is lost, the data on tape cannot be restored anymore.
 Unlike the configuration data, this key is not automatically saved on tape.

6.4 Volumes

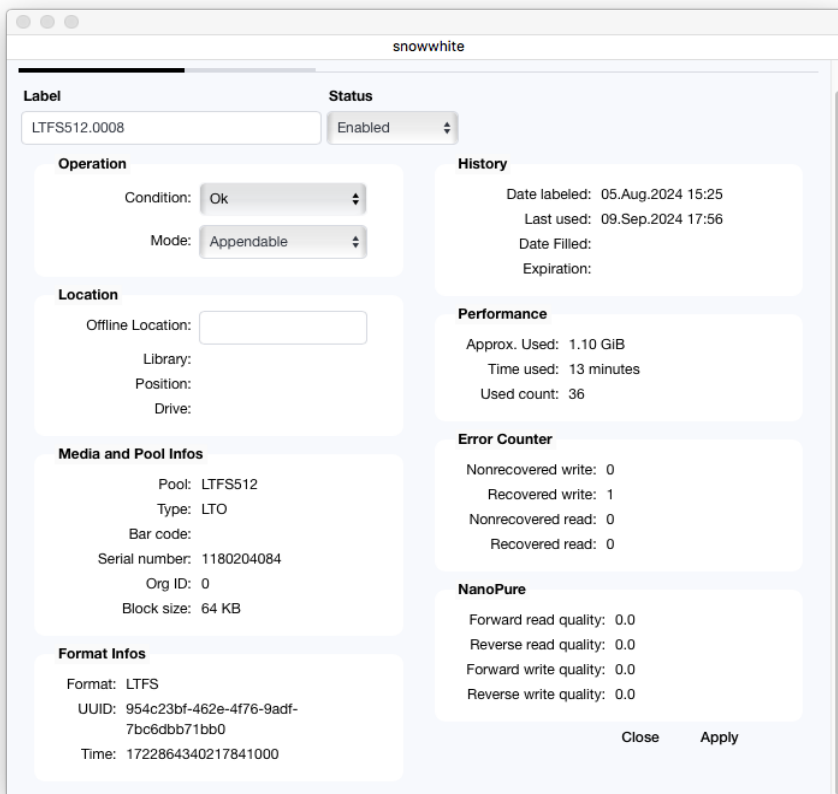
Backup is usually done on mass media such as tapes. P5 manages media as volumes. Before a media can be used in P5, it must be labeled. Labeling involves physically writing a label on the media and adding a volume into the configuration database. The label and the volume both carry diverse configuration information. A name and a sequence number uniquely identify each volume with the name of the pool it belongs to.

6.4.1 Creating a Volume (Labeling)

A volume is created automatically during labeling. The labeling function is available in the device and in the jukebox configuration areas. See section 10.1, Labeling Media.

6.4.2 Modifying a Volume

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the Archive or Backup module from the top menu and select Volumes.
3. Fill in the required fields of the Volume Manager 's dialog window:



Volume Label:

The volume name is automatically generated during labeling. It consists of the pool name and a sequence number.

Status:

The status flag defines the availability of this volume. The possible states are:

Enabled: The volume is available for use by the system.

Disabled: The volume will not be used by the system.

Usage:

Usage advises if a volume can be safely used for writing or if further writing should be avoided on this volume. In case non-recoverable read or write errors occur, 'Condition' is set to 'Suspect'.

The possible values are:

Ok The volume is fully functional.

Suspect: The volume will not be used for further writing due to detected malfunctions or irregularities.

Mode:

The mode tells which operations are allowed with this volume:

Appendable: The volume can be read and data can be appended to it.

Readonly: The volume may only be read.

Recyclable: If new volumes are required, this volume may automatically be re-labeled and reused.

Closed: The volume has been closed following a backup or an archive and will not be used any further.

Full: The volume is full.

Media and Pool Info

Pool Name of the pool the volume belongs to.

Type Type of the media.

Kind Further specifies the media (for tapes only).

Bar code In case the changer has a bar code reader, the bar code label attached to the media will be scanned and stored in this field. This enables quick identification of media replaced in the changer.

Org ID In case the volume was imported from another P5 system, this field contains the original volume ID from that system.

Location

Offline Location	If a description is given, it is shown as location when the tape is offline (not in the library).
Location	Description of the jukebox in which the volume is currently located.
Jukebox	Name of the jukebox in which the volume is currently located.
Position	Slot or drive number in which the volume is located within the jukebox.
Drive	If the volume is located in a drive, this holds the drive's name.

History

Date labeled	Date when the media was labeled.
Last used	Date the volume was used last.
Date filled	Date when the volume was filled.
Expiration	The pool defines the retention time for the volumes. This value is calculated when the backup is started and shows the earliest date when the volume may be recycled. However, recycling may be prevented due to some files not yet been saved for a second time.

Performance

Used size	The number of kilobytes of payload data stored on the volume.
Time used	The cumulative amount of time the volume was loaded and used in a drive.
Used count	The number of times the volume has been mounted into a drive.

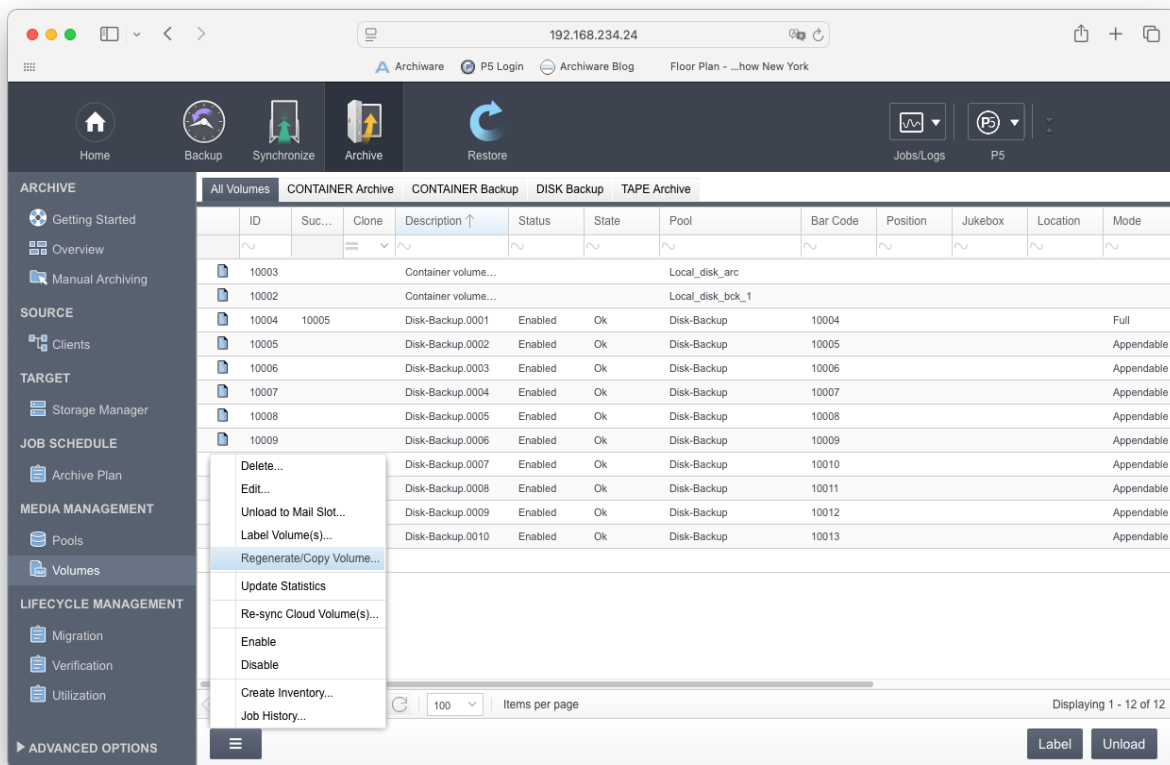
Error Counter

Non-recovered write	The number of write errors that could not be corrected. Warning: This signals that this volume is significantly damaged. It is recommended not to use this volume for further writing. The volume's condition field will automatically be switched to 'Suspect'. This prevents the system from using it for writing. If more than one volume shows this error in short sequence, the drive should be inspected.
Recovered write	The number of writing errors that were successfully corrected.
Non-recovered read	The number of reading errors that could not be corrected. Warning: See 'Non-recovered write'
Recovered read	The number of read errors that were successfully corrected.
<i>NanoPure</i>	This section shows the media quality, as read from the cartridge chip when NanoPure Media are used.

To confirm your input, click Apply. In case all the entered data is correct, the volume will be updated.

6.4.3 Volume Manager: Gear Menu

For each of the configured volumes, there are several options you can access from the Gear menu. Alternatively, you can right-click on a volume entry in the main tab to access available options.



6.4.3.1 Edit ...

Click here to get the volume's detail window with editing capability.

6.4.3.2 Unload to Mail Slot...

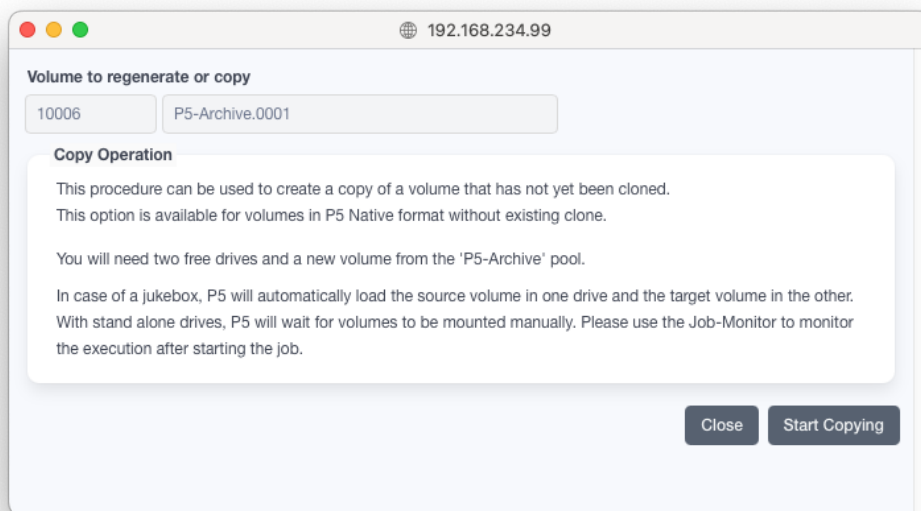
This function is available only where the selected volume reside in a tape library with addressable mail slots. The function allows the selected volume to be moved into a mail slot in the library.

6.4.3.3 Regenerate or Copy Volume...

This option allows you to recreate a clone volume of a backup or an archive pool and also create a clone of a volume where a clone does not already exist. To start this process use the context or the gear menu.

To regenerate a volume from its clone, the volume must have been written with cloning enabled, and the clone volume must be made available. Use this function when one of the two clones gets damaged.

In case the volume is not already a member of a clone pair, but a tape volume in P5 format, this methods allows to create a copy of the volume's data to another volume. The target volume must be a new, empty tape, labelled for the same pool as the original volume. The copy process then creates a clone tape of the original one.



Before starting a volume regeneration from a clone, please stop all archive jobs that write to the same media pool, then navigate to the affected pool and turn off the cloning option before starting the regenerate job.

See section 6.3.1 Creating a Pool for where to set the cloning option.

Both, the regeneration from clone and the creation of a clone tape requires two tape drives. It is recommended to use the same LTO generation for the clone tape as for the original tape in order not to waste tape space.

6.4.3.4 Re-sync Cloud Volume ...

This function is available for cloud-volumes only, that may either be container volumes stored in a cloud storage or VTL disk volumes associated with a cloud storage. Use this function when a volume is out-of-sync between the local and the cloud storage. The volume is then marked with a yellow sign in the Volume Manager. This function re-synchronizes both instances of the volume, the local one and the one in the cloud. This function is also capable to repair cloud related container volumes.

6.4.3.5 Enable

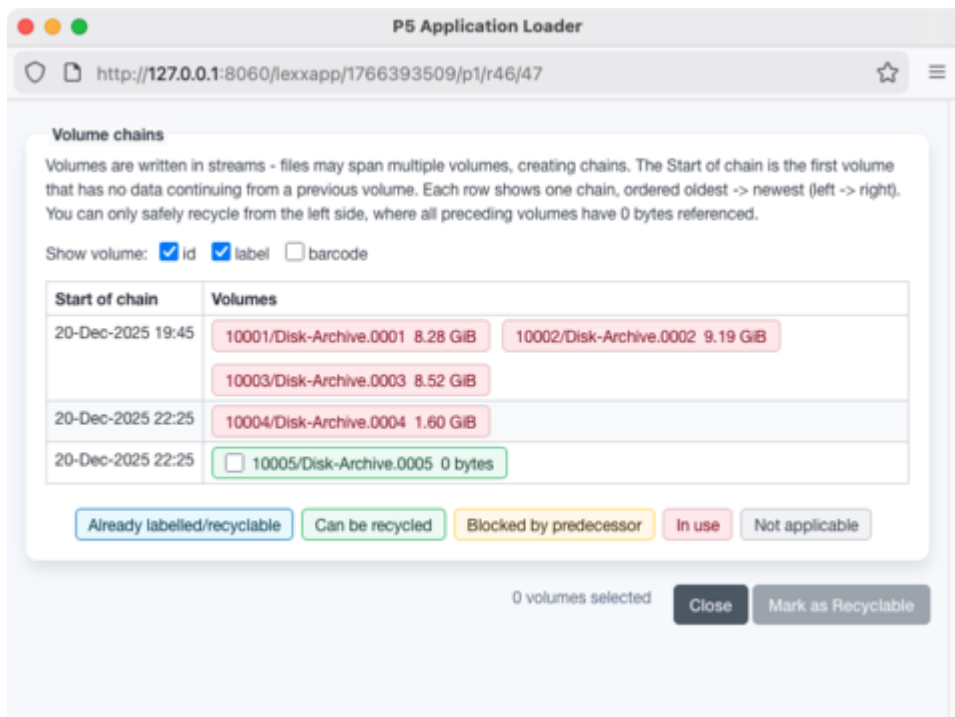
Click here to enable a previously disabled volume again.

6.4.3.6 Disable

Set the volume to disabled. The volume will remain unchanged but is not used anymore by any job.

6.4.3.7 Recycle Volume(s) ...

This method is not applicable for container/cloud storage volumes. It helps deciding whether an archive volume can be recycled, for instance after deleting folders from the index.



Volumes states:

Already recycled or empty	Nothing to do, this volume is usable
Can be recycled	This volume can be set to recyclable, it does not contain files anymore. A later archive job to this pool may then recycle and reuse the volume.
Blocked by predecessor	This volume may appear as empty, but it may contain a partial file that begins on the volumes predecessor
in use	This volume contains files that are referred to in the index
Not applicable	This is a container/cloud volume or a backup volume

6.4.3.8 Label Volume(s) ...

This option allows the relabeling of a tape volume, and thus erase an existing disk volume or an existing tape volume in a tape library. When using a single tape drive, please open the tape drive section in the storage manager to do so, see 6.2.2, Single Tape Drives. For a full description of the labeling process, please refer to section 10.1, Labeling Media

Labeling a media will erase it's contents and delete the contents from the according index. The volume can then be assigned to a new (or the same) media pool.

6.4.3.9 Delete Volume(s) ...

Volume deletion means it is removed from the volumes table, its data is removed from the according index. The physical volume is from then on treated as unknown. The volume itself is not accesses nor erased.

6.4.3.10 Create inventory ...

This function exports a volume inventory into an TSV (tab-separated) file. The file can then be read with a normal ascii-editor or imported into, for instance, Libre-Office Calc or MS Excel.

6.4.3.11 Job History

The Volume Job History shows a list of all jobs that affected the particular volume since it has been labeled. This way an overview is given of what is on the tape.

Note that single jobs can write to multiple tapes, so it cannot be assumed that the complete data written by the named job resides on this tape.

192.168.234.197							
Plan Name/Operation	Client	Level	Size	Start Time	End Time	Duration	Job
Backup Plan	localhost	incr		11-Nov 16:33	11-Nov 16:34	42 secs	10040
Save backup index 'localhost' to pool(s) 'Disk-Backup'	localhost			11-Nov 16:33	11-Nov 16:33	4 secs	10041
Backup Plan	localhost	incr		07-Nov 22:30	07-Nov 22:32	2 mins 3 secs	10028
Save backup index 'localhost' to pool(s) 'Disk-Backup'	localhost			07-Nov 22:31	07-Nov 22:31	30 secs	10029
Backup Plan	localhost	incr	190.00 KB	07-Nov 17:57	07-Nov 17:57	39 secs	10023
Save backup index 'localhost' to pool(s) 'Disk-Backup'	localhost			07-Nov 17:57	07-Nov 17:57	8 secs	10024
Labeling for pool Disk-Backup	localhost			07-Nov 17:55	07-Nov 17:55	1 sec	10022

Note: Job history logged beginning with P4.2

6.4.3.12 Update Statistics

This function checks and corrects (if required) the volume statistics, that is the amount of written data on the tape. Volumes where this is required are marked in the volume table with a blue circle. Calling this function is required only in case something went wrong during a job that wrote to this volume.

6.5 Archive Plans

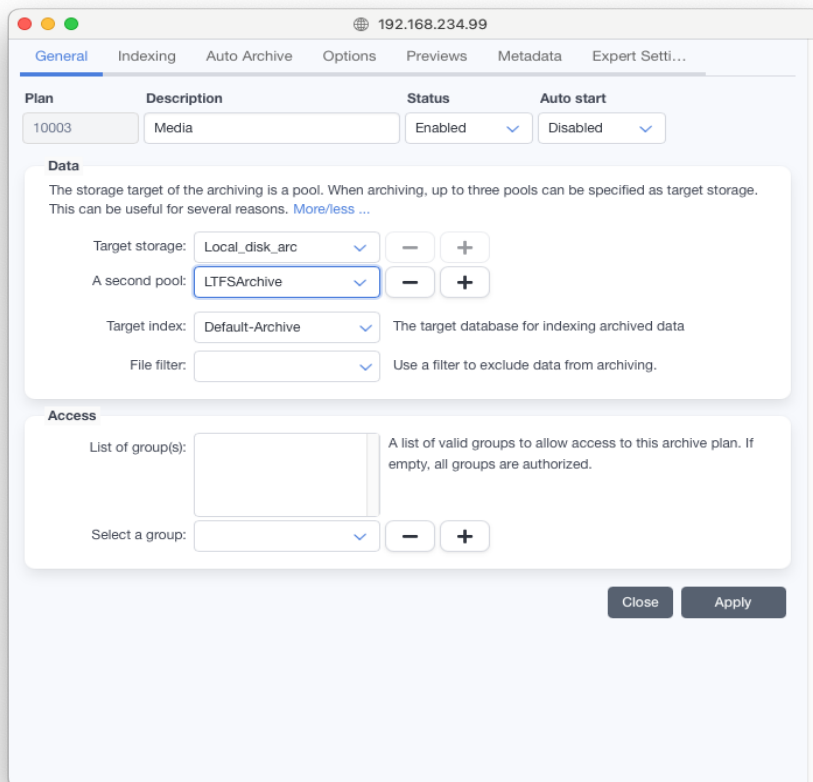
Data from the P5 server or one of its clients are archived by means of an archive plan. To archive data select one or more directories from the manual archiving area and submit them for archiving by using the context sensitive menu shown when clicking the right mouse button or Archive from the bottom bar. When initiating the archiving, you may choose from all the available archiving plans.

The purpose of the archive plan is to hold the common criteria for a specific archiving operation. The two important parameters are: The target 'Pool' where to archive the data, and the target 'Index' where to keep the file list.

Alternatively, data can be submitted for automatic archiving by copying it to the drop folder of the desired Archive Plan.

6.5.1 Creating an Archive Plan

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the Archive module from the top menu and select Archive Plan.
3. Choose New from the bottom bar and fill in the required fields



Description:

Enter the desired plan name here. The name must be unique

Status:

The status flag defines the availability of this plan. The possible states are:

Enabled: The plan is available for use by the system.

Disabled: The plan will not be used by the system.

Target storage:

The pool defines which group of volumes the data will be saved onto.

See section 6.3, The Media Pool.

A second pool:

Up to three target pools can be defined. P5 will write to the defined pools by creating multiple archive jobs, one for each target pool. Whether or not these jobs run in parallel depends on the availability of the required resources.



Note: By writing to different pools, P5 creates multiple independent copies of the data. In order to create clone tapes in the same pool, see section 6.3 The Media Pool. Cloning however requires two equal tape drives.

Target index:

Every time data is archived P5 creates a list of the archived files and folders. Using this list, it mirrors the folder structure of the on-line storage in the index. You can add new indexes from the 'General Setup' administration area, the 'Database' resource area. See section 6.12, Index Databases.

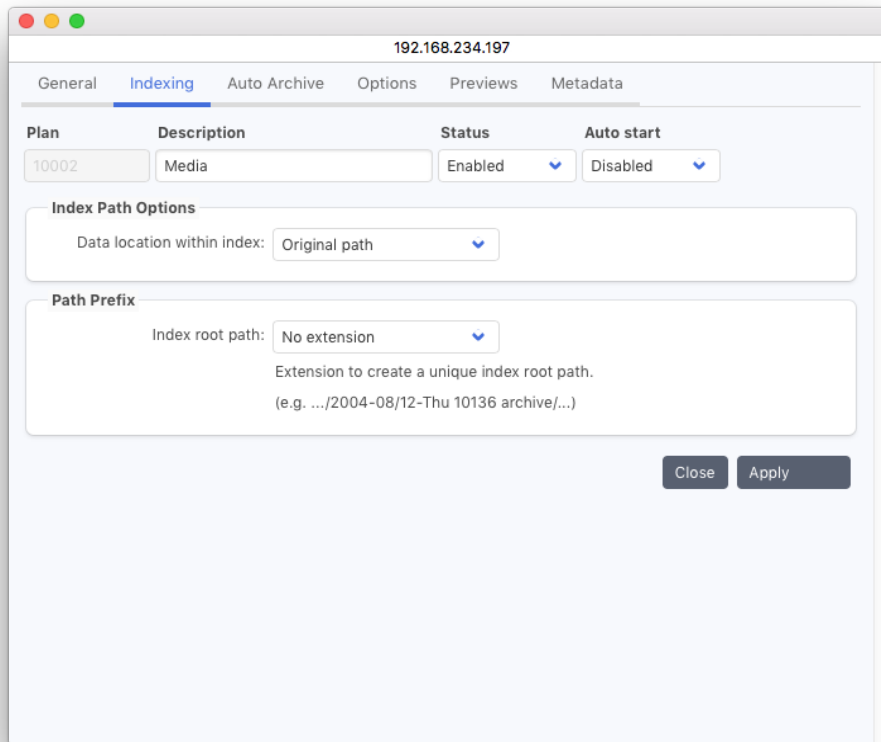
File filter:

If you wish to filter out certain data from the storage process, you can choose the appropriate filter here. See section 6.9, Filters.

Access:

Only users who are members of the given groups can use this archive plan to submit data for archiving. If no group is given, any user who is allowed to archive can use this plan. Groups are maintained in P5, see section 6.15, Group Preferences.

Index Options



When archived, each file will be entered into the index along with its directory path.

Data location within index:

Each archived file and its directory path have a corresponding index entry. To enable you to structurally organize your data, you can define your own hierarchical structure in the index.

There are three possibilities for this setting:

Original Path:

The path from disk is copied 1:1 into the index

Predetermined for user: Absolute Path:

A single path is defined under which all the data archived with this plan will be kept in the index. The user cannot change this path or choose an alternative when archiving.

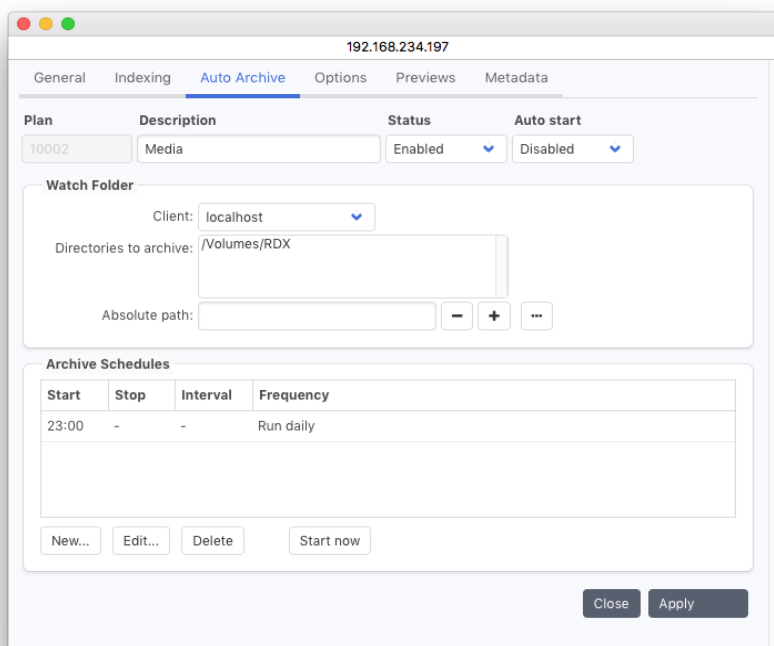
User selectable: Suggested Locations:

A list of paths is suggested to the user when starting an archive. The user can select under which of these paths the data is to be stored. A path not appearing in the list of suggestions can also be entered.

Index Root path:

By default, P5 does not add an automated prefix to the path, as separate archive jobs are in any case separated by snapshots. A unique prefix may make sense in case the archive is set up using a hot folder, archiving all files to the same directory in the index. In that scenario, big index directories may be created which are slow when browsing over all snapshots.. When archiving from different source folders (usually used when manually archiving), this option is not required.

Automatic Archive Options



P5 can automatically archive data by defining a Watch Folder. This is a list of directories on a P5 client where the data is to be found. By means of filters, one can further specify which files in the given directories to include or to exclude.

Client:

Holds the name of the P5 client from where the directories are to be archived. The P5 server is always named 'localhost'. See section 6.11, P5 Clients.

Directories to archive:

All the folders that you wish to archive have to appear here.



Note: By using filters you can selectively choose the data to be archived, see section 6.9, Filters.

Absolute path:

To add or remove a path in the list, enter it here and use the + or - button.

- + The entry will be added to the list.
- The entry will be removed from the list.

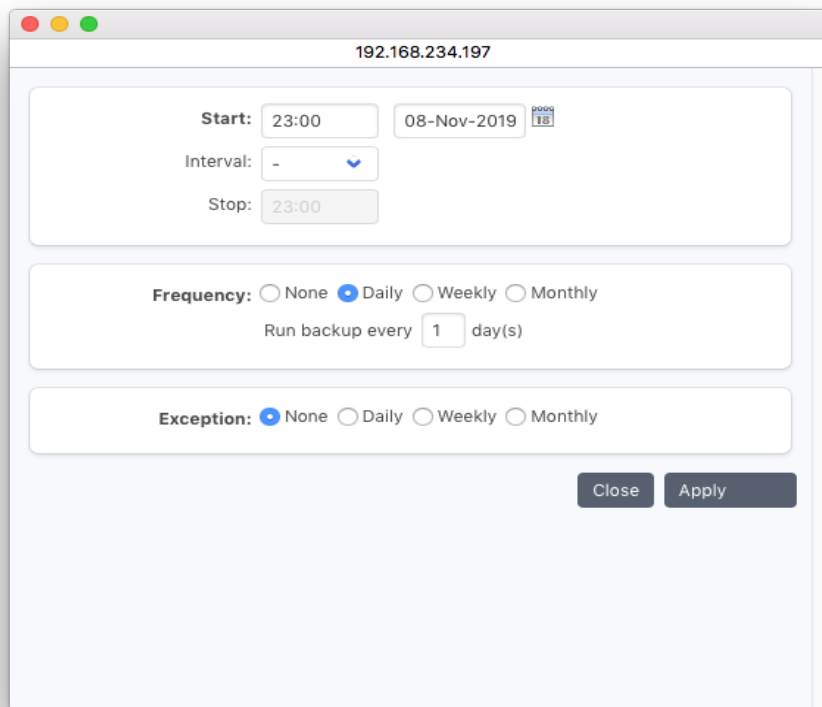
Auto start:

Enabled: The plan will start automatically at the given time.

Disabled: The plan will not start automatically.

It can be started manually by clicking "Start now".

Archive Schedules



Start:

Define the start date and time. The date defines when the plan should be run for the first time. The time defines at which time it should be run. It is possible to re-run the archive in intervals by specifying an interval and an end time. According to the frequency given below, the plan will be executed at the given time.



To ensure that only files that are not changing are archived, use a filter to select only those files that have not changed for some length of time.

Interval:

If set, the plan execution is repeated within the specified time range in the given interval. So it is possible to archive files (with a short delay) as soon as these are changed or appear in the given folder. Note that for files that are still subject to changes, this setting may cause the same file to be archived several times.

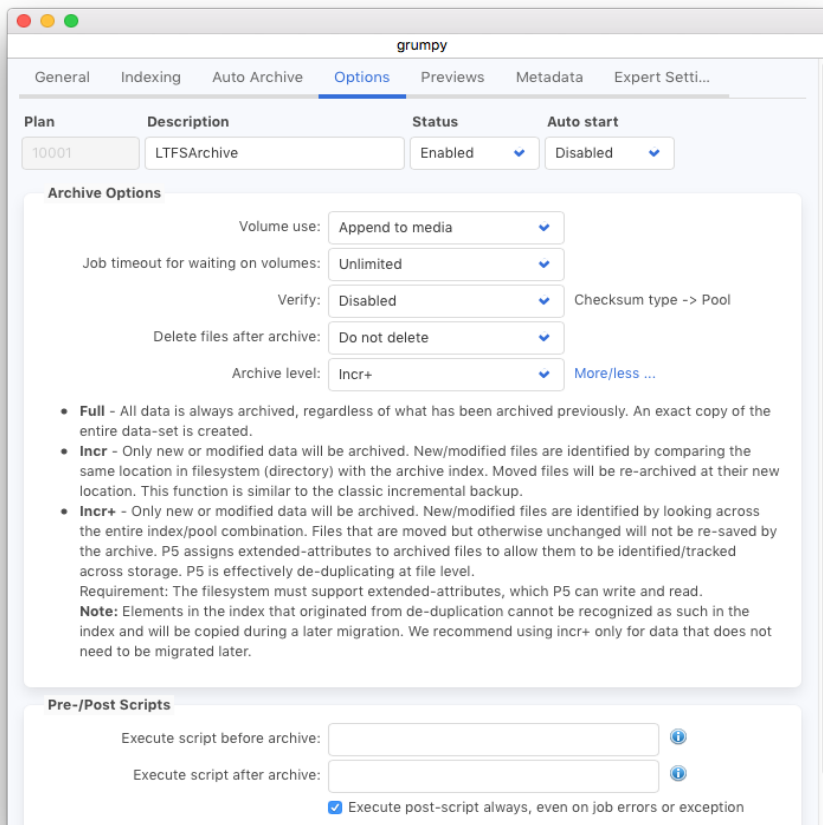
Frequency:

The frequency defines how often the plan should be executed. It is possible to set complex frequencies such as "Wednesday of every third week" or "Saturday of each month". The start date allows for complex planning such as two plans that alternate weekly.

Exception:

Exceptions can be set the same way as the frequency to disable the execution at specific times, for example if the drive is required for some other purpose at the time. Care has to be taken not to define an impossible schedule through adding an exception.

Archive Options



Volume use:

Normally the data of consecutive archiving runs are stored sequentially (Append to media) onto the volumes of the given pools. If each archive run needs to be stored on its own volume, set 'Use new media' for 'Volume use'.

Append to media: New data will be appended to the end of the media.

Use new media: Each save will start on a new media.

Job timeout for waiting on volumes:

This field controls timeout parameter for waiting on files. Timeout values are selected from the drop-down list and range from 1 hour to 4 days. Timeout can also be overridden by selecting the 'Unlimited' value.

Verify After Write

It is not required to start a separate verification of the written data since the tape drives detect errors automatically while writing anyway.

Yes: Start a verification after the job completes.

No: Do not start a verification.



When verification is enabled, all the tapes written by the archive job are required once more for verification. Removing the media will cause the job to wait for tapes.

When verification is enabled, do not run multiple archive jobs to the same media pool at the same time since that might cause a deadlock.

Delete files after archive:

Do not delete:	Keeps the files.
Delete files and folders:	Remove files and folders.
Delete files only:	Remove the files only.
Delete files and create stubs:	Remove the file but replace original file with a stub file that can be used by the P5 Companion to restore the file.

Archive Level:

Full: Archive all files even if already in the index

Incr: Archive files only in case these are not yet contained in the archive.

Incr+: While archiving, files are tagged with an extended attribute. In case a tagged file is moved elsewhere in the file system and then archived again, this mode will save a reference to the existing saved file and will not save the file's data again.



Note that if a file is explicitly selected, it is added even if it already exists in the archive and an incremental option is set.

In order for `Incr+` to work as expected, the file system where the data resides must support extended attributes.

When running incremental archives, P5 compares the file by its path. As a consequence, when the option to use an index path prefix is selected, the incremental archive **always** runs as full.

Execute script before archive:

The script or command to execute before starting the archiving.

Execute script after archive:

The script or command to use after completing the archiving.

Specify the full path to your script to execute it on the source client or prefix the path with the P5 client name separated by a colon to execute the script on a specific client.

Example:

on the source client:

```
/usr/local/scripts/db-stop
```

on a specified client:

```
dbserver:/usr/local/scripts/db-stop
```

on a windows machine:

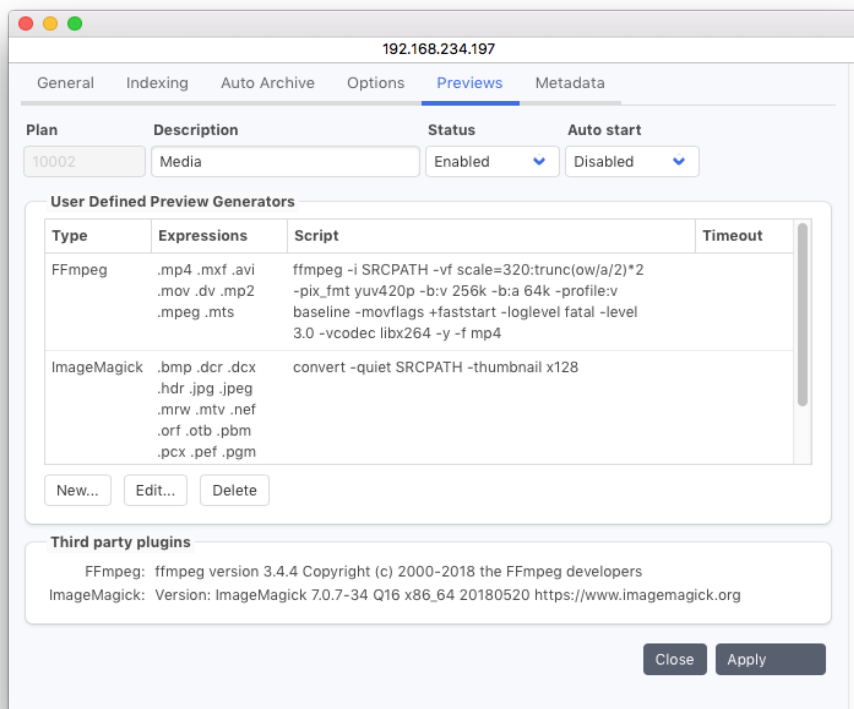
```
/c/database/scripts/db-stop.bat
```

Default Preview Generation

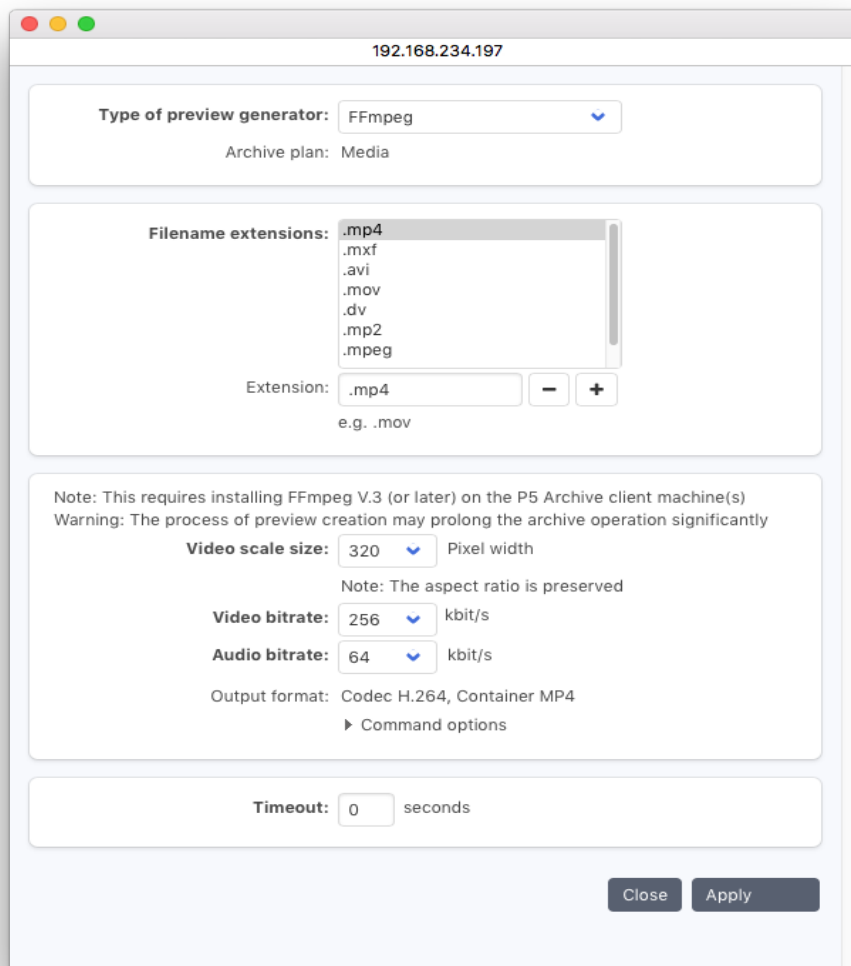
Previews can be automatically generated during archiving. Please note that previews are saved together with the archive index on disk and may require considerable space depending on the resolution.

The preview generation is not done internally, but requires external tools that are not included. For still images, an interface to ImageMagick is predefined, for movie formats an interface to FFmpeg is prepared.

Starting with version 6.1, P5 includes a build of ImageMagick and FFmpeg that cover the most common Image and movie formats, but cannot be extended. For extended formats, please install ImageMagick and FFmpeg and the required extensions. In order to use the installed versions of both programs, please make sure these can either be located via the PATH environment variable or edit the command in the P5 preview definition and add the full path to the command.



As an alternative to the default preview converters, it is possible to configure your own converter commands. For each converter definition, a case sensitive list of extensions and a command can be given. For files with one of the listed file name extensions, P5 will call the converter command. The converter is expected to accept the full file name of the image or video to convert as a parameter. It must return the full name (including path) of the generated preview via its standard out channel. Preview generators must not create other output on the standard out or standard error channel.



The predefined preview generators offer the most common parameters during setup. To fine tune the generation process, open the *Command option* and adapt the parameters to your need.

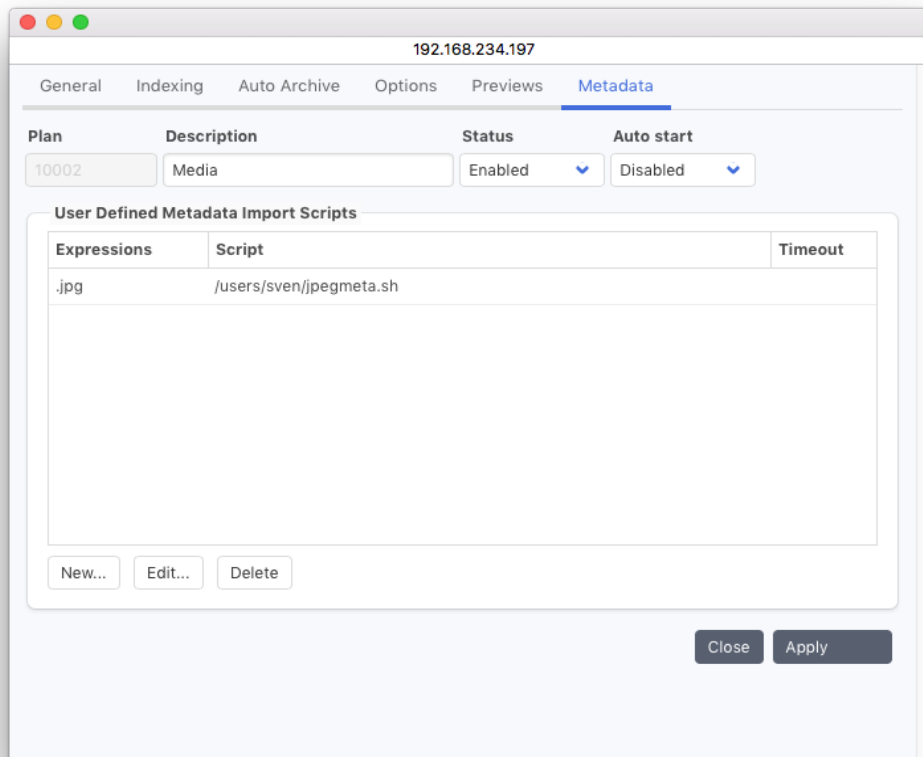
It is possible to reduce or extend the file name extension list to apply specific options to specific files only or to define different preview setting for different file type.

Both, ImageMagick and FFmpeg use plugins to support several file formats and encodings. In case your specific format is not supported, please check whether there is a plugin available for it. If none are available, you may need to define your own preview generator using a different tool.

The *timeout value* defines how long at most the preview generation of a single file may take. When the given amount of time expires, P5 will cancel execution of the command.

Automatic Metadata Import

P5 provides an interface to add metadata information to archived files using external commands.



Multiple scripts can be used depending on the file name. Click **New . . .** to add a metadata script.

The screenshot shows a web interface for configuring metadata scripts. The window title is "192.168.234.197". It features three main sections:

- Filename extensions:** A list box containing ".jpg". Below it is an "Extension:" input field with minus and plus buttons, and an example "e.g. .jpg .mov".
- Metadata import script:** An input field containing "/users/sven/jpegmeta.sh". Below it is an example command "e.g. /my/converters/jpeg -resolution 128".
- Timeout:** An input field containing "0" followed by the text "seconds".

At the bottom right, there are two buttons: "Close" and "Apply".

Enter an expression that matches the file names to which this script shall be applied. Enter the command or script name and a timeout value after which the script shall be canceled in case it does not terminate.

The script itself must accept the complete file name including path of the file to read and must echo the values as key-value-pair. Expressions containing blanks can be parenthesized in braces to read them as one value.

For example, an output line like

```
imagesize {200x400} findercomment {NAB show 2013}
```

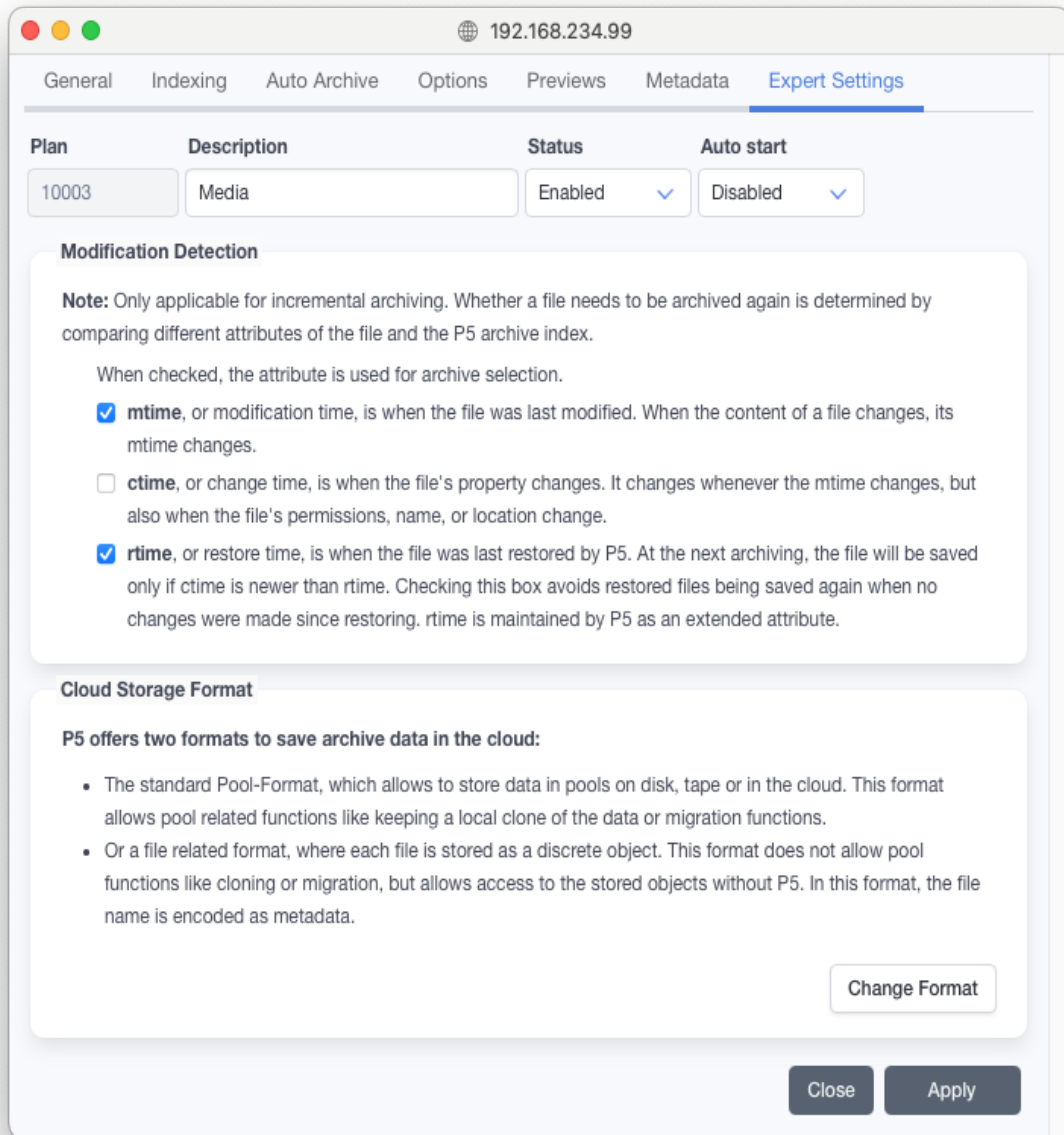
is interpreted to set the archive metadata fields `imagesize` and `findercomment`.

See section 6.12.5.1 `Fields ...` for a description of how to add your own metadata fields to an archive index database.

To confirm your input, click **Apply**. In case all the entered data is correct, the plan will be updated.

Archive Plan Expert Options

The expert options allow fine-tuning of how an incremental archive determines whether a file has changed or not.



***mtime* :**

The mtime or modification time changes when a the file content is changed. The mtime should always be compared. During restore, the mtime is reset to the original time the file was saved with.

ctime:

The ctime or inode change time changes, when a file attribute is changed, like owner or permission. It also changes when a hardlink to that file is created or removed. Disabling the ctime will make the incremental archive job ignore owner and protection changes. Note that the ctime it also set when the file is restored.

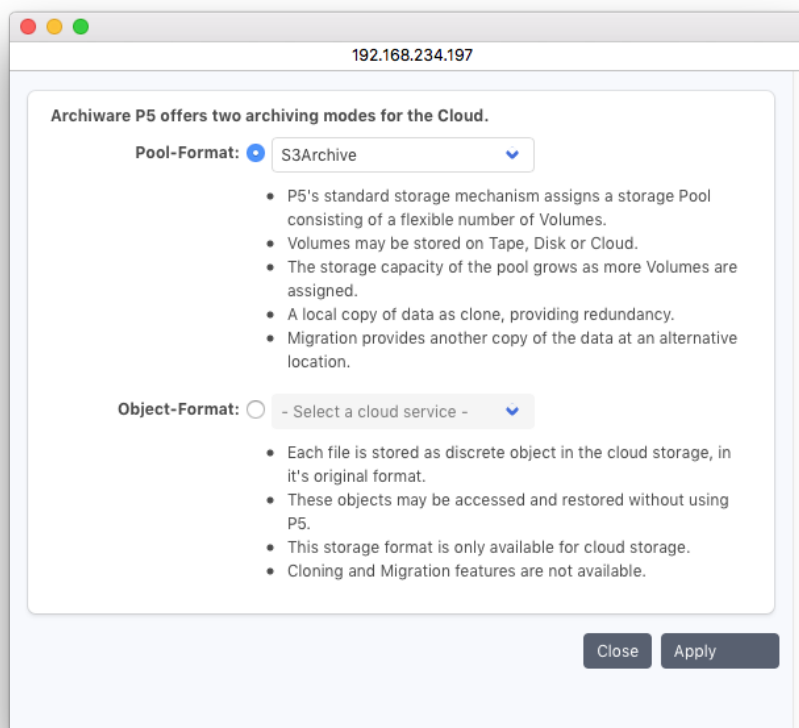
mtime:

The mtime or restore time is set by P5 when the file is restored. When this flag is set, P5 saves the file in case the file's ctime is newer than the restore time, or in other words when the file was changed after it was restored. Note that mtime comparison requires support for extended attributes on the filesystem.

Cloud Storage format:

P5 offers two formats to save archive data in the cloud:

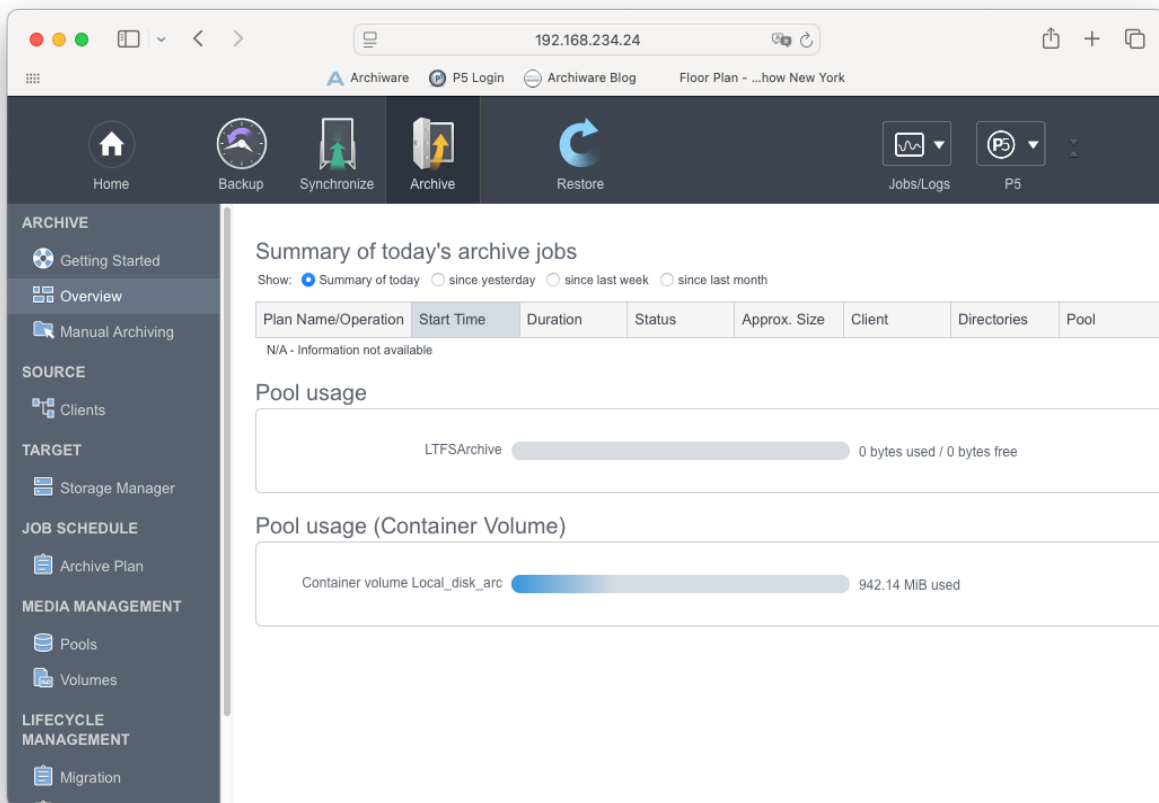
- The standard Pool-Format, which allows to store data in pools on disk, tape or in the cloud. This format allows pool related functions like keeping a local clone of the data or migration functions.
- Or a file related format, where each file is stored as a discrete object. This format does not allow pool functions like cloning or migration, but allows access to the stored objects without P5. In this format, the file name is encoded as metadata.



6.5.2 Archive Overview

The archive overview shows a summary of the latest archive jobs. It shows the status and duration together with the archived folder.

The overview window is helpful to get an idea about the archiving state at a glance.



6.5.3 Archive to LTFS Tapes

Starting with version 6.1, it is possible to archive to LTFS tapes as an alternative to the P5 proprietary format.

Some functions are not possible on LTFS tapes:

- Cloning, i.e. writing to two tapes in parallel, is not possible
- The client-based data encryption is not supported, LTO Hardware Encryption is possible, see 6.17, LTO Encryption Key Management.
- Only checksum verify, no content verify after writing
- Multiplexing, i.e. the parallel execution of multiple jobs writing to one media pool
- The configuration database is not automatically saved to LTFS tapes, that means a configuration restore from LTFS tapes is not possible

Also, LTFS writing cannot be interrupted to label new tapes. As a consequence, When archiving to LTFS, a sufficient quantity of tapes must be labeled before starting to archive, the archive job will end incomplete when it runs out of tapes.

Regarding the written files, there are differences in what can be stored on an LTFS tape:

- Owner and group attributes are ignored and not stored.
- OSX finder info and resource forks are ignored and not stored.
- Extended file attributes (xattrs) are stored when these are string attributes. Binary attributes are ignored.
- The only time stamp stored on LTFS tapes is the file modification time (mtime).

File names are stored in a translated format on LTFS tapes, several characters are translated to "percent encoding", same as for instance in URLs in a browser, for instance a question mark is translated to the sequence %3F. This affects:

- Characters above code 127, this includes all umlauts and accent characters.
- The Characters * ? < > " | \ : / and the percent character itself.
- Control characters with codes below 32.

When restoring files from an LTFS archive tape, the file names are translated back to the originals, as stored in the archive index. When reading files with other LTFS compliant software, please make sure the LTFS standard encoding is supported. In that case, control characters should be avoided in filenames.

The LTFS implementation corresponds to the

Linear Tape File System (LTFS) Format Specification Version 2.5

https://www.snia.org/sites/default/files/technical_work/LTFS/LTFS_Format_v2.5_Technical_Position.pdf

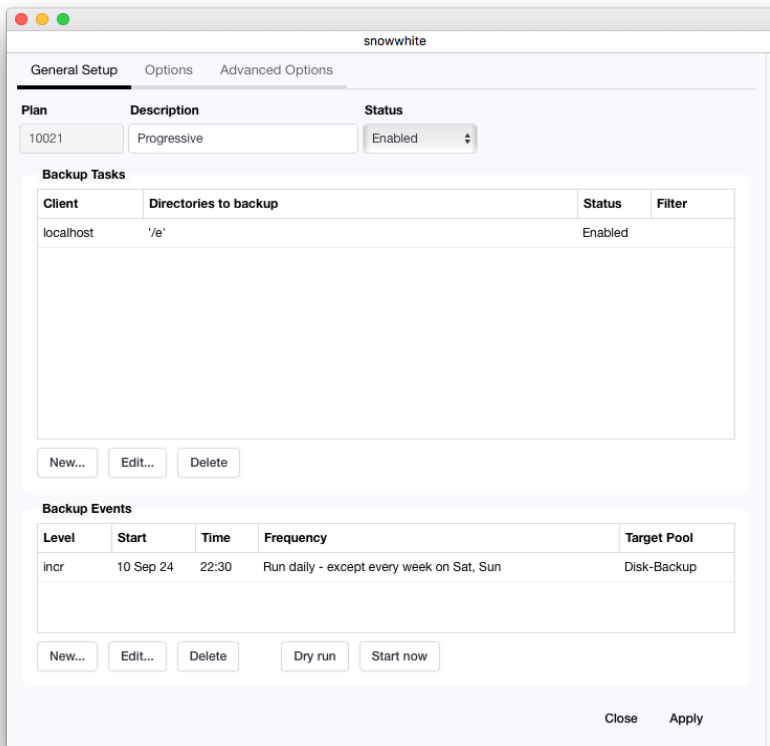
6.6 Backup Plans

Data from the P5 server or one of its clients can be backed up by means of a backup plan. To backup data, configure one or more directories in a backup plan. The target where the data will be saved is a volume pool. A backup run can save either entire directory structures or only data (files) that has been changed since the last backup. Each backup plan is executed according to the defined schedule.

A backup plan can backup several clients. Backing up each client is considered as one task of the backup plan. The client's directories to be backed up are defined in the task.

6.6.1 Creating a Backup Plan

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select Backup from the top menu and select Backup Plan.
3. Fill in the required fields.



Plan Name:

Enter the desired plan name here. No other plan may carry the same name.

Plan Status:

The status flag defines the availability of this plan. The possible states are:

Enabled: The plan is available for use by the system.

Disabled: The plan will not be used by the system.

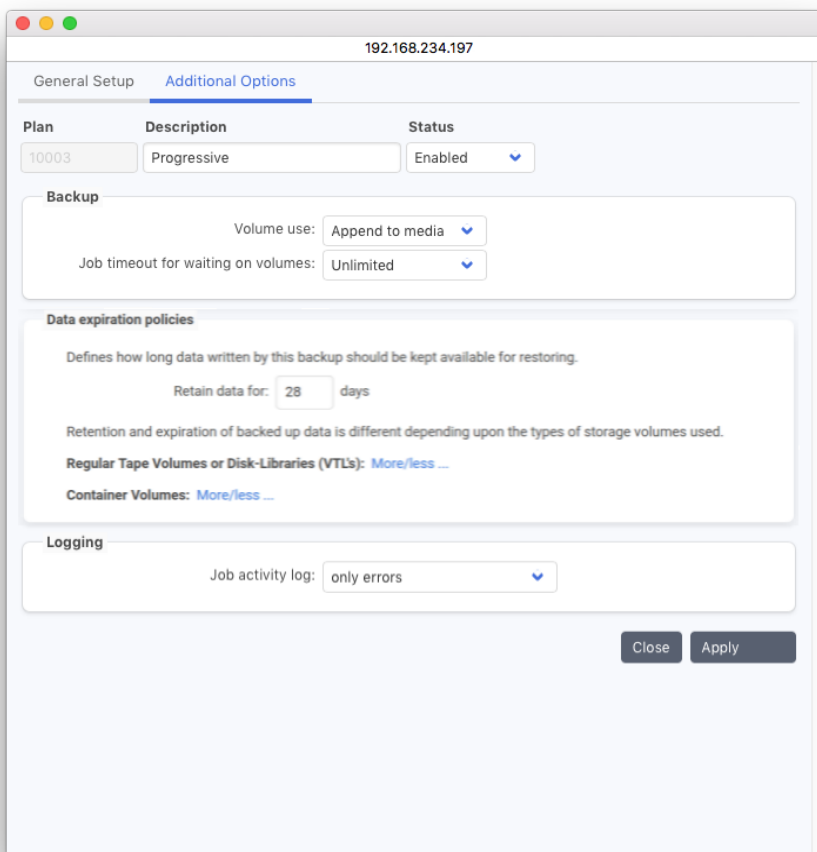
Start Buttons:

The buttons in the lower half of the screen allow starting the configured backup for one or for all clients manually:

Dry run: The plan is started but does not save any files to the storage. This option can be used for instance to check how much data would be saved.

Start now: The plan is started regular.

Backup Plan: Options



Volume use:

Normally, the data is stored sequentially on a set of volumes of a given pool (Append to media). Should each new backup cycle be done on a separate set of volumes, set this field to 'Use new media'.

Append to media: New data will be appended to the end of the media.

Use new media: Each full backup will start on a new media.

Job timeout:

By default, a backup plan will stop and wait for new volumes in case all volumes of the given pool are full. While the jobs are in such a waiting state, the user can add media to the jukebox, label media or cancel the job. In case of an unattended backup, the job may be configured not to wait for a manual intervention or to wait for a limited time only before it stops and issues an error message. This field allows to limit the time the backup job will wait.

Skip failed mount point:

In case a mount point is part of the backup and the mounted file system is not present/mounted, P5 would save an empty folder. That implies that the previous content of that folder is marked as no longer on disk in the index, and the next backup with that file system mounted again would save the entire filesystem as new. By checking this option, P5 would skip the folder instead and keep its content.

Recycle Policy, data retain time:

This value defines for how long the data on the media must be retained. When the period expires for the newest data on the volume, it is either recycled or marked for recycling, depending on the chosen recycling strategy.

Job activity Log:

By default, only errors are kept in the activity log. Select here *Errors and file operations* in order to also log each file that is added to the backup. This setting will be applied for a single job execution and reset afterwards in order to avoid huge logs.

To confirm your input, click Apply. In case all the entered data is correct, a backup plan will be created.

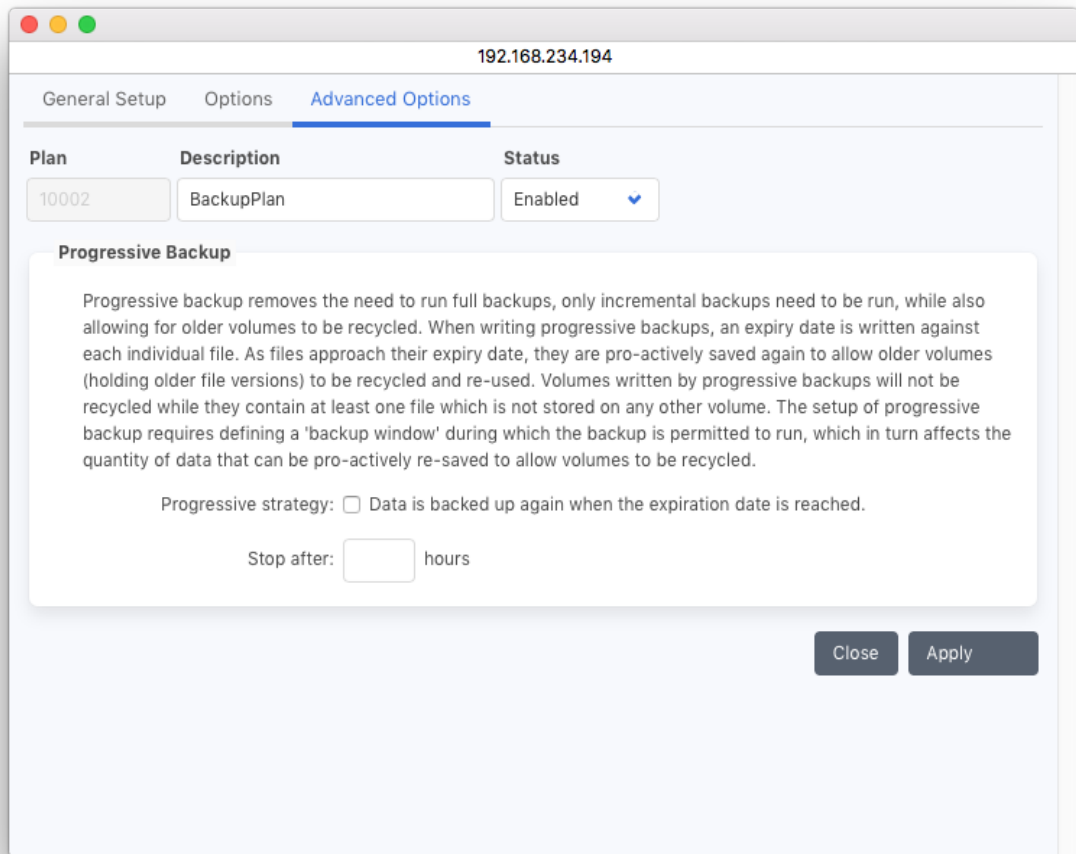


The Backup is not foreseen to keep the data unlimited on tape. Doing so means the backup grows over time and gets slower while doing so, at some time, browsing or searching the backup index may become sluggish.

Another side effect is that folder that contain changes may be spread over many tapes, that implies a restore of that folder must read from many tapes, which requires additional time to mount and position the tapes.

We thus suggest to renew the backup in regular intervals, for instance monthly or every few months, and to recycle and reuse tapes that expired.

Backup Plan: Advanced Options



Progressive Backup:

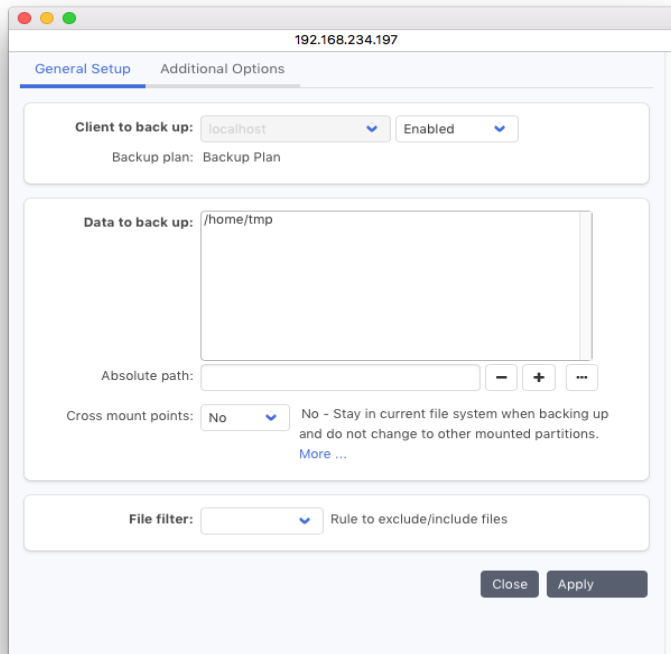
The progressive mode is a special mode to backup files. in the backup module. P5 then maintains an expiration time for each file and saves the file again after that time. The progressive mode is intended for big data sets where the backup does not allow running a full backup. In progressive mode, only the last state of the file system can be restored completely, older states may be incomplete due to recycling. Note that when enabling progressive mode, a stop time must be defined to limit the time that the backup runs. See also section 13.2.2, Progressive Backup for Huge Amounts of Data .



Activate the progressive mode only when writing to a tape library or virtual tape library. For other media, progressive mode is not applicable and does not make sense. Specially, do not set progressive mode when writing to Container Storage.

Backup-Task: General Setup

By selecting a client from the list of clients and clicking 'edit', you can configure the Backup Task. In this window you can define the directories to be backed up for this client and the filters to be applied.



Client to back up:

Select the client whose directories are to be backed up. The P5 server is always listed as localhost. In order to define a client that is not in the list, see section 6.11, P5 Clients.

Data to back up:

Add all the client's folders to be backed up to this list

Absolute path:

Enter the path of the directory you want to backup. Add or remove it in the list by using the '+' or '-' buttons

- + The path will be added to the list
- The selected path will be removed from the list

Cross mount points:

No By default, P5 does not traverse mount points, each file system to include in the backup must be specified

Yes The selected paths and file systems mounted below will be included in the backup.



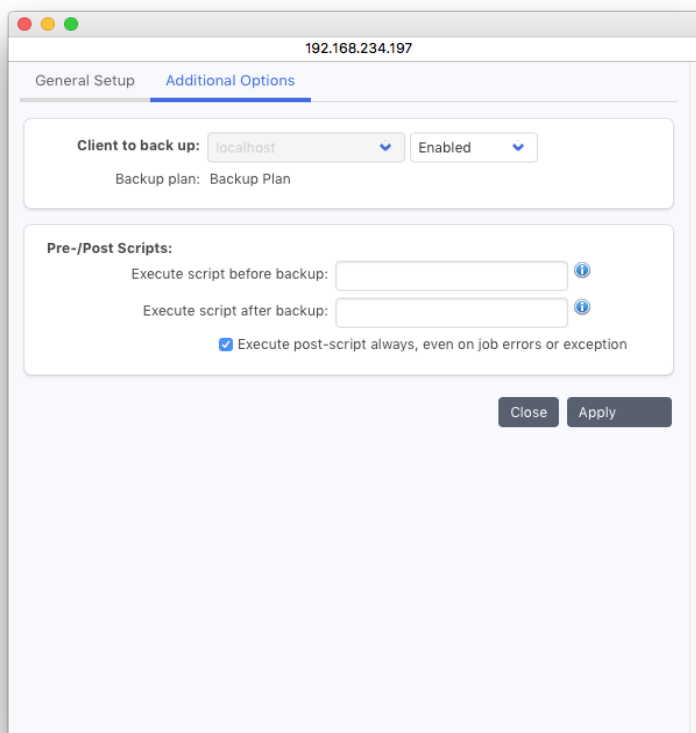
All is a placeholder for hard disk partitions mounted at the start time of the backup job. Dynamically mounted partitions such as removable disks, network shares or Btrfs sub-volumes are not included with the **All** placeholder.

In order to backup dynamic volumes, either that volumes must be specified, or a parent directory must be specified as path together with the **cross mount points** option.

Please use the **All** placeholder in preference to specifying the root file system `/` and setting the option Cross mount points to **Yes**. The later would include USB drives, CDs or network shares when these are mounted, which is not recommended.

In order to exclude certain files or folders use a filter. See section 6.9, Filters.

Backup-Task: Additional Options



Prescripts and Postscripts:

It might be necessary to stop some programs, for example a database, before starting a backup of its files. After completing the backup, you can start the stopped programs again.

Execute script before backup:

The script or command to execute before starting the backup.

Execute script after backup:

The script or command to use after completing the backup.

Specify the full path to your script to execute it on the source client or prefix the path with the P5

client name separated by a colon to execute the script on a specific client.

Example:

on the source client:

`/usr/local/scripts/db-stop`

on a specified client:

`dbserver:/usr/local/scripts/db-stop`

on a windows machine:

`/c/database/scripts/db-stop.bat`

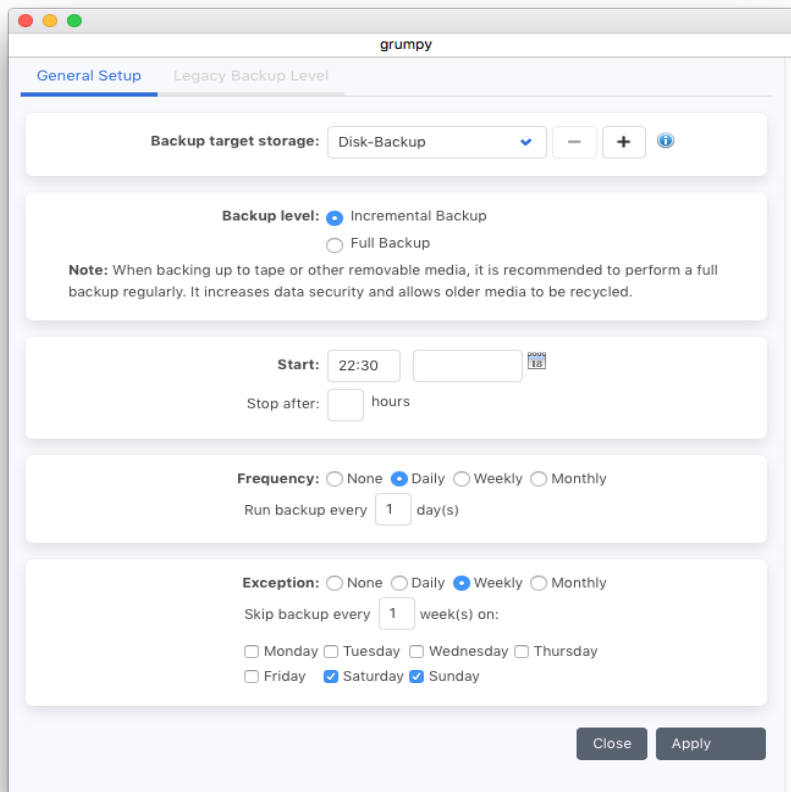


Note that when giving relative paths, these are relative to the P5 install folder.

Pre- and Postscripts are executed only in case the according job is executed by the scheduler, they are skipped on manual execution.

Backup – Schedule Events

The Backup schedule consists of one or more backup events. Each event holds a backup level, a repetition rule, a possible exception for that rule and a target backup pool. By creating multiple schedule events within a backup plan, different actions can be defined with different time schedules.



Backup level:

You can choose between three levels:

Full Backup	A full backup of the configured clients will be done.
Incremental	Only those files that changed since the last full or incremental backup are included in the backup.

Start backup

Specify the date and time when to run the backup plan. The start date is used as the calculation base for the frequency (see below).

Stop after:

If a value is given, P5 stops running the backup plan after the given number of hours. Data that has not been written within the given time period is not saved by the current run of the backup plan and can be saved by a subsequent (incremental) backup.

Frequency:

You can choose between four different calculation bases:

None	The schedule event will appear only once, at the specified start time
Daily	The event will be triggered every n days
Weekly	The event will be triggered every n weeks, the day(s) of the week can be specified.
Monthly	The event will be triggered every n months, the week and the day of the week can be specified.

Exception:

An exception can be specified the same way as the frequency. If for example an incremental backup must be skipped when a full backup is performed, configure an exception to the incremental backup event.

Target pool:

The target pool for this event. Choose a pool from the pulldown menu. Usually one media pool to write to is sufficient. It is possible to define up to three media pools to write to. In case multiple pools are specified, the backup will switch to an alternate backup mode. In this mode:

- the first of the specified pools that is available for backup will be selected
- the backup is then executed referencing the media pool instead of the backup plan, this implies that all files that are not saved in the corresponding media pool are saved again.

In general, it does not make sense to specify backup events with one target pool and those with multiple target pools within the same backup plan.



When defining multiple events, make sure that these do not overlap to avoid unexpected results or errors.

Care has to be taken not to define an impossible schedule through adding an exception.

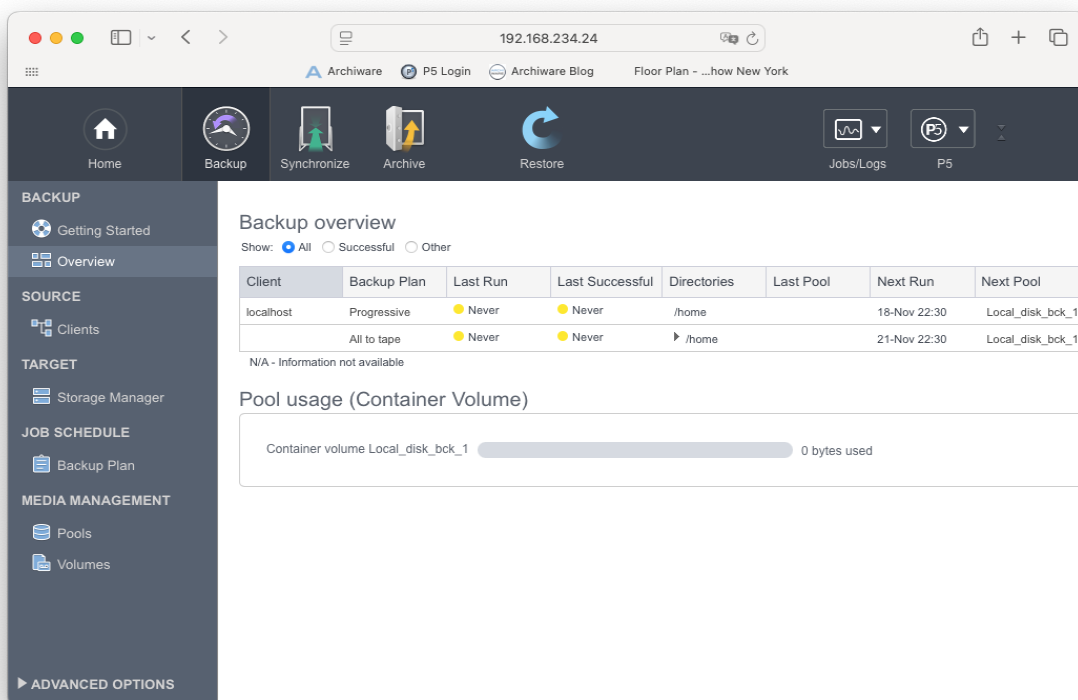
To confirm your input, click **Apply**. In case all the entered data is correct, a backup plan will be created.

6.6.2 Backup Overview

To obtain a status report of all the active backups at a glance, visit the Backup Overview.






To reach the backup overview:

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the Backup module from the top menu and select Backup Overview in the sidebar






The display shows the plan name, the clients as well as the folders which are being backed up for each backup plan. Each line also shows the result of the last execution or, if applicable, indicates that the plan is currently running.

The Last Backup column shows the return state of the last backup execution:

-  Green: The job finished successfully
-  Orange: The job terminated with warnings. See the file access log in the job monitor.
-  Red: The job terminated with error
-  Blue: The job is currently running
-  Yellow: The job has never been started

The last successful column shows when the job ran successfully the last time:

-  Green: within the last 24 hours
-  Red: more than 24 hours ago
-  Yellow: Never

The Next Run shows when the plan is scheduled to run next.

6.7 Host Sync Plans

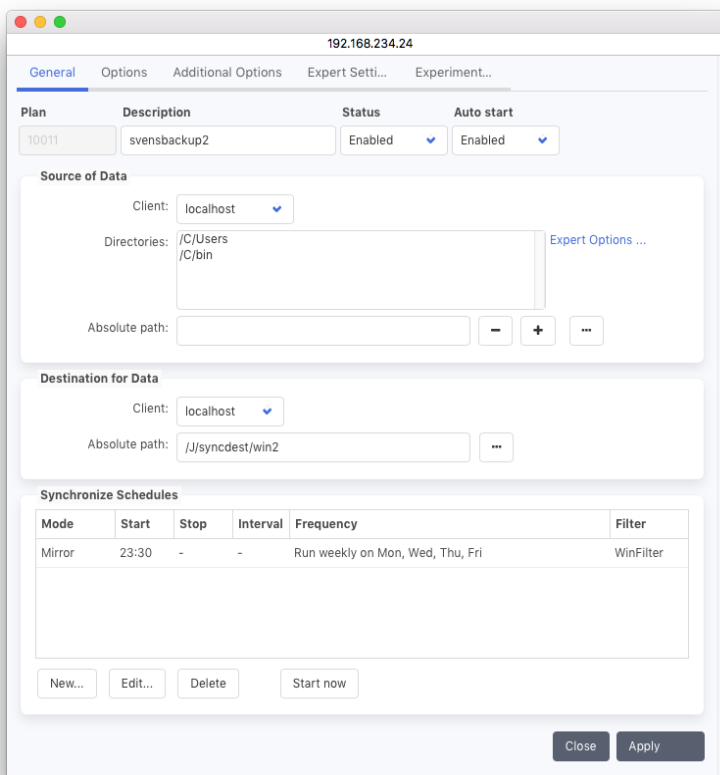
With the Synchronize plan, P5 keeps a target directory in sync with a source directory in defined time intervals. Any files available in the source directory at the moment of running a synchronize plan are copied over to the target. Files deleted from the source directory do not have to be deleted on the target at the same time.



Synchronization is not a two way process, so the source will not be synchronized back from the target. The purpose of synchronization is primarily to increase the availability and security of the data since no time-consuming restore operations are required.

6.7.1 Creating a Sync Plan

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select Synchronize from the top menu and then Synchronize Plan.
3. Select New to create a new Sync Plan.
4. Fill in the required fields.



Description:

Enter the desired plan name here.

Status:

The status flag defines the availability of this plan. The possible states are:

Enabled: The plan is available for use by the system.

Disabled: The plan will not be used by the system.

Auto start:

Enabled: The plan will be started automatically at the given time.

Disabled: The plan will not be started automatically.

Source:

In this area, all the information relevant to the source is given. This includes:

Client:

The P5 client on which the source folder(s) will be found.

Directories:

The absolute path(s) of the folder(s) on the given client.

If a single directory is given, the files from that directory are synchronized to the target.

If multiple directories are given, these are created as sub-directories of the directory on the target and serve as targets for the corresponding directories on the source.

In order to create shorter target paths and to create more appropriate target folder, it is possible to put a "dot-folder" as marker into the source path(s). P5 will then omit the part of the paths left of the marker.

Example:

Synchronizing	Will create
/volume/data1/project1	/volume/backup/volume/data1/project1
and	and
/volume/data2/project2	/volume/backup/volume/data2/project2
into	
/volume/backup	

Synchronizing	Will create
/volume/data1/./project1	/volume/backup/project1
and	and
/volume/data2/./project2	/volume/backup/project2
into	
/volume/backup	



Note:

The target folders resulting from this path modification must not overlap, otherwise the results will be unpredictable.



Note:

Starting with version 5.5.0, Synchronize supports the direct transfer of data from one client to another. It is no longer required that the source or the target folder must be on the same machine running the P5 server.

Destination for Data:

In this area, all the information relevant to the target directory is given. This includes:

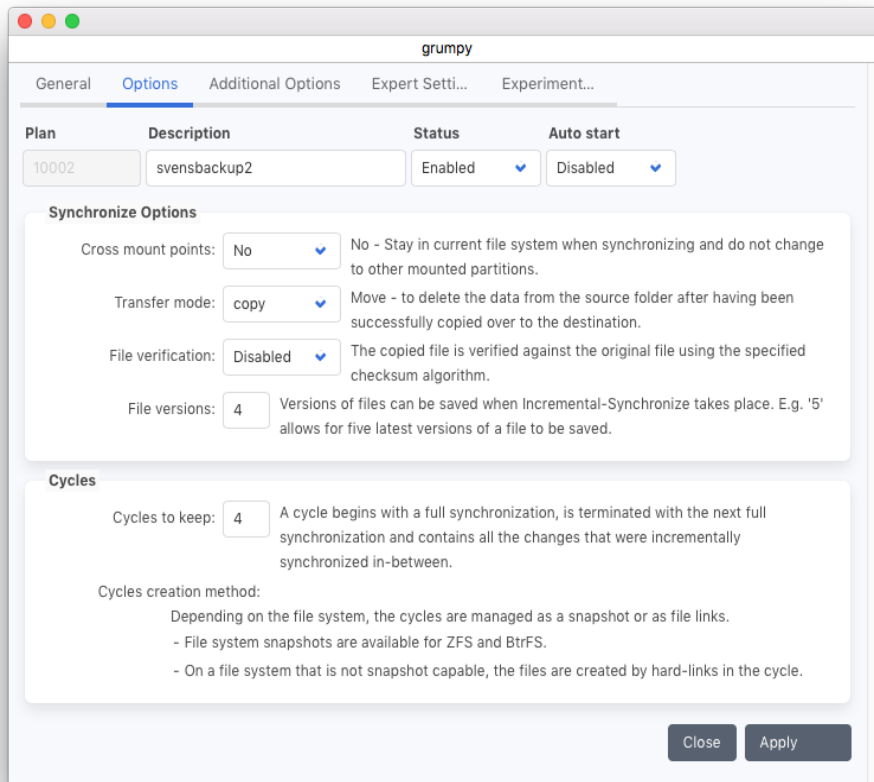
Client:

The P5 client on which the target folder will be found.

Path:

The absolute path of the folder on the given client.

Synchronize Options:



Cross mount points:

In Unix, file systems are mounted onto mount points, which are directories in the directory tree. This option shows if such mounted file systems should also be traversed and synchronized:

Yes : Synchronize mounted file systems.

No : Do not synchronize the mounted files systems.

Transfer mode:

The transfer mode allows using P5 synchronize as an automated file transfer mechanism:

Move : Delete files after copying ("transfer").

Copy : Do not delete files on the source ("copy", regular mode).

File Verification:

While archiving, a checksum is created when the file is read on the source side. When enabled, the written file on the target system is checked against that checksum. In addition, the checksum is stored in an extended attribute of the written file, to enable checks outside of P5.

XXH3 checksums are the fastest, but different checksum types are possible.

File versions:

This specifies how many of the latest versions of a file must be kept. If not zero, the previous versions of the file are stored in a `__VERSIONS` sub-directory, within the same directory where the original file is stored.

Later, in the Synchronize Schedule events, you can specify if you would like to keep the versions while executing the event.

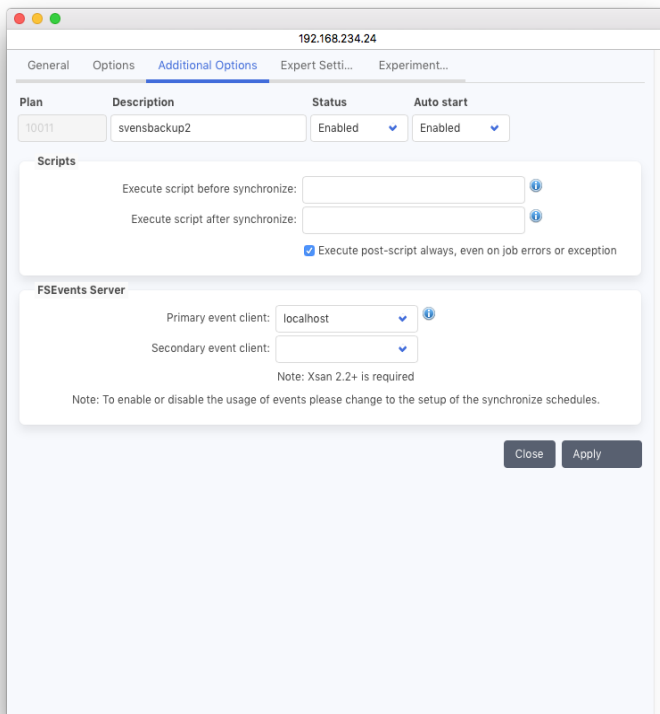
Cycles to keep:

When synchronizing in mirror (full) mode, you may choose to retain the previous data on the target in the form of an old cycle. The chosen number of cycles will be retained. In this case, a new copy of the data will be kept on the target.

When retaining a cycle, P5 creates hard links of the files that are kept unchanged.

In case the target filesystem is ZFS or BTRFS, P5 creates filesystem snapshots instead.

Additional Options:



Scripts

It might be necessary to stop some programs, for example a database, before starting a synchronization of its files. After completing the synchronization, you can start the stopped programs again.

Pre-Script:

The script or command to execute before starting the synchronization.

Post-Script:

The script or command to use after completing the synchronization.

Specify the full path of your script to execute it on the source client or prefix the path with the P5 client name separated by a colon to execute the script on a specific client.

Example:

on the source client:

```
/usr/local/scripts/db-stop
```

on a specified client:

```
dbserver:/usr/local/scripts/db-stop
```

on a windows machine:

```
/c/database/scripts/db-stop.bat
```

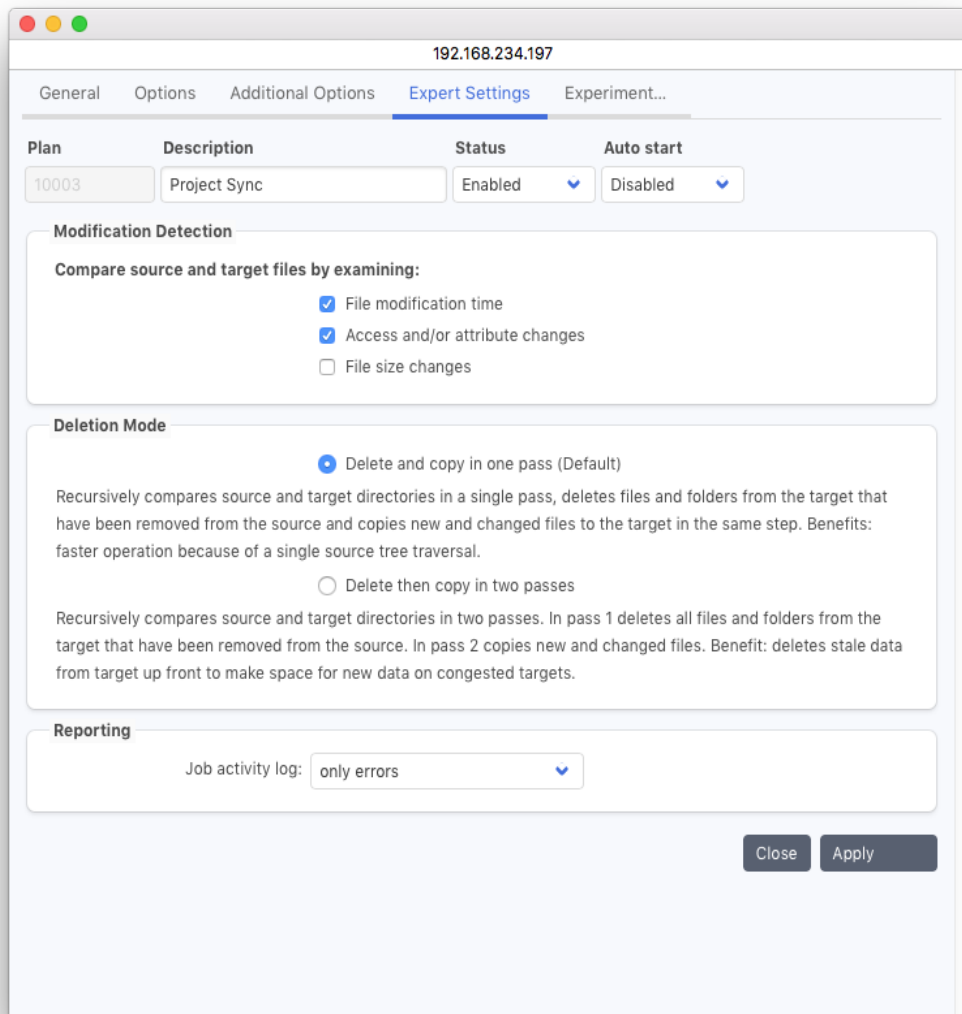
FSEvents Server:

This setting is specific to the XSan 2.2 file system. Use this only on Macs when the source of the synchronize plan is on an Xsan. Xsans support fsevents on the meta data controller MDC. In order to use file system events, a P5 client must be installed on the MDC and the P5 client's name must be set in the synchronize options. You can chose whether or not to use file system events. They can be set in the synchronize job schedule below.



Please do not specify an FSEvent server for file systems local to the selected client.

Synchronize Plan Expert Settings



Modification Detection

These settings are used to tell Synchronize which attributes of the file to use when deciding whether a file should be synchronized or not:

- | | |
|--|---|
| <i>File modification time</i> | Compare the modification time of the file on the source with that of the target |
| <i>Owner and/or protection changes</i> | Compare whether the ownership or the protection mask changed |
| <i>File size difference</i> | Compare whether the size changed |



Notes:

Use these settings only if you understand their impact. By default all the options are enabled (checked).

When using the "Owner and/or protection changes" for file comparison, the target file system **MUST** support extended attributes (xattr). Current standard file systems (like HFS, EXT3/4, NTFS) do so.

Deletion Mode

Delete and copy in one pass

This is the default behavior. P5 will traverse all folders in their natural order and copy or delete files.
In case copying is done first and deletion later on (for instance because a folder was moved and the new location is read before the old one), then the sync target folder requires additional space during the sync job.

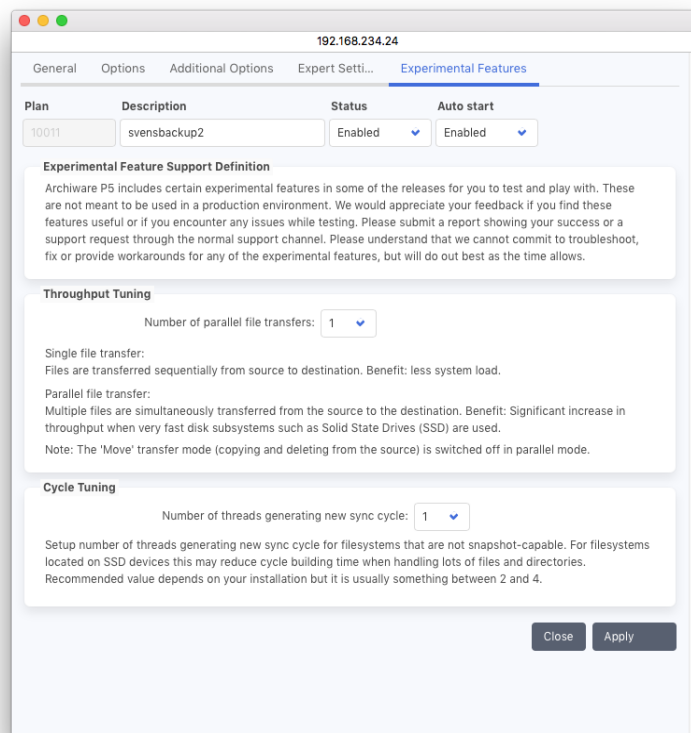
Delete and copy in two passes

With this setting, the sync job will traverse the target twice. In the first pass, deletion takes place for files no longer on the source. This processing is slower but requires less space on the target folder.

Reporting – Job activity Log

By default, errors only are kept in the activity log. Select *Errors and file operations* in order to also log each file that is updated by the synchronize job. This setting will be applied to a single job execution and reset afterwards in order to avoid huge logs.

Synchronize Plan, Experimental Settings



Throughput Tuning

This setting allows to define multiple parallel transfers in one sync plan. Parallel transfers require additional threads, so the setting should not exceed the number of CPU cores. The file transfer may then be sped up on fast disks with short latency like SSD. On standard single HDD storage, increasing the value usually yields no advantage.

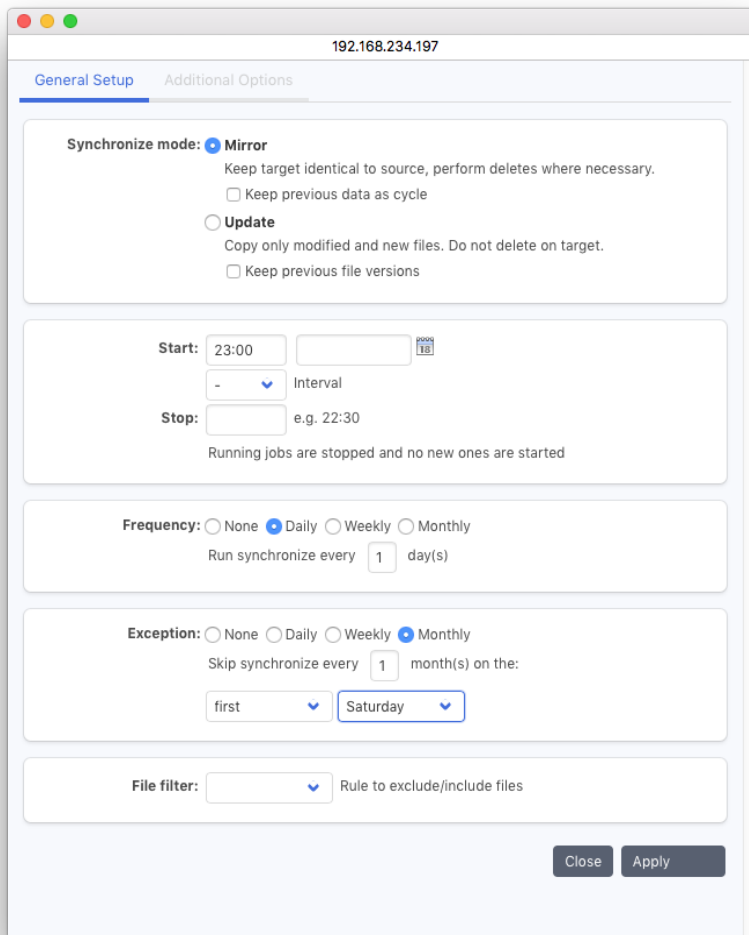
Cycle Tuning

This setting allows multiple parallel threads to be created when creating Snapshots in the __CYCLES folder. This works with non-snapshotting file systems where snapshots are created using filesystem links or copies. The setting should not exceed the number of CPU cores. The cycle creation may then be sped up on fast, low latency disks (SSD). On standard single HDD storage, increasing the value usually yields no advantage.



Please do not use experimental features in production environments without extensively testing first.

Synchronize Plan Schedule, General Options



Synchronize Mode:

Define the mode for this event:

Mirror: Any files on the target, which are no longer present on the source, will be removed from the target. Upon completion of the synchronization, the source and target will carry identical files. This mode was previously named "Full Sync". When using Synchronize as backup to disk, this mode is comparable to a full backup. Mirror mode does not transfer files that are already on the target. In mirror mode, the previous state on the target can be preserved in snapshots. The sync plan options define how many such snapshots will be kept and how to generate them.

Update: No files will be removed from the target. Only new files and changes will be copied. This mode was previously called "Incremental Sync". It is comparable to an incremental backup. Versions of changed files can be maintained. The sync plan options define how many versions will be kept.

Start:

Define the start date and time. The date defines when the plan should be run for the first time. The start time defines at which time it should be run. It is possible to rerun the sync in intervals by specifying an

interval and an end time. According to the frequency below, the plan will be executed at the given time.

Stop and Interval:

P5 can repeat the Synchronize Job in short intervals, for instance hourly, until the given stop time is reached. Such short time repetitions can be used to keep a folder synchronized during the day. For these repetitions, it is recommended to make use of the File System events to detect changes and to avoid generating saved cycles by keeping the old state as snapshot..

In case a stop time is set, a possibly running sync job will be canceled when the Stop time is reached.

Frequency:

The frequency defines how often the plan should be executed. It is possible to set complex frequencies such as "every third week on Wednesday" or "monthly on the first Saturday". The start date allows for complex planning such as two plans that alternate weekly.

Exception:

Exceptions can be set in the same way as the frequency in order to disable the execution at specific times, for example if the drive is required for some other purpose at the time.

File Filter:

If you wish to exclude or include specific files to synchronize, choose the appropriate filter. See section 6.9, Filters.

Synchronize Plan Schedule, Additional Options

Synchronize Policy:

Choose whether to let P5 use file system events or to scan the file system in order to determine which files it needs to synchronize. File system events are supported when the data source is

- on a Windows host with the NTFS file system.
- on a macOS the HFS+ file system.
- on an Xsan version 2.2 or later (via a P5 client running on the meta data controller).
- On a Linux system using the inotify interface.
-
- Combining multiple synchronize schedules optimizes the power of P5 Synchronize as a replication mechanism. For example, using file system events during business hours to periodically run an Update mode sync with a short interval filter to exclude files that are currently being worked, allows to keep track of current changes. A sync run in Mirror mode in the evening with file system scan ensures no files are lost or missed by a missed file system event.

Load limitation:

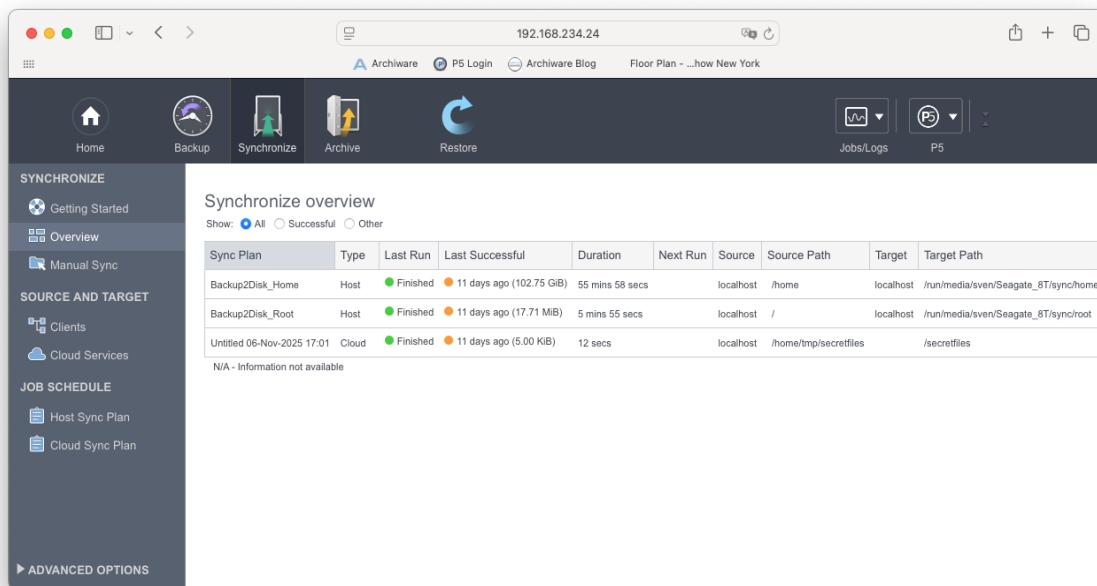
When synchronizing to or from a P5 client, this setting allows to limit the network bandwidth used by the job, this may make sense in case there are parallel jobs running over the same line or in case the job to or from a remote client fails with too much network load.

6.7.2 Synchronize Overview

To obtain a status report of all the active synchronize plans at a glance, visit the Synchronize Overview.






To reach the synchronize overview:

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the Synchronize module from the top menu and select Synchronize Overview in the sidebar






The display shows the plan name and the source and target folders which are being synchronized for each sync plan. Each line also shows the result of the last execution or, if applicable, indicates that the plan is currently running.

The last run column shows the return state of the last backup execution:

-  Green: The job finished successfully
-  Orange: The job terminated with warnings. See the file access log in the job monitor.
-  Red: The job terminated with error
-  Blue: The job is currently running
-  Yellow: The job has never been started

The last successful column shows when the job ran successfully the last time:

-  Green: within the last 24 hours
-  Red: more than 24 hours ago
-  Yellow: Never

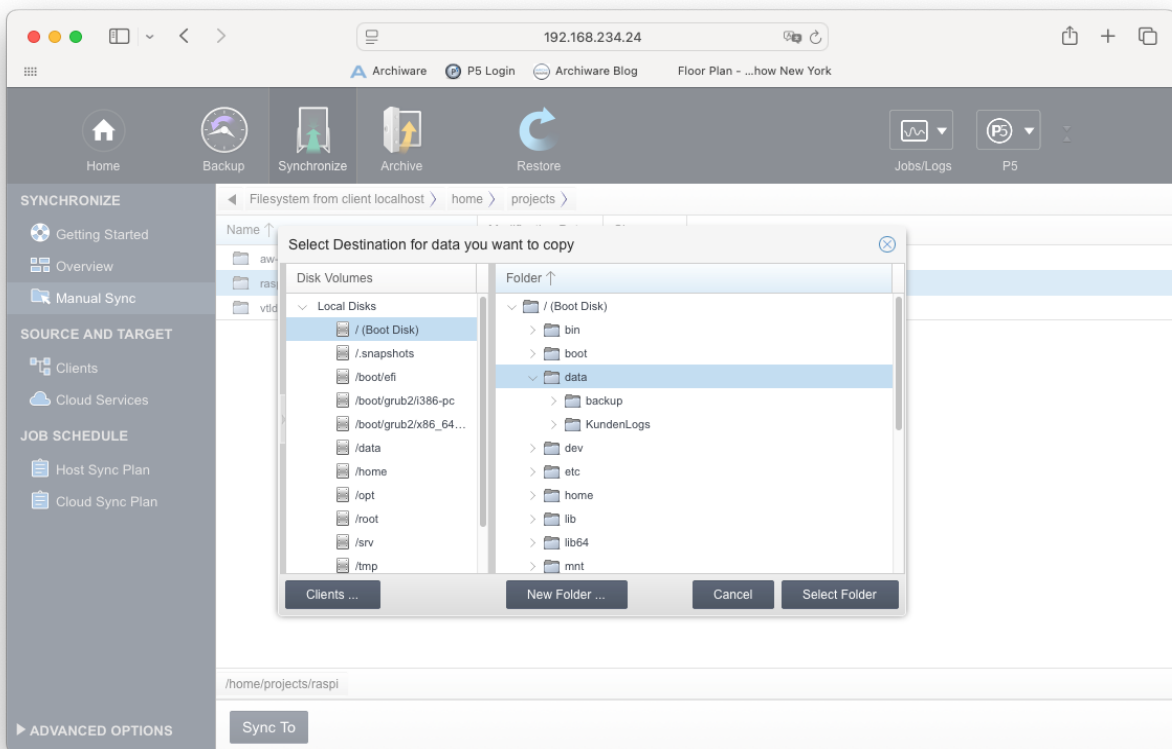
The Next Run shows when the plan is scheduled to run next.

6.7.3 Manual Synchronize

For a single, spontaneous synchronize job, use the manual synchronize section.

To reach the manual synchronize interface:

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the Synchronize module from the top menu and select Manual Sync in the sidebar



The Interface will then ask for source and target folders and define temporary synchronize jobs which are then started immediately. In case multiple source files or folders are selected, there will be a separate job for each selected item; on the target side these will be placed in the same target folder.

6.8 Cloud Sync Plans

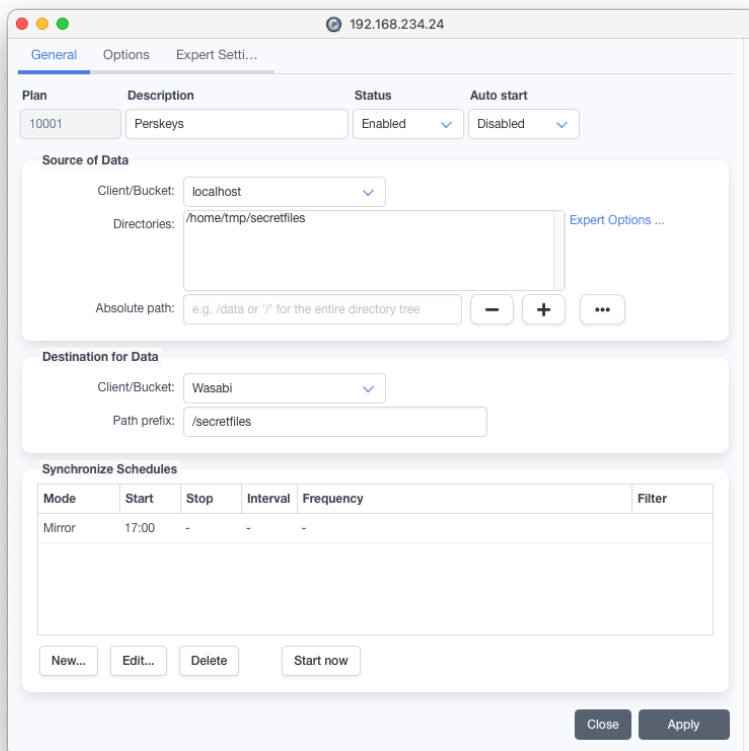
With the Cloud Sync plan, P5 keeps a target in sync with a source in defined time intervals. Any files available in the source at the moment of running a cloud sync plan are copied over to the target. The difference to the Host Sync Plan is that one side, either source or target, is a cloud storage. That in turn also requires a different implementation.



Synchronization is not a two way process, so the source will not be synchronized back from the target. The purpose of synchronization is primarily to increase the availability and security of the data since no time-consuming restore operations are required.

6.8.1 Creating a Cloud Sync Plan

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select Synchronize from the top menu and then Cloud Sync Plan.
3. Select New to create a new Cloud Sync Plan.
4. Fill in the required fields.



Description:

Enter the desired plan name here.

Status:

The status flag defines the availability of this plan. The possible states are:

Enabled: The plan is available for use by the system.

Disabled: The plan will not be used by the system.

Auto start:

Enabled: The plan will be started automatically at the given time.

Disabled: The plan will not be started automatically.

Source:

In this area, all the information relevant to the source is given. This includes:

Client/Bucket:

The P5 client on which the source folder(s) will be found.

In case the source side is a cloud storage, specify a cloud service here.

Directories (if the source side is local):

The absolute path(s) of the folder(s) on the given client.

Path prefix (if the source side is a cloud service):

The path(s) to the folder in the bucket specified by the cloud service.

Destination for Data:

In this area, all the information relevant to the target directory is given. This includes:

Client/Bucket:

The P5 client on which the target folder(s) will be found.

In case the source side is a cloud storage, specify a cloud service here.

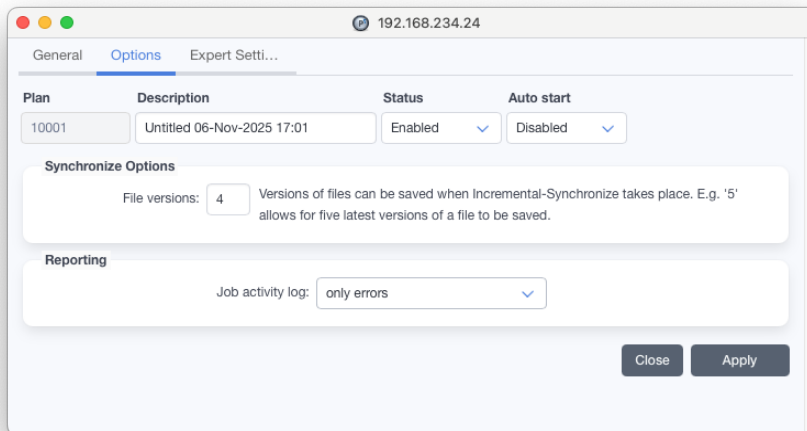
Absolute path (if the target side is local):

The absolute path(s) of the folder(s) on the given client.

Path prefix (if the target side is a cloud service):

The path(s) to the folder in the bucket specified by the cloud service.

6.8.2 Cloud Sync Options



File versions:

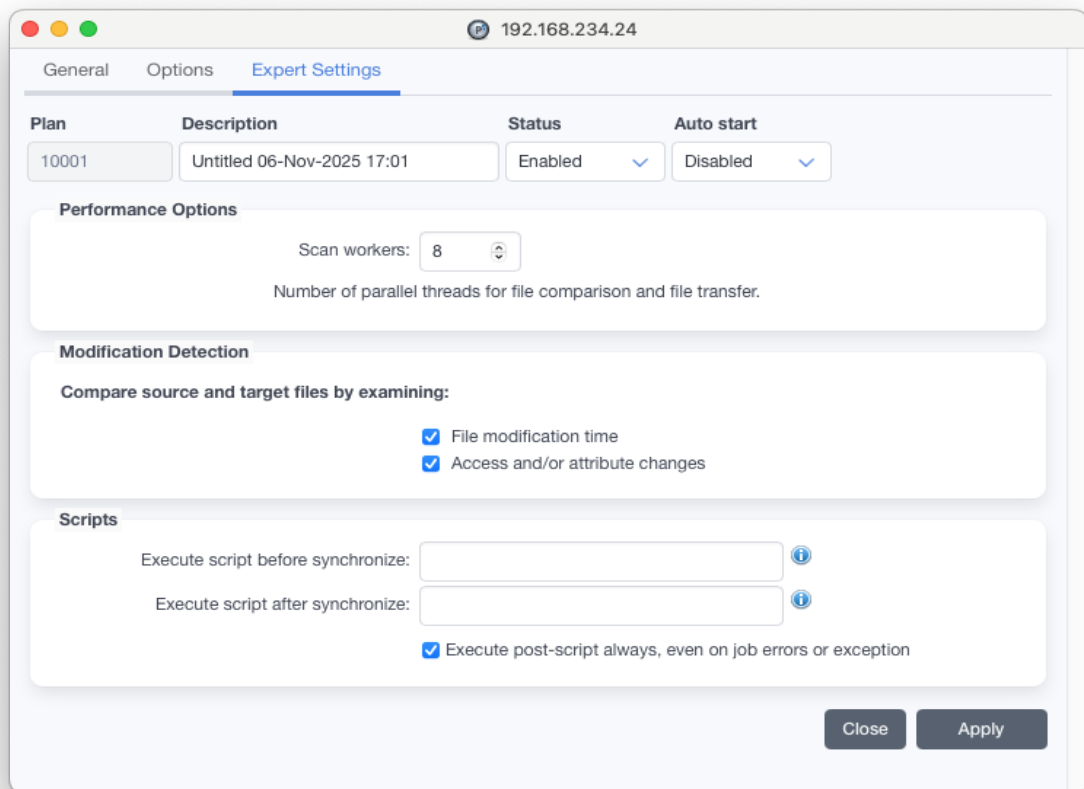
This specifies how many of the latest versions of a file must be kept. If not zero, the previous versions of the file are stored in a `__VERSIONS` sub-directory, within the same directory where the original file is stored.

Later, in the Synchronize Schedule events, you can specify if you would like to keep the versions while executing the event.

Reporting – Job activity Log

By default, errors only are kept in the activity log. Select *Errors and file operations* in order to also log each file that is updated by the synchronize job. This setting will be applied to a single job execution and reset afterwards in order to avoid huge logs.

6.8.3 Cloud Sync Expert Settings



Performance Options

Scan Workers

The number of threads that scan the folder tree for changes

Modification Detection

These settings are used to tell Synchronize which attributes of the file to use when deciding whether a file should be synchronized or not:

File modification time

Compare the modification time of the file on the source with that of the target

Access and/or attribute changes

Compare whether ownership or protection mask changed

Scripts

It might be necessary to stop some programs, for example a database, before starting a synchronization of its files. After completing the synchronization, you can start the stopped programs again.

Pre-Script:

The script or command to execute before starting the synchronization.

Post-Script:

The script or command to use after completing the synchronization.

Specify the full path of your script to execute it on the source client or prefix the path with the P5 client name separated by a colon to execute the script on a specific client.

Example:

on the source client:

```
        /usr/local/scripts/db-stop
```

on a specified client:

```
        dbserver:/usr/local/scripts/db-stop
```

on a windows machine:

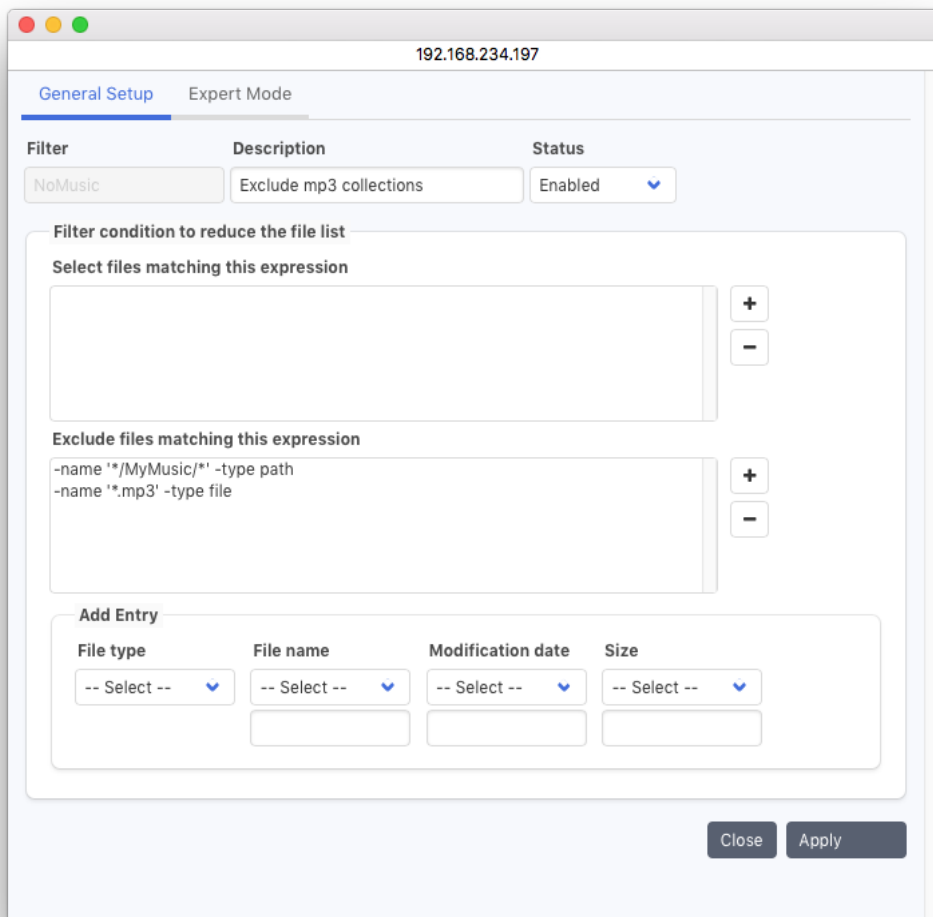
```
        /c/database/scripts/db-stop.bat
```

6.9 Filters

While executing a backup, a synchronize or an archive plan, the entire folder tree, including all the files and subfolders within the selected folder, will be saved through an in-place recursive folder tree traversal. If thereby some files are to be omitted, they can be filtered out.

6.9.1 Creating a Filter

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select **File Filter** from the **Advanced Options** sidebar.
3. Select **New** from the **bottom bar** to create a new filter.
4. Fill in the required fields.



Filter Name:

Enter the desired filter name here. The following is checked:

1. Uniqueness: No other filter may carry the same name.
2. Validity: The name is at most 32 characters long and can be composed of ' [A-Z][a-z][0-9]-_: . '. White space is not allowed.

Description:

You can describe the filter here.

Filter Status:

The status flag defines the availability of this filter. The possible states are:

Enabled: The filter is available for use by the system.

Disabled: The filter will not be used by the system.

Filter Condition to Reduce the File List:

In the course of executing a task, every file and/or path name (element) from the list of files to be processed is checked against the defined filter in two steps:

Select files matching this expression:

If the check against any of the defined conditions succeeds, the process will continue.

Exclude files matching this expression:

This step works the other way around. If a check against any of the defined conditions succeeds, the element will be disregarded and will not be processed. If both conditions apply, exclusion is stronger.

Add Entry:

By means of the + button on the right, the filter definition set up below is added to the list. The - button removes the selected definition from the list.

File type

With the *file type* field, you define to which element the given pattern should be applied.

File: An ordinary file.

Directory: A directory and all its elements.

Path: An absolute path and all its elements.

File name

You can choose to check the entered *file name* against different patterns.

ends with: Requires the element to end with the given string.

starts with: Requires the element to begin with the given string.

contains: Requires the element to contain the given string.

matches: Requires the element to exactly match the given string.

Modification date

With the *modification date* field, you can define a relation to the file's modification date. An element of the chosen type will apply if its modification date:

- is older: than the given number of days.
- is newer: than the given number of days.
- is exactly: on the given date.
- is before: the given date.
- is after: the given date.

Size

With the *Size* field, you can define a relation to the file's size. An element of the chosen type will apply if its size is:

- smaller than: the given number of kilobytes.
- larger than: the given number of kilobytes.

Expert mode:

The expert mode for configuring a filter gives you the means to extensively define the data to be included or excluded from the backup by using the same syntax as for the Unix `find` command.

The `General Setup` window allows to define filters for the most common cases without requiring expert knowledge.

To confirm your input, click `Apply`. In case all the entered data is correct, a filter will be created.

6.10 Logs

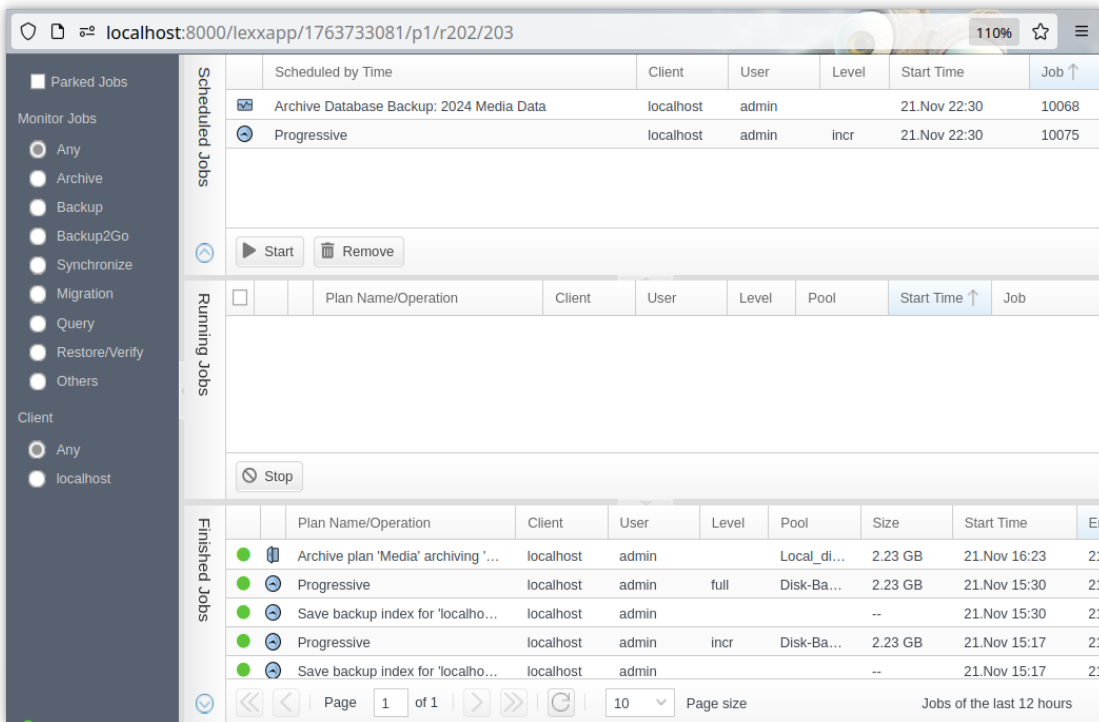
P5 automatically creates and manages logs and protocols. Every job results in a log entry upon completion. You can view this entry with the data browser.

Job results can also be e-mailed to selected user(s), see 11.3, Customizing Reporting.

6.10.1 The Job Monitor

The Job Monitor shows the currently running jobs and monitors the current activities. To open the Job Monitor:

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the Job Monitor from the top menu's Job/Logs section



By default, the job monitor shows the current jobs. The sidebar allows to filter the jobs by:

- Jobs:** Filter the visible jobs to show only jobs of the selected type(s).
- Client:** Filter the visible jobs to show only jobs of the selected client(s)

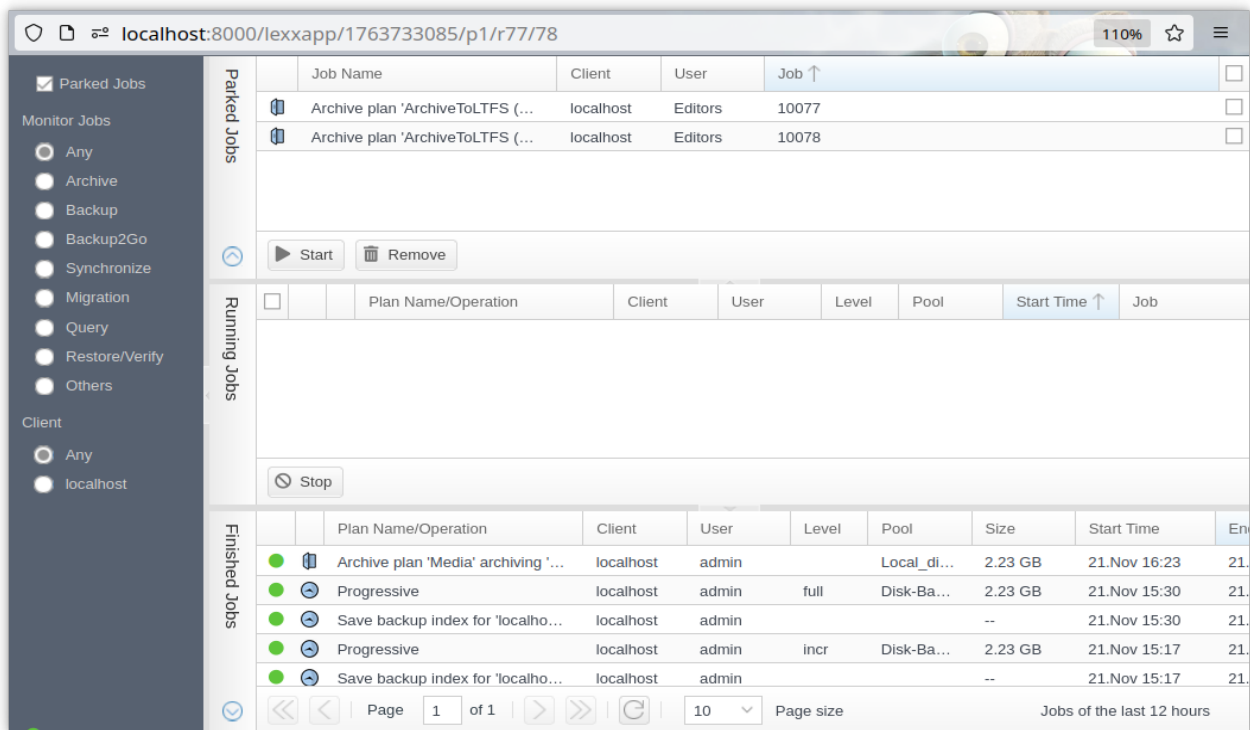
Click the check box to the left of each job for further actions, such as starting or stopping jobs or getting a detailed job report.

Parked Jobs

In the user preferences it is possible to define whether archive and restore jobs shall be started right away or enqueued as parked jobs, see section 6.14, User Preferences. Also jobs that are running but is still waiting to access a media to write to can be manually parked in the job monitor.

Parked jobs do not start automatically but must be started manually by an operator. To toggle the job monitor between parked and scheduled job, click at **Parked Jobs** in the upper left edge.

This mechanism allows a higher control when and in which order jobs shall be started.



Keeping the job monitor window open will also prevent the current login session from expiring.

6.10.2 The Log History

The Log History window shows results of past jobs. To open the log history:

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the Log History from the top menu's Job/Logs section

The screenshot shows a web browser window displaying the Log History interface. On the left is a sidebar with filters for Job Log Query (Last 12 hour, Today, Yesterday, Past week, Past three months, Enter period), Result (Any, Completed, Warning, Interrupted, Error), Type (Any, Archive, Backup, Backup2Go, Synchronize, Migration, Query, Restore/Verify, Others, Search), and Client (Any, localhost). The main area is a table with columns: Plan Name/Operation, Client, User, Level, Pool, Size, Start Time, and End Time. The table contains 34 rows of job logs, including operations like 'Archive plan', 'Progressive', 'Save backup index', 'Labeling for pool', 'Archive Database Backup', 'Restore Job', and 'Projects2024'. The bottom of the window shows pagination controls: Page 1 of 1, 50 items per page, and Displaying 1 - 34 of 34.

Plan Name/Operation	Client	User	Level	Pool	Size	Start Time	End Time
Archive plan 'Media' archivin...	localhost	admin		Local_di...	2.23 GB	21.Nov 16...	21.Nov 16...
Progressive	localhost	admin	full	Disk-Ba...	2.23 GB	21.Nov 15...	21.Nov 15...
Save backup index for 'localh...	localhost	admin			--	21.Nov 15...	21.Nov 15...
Progressive	localhost	admin	incr	Disk-Ba...	2.23 GB	21.Nov 15...	21.Nov 15...
Save backup index for 'localh...	localhost	admin			--	21.Nov 15...	21.Nov 15...
Labeling for pool Disk-Backup	localhost	admin			--	21.Nov 15...	21.Nov 15...
Archive Database Backup: 2...	localhost	admin			--	20.Nov 22...	20.Nov 22...
Restore Job: The Muppets_ ...	localhost	admin			--	19.Nov 10...	19.Nov 10...
Archive Database Backup: 2...	localhost	admin			--	18.Nov 22...	18.Nov 22...
Archive Database Backup: 2...	localhost	admin			--	17.Nov 22...	17.Nov 22...
Archive Database Backup: 2...	localhost	admin			--	16.Nov 22...	16.Nov 22...
Archive Database Backup: 2...	localhost	admin			--	15.Nov 22...	15.Nov 22...
Archive Database Backup: 2...	localhost	admin			--	13.Nov 22...	13.Nov 22...
Projects2024	localhost	admin			64.34 MB	13.Nov 18...	13.Nov 18...
Projects2024	localhost	admin			942.14 MB	13.Nov 15...	13.Nov 15...
Projects2024	localhost	admin			942.14 MB	13.Nov 15...	13.Nov 15...
Projects2024	localhost	admin			942.14 MB	13.Nov 15...	13.Nov 15...
Archive Database Backup: 2...	localhost	admin			--	12.Nov 22...	12.Nov 22...
Archive Database Backup: 2...	localhost	admin			--	11.Nov 22...	11.Nov 22...
Archive Database Backup: 2...	localhost	admin			--	10.Nov 22...	10.Nov 22...
Archive Database Backup: 2...	localhost	admin			--	09.Nov 22...	09.Nov 22...
Archive Database Backup: 2...	localhost	admin			--	08.Nov 22...	08.Nov 22...
Archive Database Backup: 2...	localhost	admin			--	06.Nov 22...	06.Nov 22...
Backup2Disk_Root	localhost	admin	mirror		17.71 MB	06.Nov 18...	06.Nov 19...
Backup2Disk_Home	localhost	admin	mirror		102.75 GB	06.Nov 17...	06.Nov 18...
Untitled 06-Nov-2025 17:01	localhost	admin	update		5.00 KB	06.Nov 17...	06.Nov 17...
Archive Database Backup: 2...	localhost	admin			--	05.Nov 22...	05.Nov 22...
Progressive	localhost	admin	incr	Local di...	--	05.Nov 22...	05.Nov 22...

By default, the job monitor shows the jobs that ended within the last 12 hours. The side bar allows to filter job results by:

Period: Show the job results from the selected period only

Result: Show only the jobs with the selected result code(s).

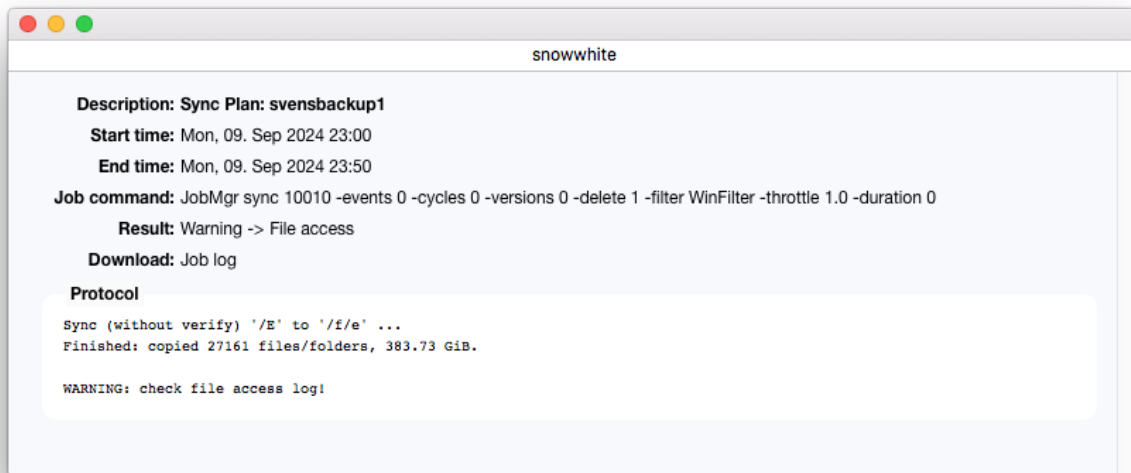
Type: Show only the jobs of the selected type(s).

Client: Show only the jobs of the selected client(s)

In case there are more job logs than fit on one page, Previous and Next let you switch pages.

6.10.3 the Job Log

Double clicking a finished job in the job history or the job monitor leads to the job detail window.



The job log contains one or multiple links, depending on the job group and the job result:

- | | |
|--------------------------------------|---|
| Job log: | This button downloads the job log as text file. |
| Job inventory:
(Archive job only) | This button downloads a list of all files that were archived by the job with a defined set of attributes like size or checksum. |
| File Access:
(Only on failures) | In the case of failures affecting single files, this button leads to a list of all file related errors. |

6.11 P5 Clients

To back up, synchronize or archive data from a machine, it must be configured as a P5 client. The P5 server is automatically configured during installation as the client named *localhost*. Each configured client has its own backup index where a list of the saved files and folders is maintained.



Each host that is to act as a client must have P5 installed. However, only the server must be configured and maintained. After installing P5 on a client computer, remember to change the default password. You will either need to specify the admin user or to create an additional user on the client's P5 web GUI to be used as account in the Client definition on the server. See section 4.4, Using the P5 Agent Setup GUI for how to set a password on the client host.

6.11.1 Creating a Client

1. Log into P5 using a web browser as an admin user.
2. See section 4.1, Starting the Browser.
3. Select the Archive, Backup or Synchronize module from the top menu
4. Select Client from the sidebar.
5. Click New in the client managers bottom bar to create a Client.
6. Fill in the required fields.

Client	Description	Status
Win2012	WS_Sven	Enabled

Host

Computer is licensed as:

TCP/IP Address: e.g. 192.168.10.50 HTTP Port: e.g. 8000

User name: e.g. admin

Password:

Communication

Compress data transfer: e.g. Yes - To compress data before sending over the network.

Encrypt data transfer: e.g. Yes - To send data encrypted over the network.

Limit bandwidth: e.g. '100%' to use the total bandwidth

Client Name:

Enter the desired client name here. The following is checked:

1. Uniqueness: No other client may carry the same name.
2. Validity: The name is at most 32 characters long and can be composed of ' [A-Z][a-z][0-9]-_: . '. White space is not allowed.

Description:

You can describe the client here.

Client Status:

The status flag defines the availability of this client. The possible states are:

Enabled: The client is available for use by the system.

Disabled: The client will not be used by the system.

Client license type:

Use Server for a server operating system and Workstation otherwise. P5 will not allow workstation license to be employed for a server or Virtual Server and Virtual Workstation licenses to be employed on a physical host.

Server: For all machines offering applications and services to other machines, for example file, mail, web or database servers.

Workstation: For machines being used exclusively as workstations.

Virtual Server: For all machines being used exclusively inside a virtual environment.

»Virtual Workstation: For machines being used exclusively as workstations inside a virtualization environment.

TCP/IP address:

Holds either the client's host name or its TCP/IP address.

Port:

When installing P5, a port number is assigned automatically. Port 8000 is offered as the default.

User name:

The login name of the user who has sufficient permissions to access the data to be saved from, usually *root*.

Password:

The password corresponding to the login. The colon character ':' must not be used in the password.

Compress data:

Depending on the quality of the connection, it may help to compress the data being transferred between the P5 client and the server, for example when transferring over a WAN. For fast networks, it is not recommended to use this option since both the client and the server must do additional processing without improving the actual throughput.

Yes: Data will be compressed prior to transfer.

No: Data will be transferred as is.

Encrypt data transfer:

Specify whether the data should be encrypted for transfer. This flag does not affect the storage, only the transfer.

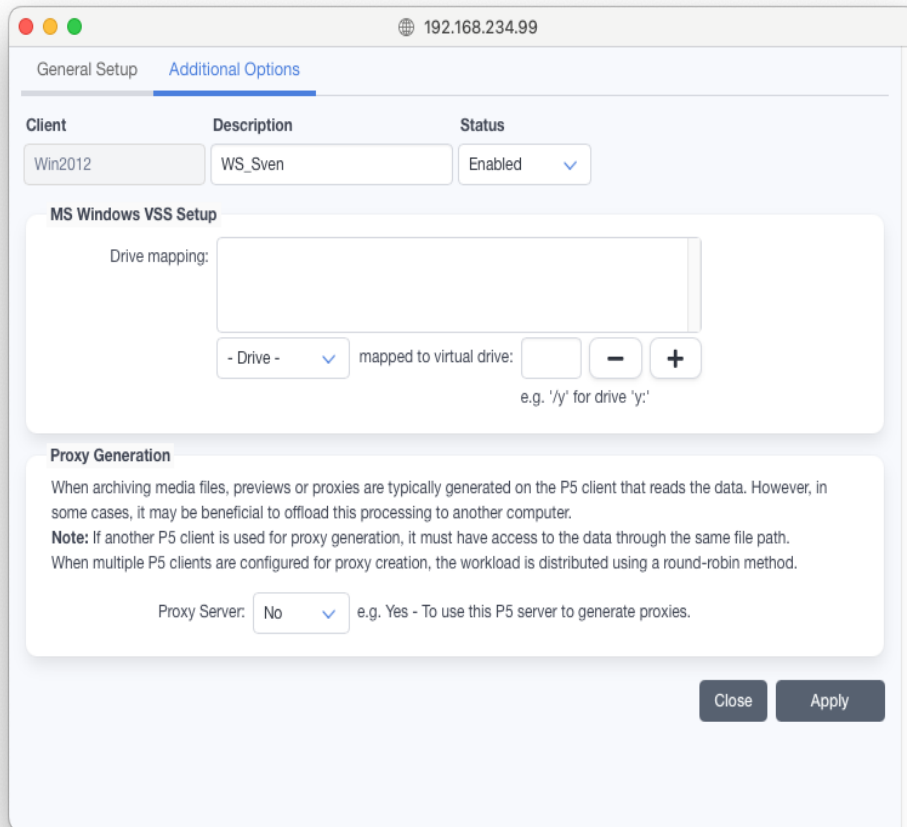
Yes: Data will be encrypted prior to transfer.

No: Data will be transferred as is.

Limit bandwidth:

Limit how much of the available line bandwidth is used. Employ this option when transmission errors would otherwise occur or in case part of the bandwidth is to be saved for other purposes.

Client Additional Options:



MS Windows VSS Setup:

Define for each Windows client's local NTFS file system whether a virtual shadow copy (VSS) should be created during processing. For each corresponding drive letter define an unused drive letter to map it to. P5 will then create a virtual shadow copy for that drive and access it with the defined mapped drive letter.

Virtual shadow copies are snapshots of the file system, created when a backup or synchronize process starts. This mechanism allows to save a consistent copy of the files even if there are opened and changing files at the time.

Proxy Generation:

Starting with version 7.4, P5 allows the use of one or more P5 clients as proxies for the preview generation when archiving. The preview generation is usually resource intensive, so it may make sense to use a fast client to generate preview files when archiving from a storage with a weak processor like for instance a NAS.

A precondition to use a proxy client for the preview creation, the proxy must have access to the archived files with the exactly the same path as the P5 client where the files are archived from. To use a P5 client as proxy for the preview generation, set this files to **yes**.

In case multiple P5 clients are defined as proxy, they are used in a round robin technique.


6.11.2 Client: Interactive Functions

From the bottom bar, several useful functions can be reached:

6.11.2.1 Ping

Ping tests whether a P5 client is available.

A click on **Ping** starts the test. The result is displayed at the bottom of the ping window. A check is made whether the client's version is compatible with the server.

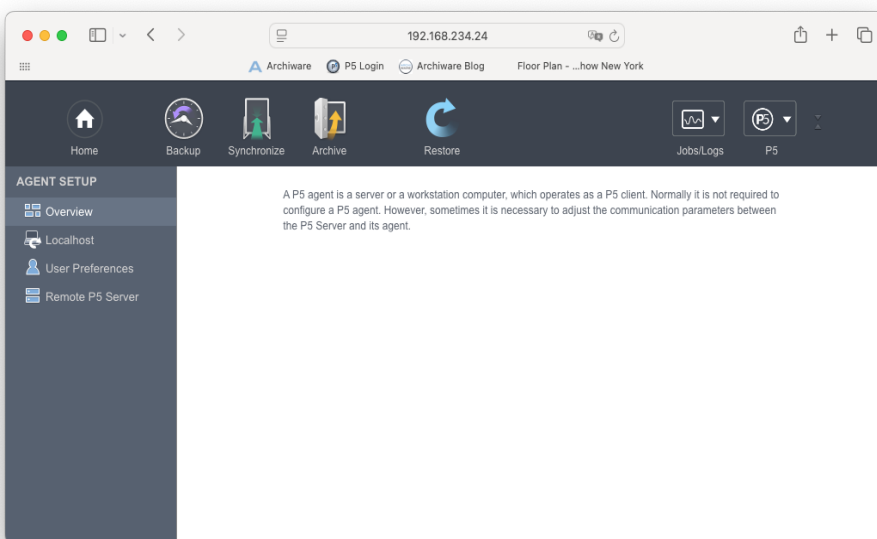


6.11.3 Setting up Encryption for a Client

In most cases, P5 clients do not need any setup or configuration and are controlled remotely by the P5 server. When data encryption is desired in storage, configuration is done on the client for security reasons.

To access the client directly:

1. Log into P5 using a web browser as an admin user. The login procedure is the same as for the P5 server, see section 4.1, Starting the Browser, except that the client's IP address must be used.
2. In the client's interface, select the Agent Setup from the P5 section in the top menu.



3. Double click the localhost entry in the main window to control the encryption of the client's data. To set up data encryption for data on the P5 server itself, select localhost client in the server's interface.



This is the only configuration item that is performed on the client. Entering a client license and configuring other client settings is performed on the P5 server.

In case the encryption for a client host is enabled and a job is run with **verify** option turned on, the P5 server must be able to decrypt the files, otherwise the verify will fail.

The key for encryption/decryption is stored as line in the ascii text file `/usr/local/aw/config/localhost.key`.

Further lines / keys from other hosts can be added to the file to be able to decrypt files from these hosts, the first line/key is used for encryption on the local host.

6.11.4 Setting up Reverse Connection for a Client

When accessing a client, for instance for Backup or Synchronize, the connection is initiated by the P5 server by creating a TCP/IP connection to the client host. That procedure requires the server to know the client's IP address, or to get it from a name server.

Sometimes this is not possible, for instance when the client is accessed via internet and has a temporary IP address only.

In that case, it is possible to setup a "reverse connection", that means:

On the server side specify the client's host-ID instead of its TCP/IP address.

On the client:

1. Log into P5 using a web browser as an admin user. The login procedure is the same as for the P5 server, see section 4.1, Starting the Browser, except that the client's IP address must be used.
2. In the client's interface, select the Agent Setup from the P5 section in the top menu.

The screenshot shows a web browser window with the URL 192.168.234.197. The page title is 'General Setup' and 'Backup2Go Options'. The 'Local Host-ID' is 299821d4-01ff and its status is 'Enabled'. The 'Remote P5 Server' section contains the following fields: 'DNS Name' (MyP5Server) or 'TCP/IP Address', 'Port' (8000), 'User name' (root), and 'Password' (masked with dots). The 'Allow reverse connection' checkbox is checked, with a dropdown menu set to 'Yes'. At the bottom right, there are buttons for 'Register on Backup Server', 'Close', and 'Apply'.

Allow Reverse Connection:

Yes: Allow the server to connect via this connection.

No: The connection will not be used. This is used on Backup2Go Workstations.

DNS name or TCP/IP address:

Holds either the servers host name or its TCP/IP address.

Port:

When installing P5, a port number is assigned automatically. Port 8000 is offered as the default.

User name:

The login name of the user who has sufficient permissions to login to the server.

Password:

The password corresponding to the login. The colon character ':' must not be used in the password.

Register on Backup Server:

After entering all values, register at the P5 server once to inform the server that in this case a reverse connection is used.

After this is set up, the client will initialize and keep up a connection to the P5 server. The server will then use this connection to connect to the client.

6.12 Index Databases

P5 manages a variable number of databases that hold all kinds of information. For example, the list of all the saved files is held in an indexed database. When clients are created, an index database is also created for this client. You can create as many archive databases as you require. When configuring an archive plan, you can assign one of these databases to it.

A Default-Archive database is automatically created during the installation and all the archive indexes will be maintained in it by default if you do not choose otherwise.

As opposed to backup indexes that are automatically saved with each backup job, archive indexes have to be explicitly saved by configuring a backup plan to do so.

6.12.1 Creating a Database

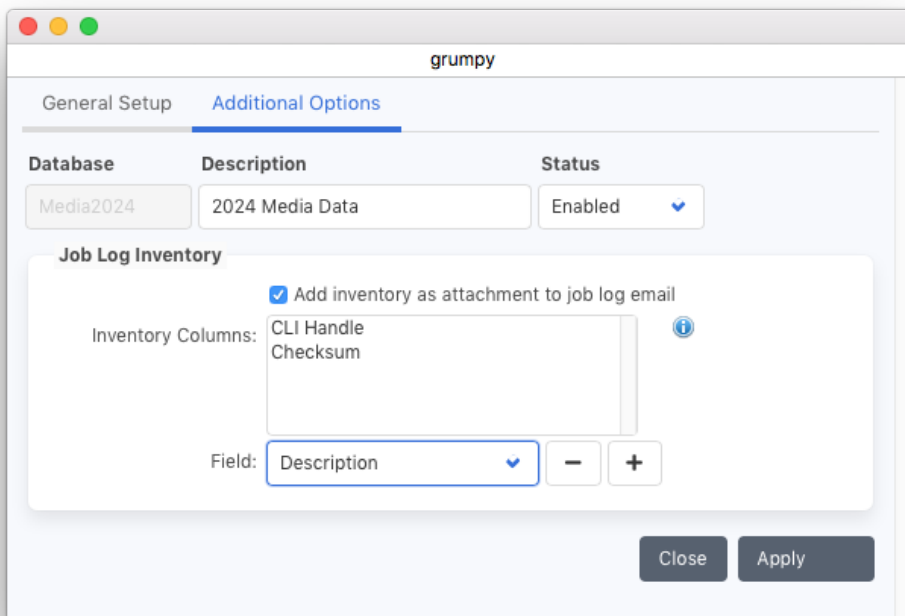
1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the Archive or Backup module from the top menu and then Manage Indexes from the Advanced Options sub-menu.
3. Click New from the bottom bar.
4. Fill in the fields of the Database dialog window.

The screenshot shows a web-based configuration window titled 'grumpy'. It has two tabs: 'General Setup' (selected) and 'Additional Options'. The window is divided into several sections:

- Database:** A table with columns 'Database', 'Description', and 'Status'. It contains one entry: 'Media2024' with description '2024 Media Data' and status 'Enabled' (with a dropdown arrow).
- Type and Usage:** Shows 'Driver type: AsxTree' and 'Usage: Archive'.
- Location:** A text field for 'Database directory:' containing 'config/index/archive/Media2024'.
- Database Backup:** Contains a checked checkbox 'Automatic backup is enabled'. Below it, 'Target client:' is set to 'Qnap'. 'Backup directory:' is '/share/Cachedev3_data/indexbackup/Media2024'. 'Cycles to keep:' is '7'. A note states: 'Note: It is recommended to keep a few cycles available, as it may be necessary to restore an older backup.'
- Scheduling:** A table with columns 'Start', 'Rebuild', and 'Frequency'. The content below the table says 'Not currently scheduled. Click New to create a new backup event.'
- Buttons:** 'New...', 'Edit...', 'Delete', and 'Start now' are at the bottom left. 'Close' and 'Apply' are at the bottom right.

6.12.2 Archive Database Options

When sending archive job mail notifications, P5 add an inventory of files archived with that job. This is maintained in the additional options.



Add inventory as attachment to job log emails:

Check this box to add an inventory to the job log. The inventory is creates as csv file, that format can for instance be imported in Excel or Libre Office Calc. Each archived file or folder is contained in a separate line.

Inventory Columns:

The inventory always contains as first column the archived file with its path. Further columns can be added:

- CLI Handle** The handle that is required to restore the file via Command line interface (CLI)
- Volumes** The list of Volumes containing the file
- Checksum** The checksum of the file. Note that this checksum can be provided only in case the Archive Plan is running with activated Verify. This checksum is consist of two parts: the checksum type as set in the media pool and the actual checksum of the file. For instance:
 xxh3_128 : 2d6b66ff8bed7318e6f7916f23e0c666
 This checksum is calculated for the data segment of the file only and can be compared to a checksum calculated with file utilities like md5sum (Linux), md5 (Mac), cksun (Linux, Mac) or certutil (Windows).
- User Metadata** Metadata added for the file

6.12.3 Archive Index Database Backup and Cleanup

P5 can create automatic backups of the Archive Index Database. Unlike Backup indexes which are saved automatically on tape, Archive indexes are saved as disk copies. If enabled, the index is saved daily at the specified time.

Target Client

Enter the client (as defined in section 6.11, P5 Clients) on which the backup shall be saved.

Backup Directory

Enter the folder where the backup shall be saved.

Start Time

Enter the time when the backup shall start. Note that there should be no archive job running which uses this index, otherwise the backup copy will be incomplete.



Starting with version 7.2, the Index Backup also triggers the archive index cleanup. After copying the index, P5 will remove obsolete index entries, for instance in case an archive tape has been deleted.

6.12.4 Archive Job Log Inventory

In this section, specify whether an inventory of the archived files will be added to the email report of after the concluded job. If enabled, P5 will add each path and file name to the inventory list. The selection box allows to add metadata fields to the inventory.



Do not forget to create a Login Area for the newly created Archive Index, otherwise the index will not be shown in the Restore section.

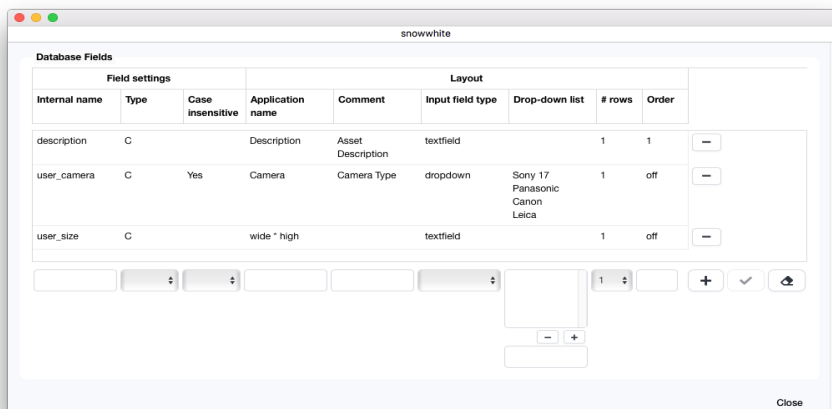
To create a Login Area, see section 6.13.1 Creating a Log-In Area .

6.12.5 Database: Functions

For each of the configured databases, there are several options you can access from the bottom bar:

6.12.5.1 Fields ...

For the selected database, this function enables you to define your individual text fields. It is only available for archive databases.



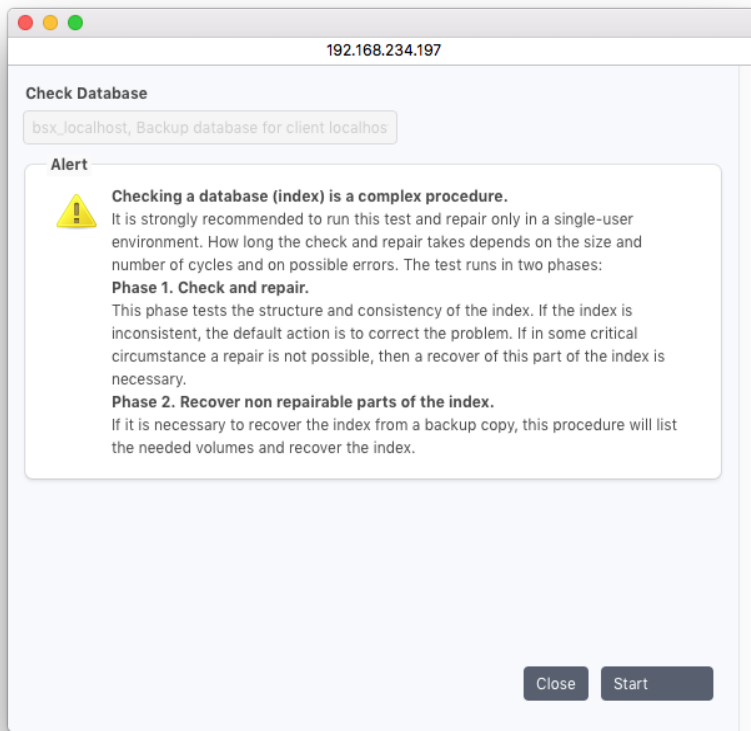
When submitting data for archiving, a plan must be selected. Among other parameters, the plan also defines the target index. Individual fields of the index are thereby available for additional tagging of the data. These fields can be used to quickly and easily locate your data within the archive index (See 9.2.4, The Search Button). The case sensitivity defines whether searching the according field is done case sensitive or insensitive.

6.12.5.2 Check Database

This function is used to check the selected backup databases, which means the backup indexes of a given client. How long it takes to complete such a check depends on the size of the index and the number of backup cycles.

It is recommended to run this process only when no other jobs which are using the selected indexes are running. In case of a backup index, problems which could not be fixed by the check function can be corrected by recovering the backup index from the volumes, since it is automatically saved with each backup.

During the second phase, the volumes required for the recovery operation will be listed in order for you to make them available to the media manager.



As opposed to the backup index, an archive index cannot be recovered automatically since it is not saved automatically. It is therefore recommended to back up the archive indexes by means of a backup plan or by the Archive Index Database Backup Function (see above), so that in the event of a corruption you can recover it from the backup.

6.13 Index Access Control (Login Areas)

P5 organizes its databases in form of a tree structure just like a directory structure of a file system on disk. Log-in areas enable access to one or more such databases. The log-in area can be compared to a Mac or a PC volume, which gives you access to a directory tree of the server. Authorization is managed through standard user/group mechanisms.

6.13.1 Creating a Log-In Area

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the Archive or Backup module from the top menu and then Access to Indexes from the Advanced Options sub menu
3. Click New from the bottom bar.
4. Fill in the required fields.

The screenshot shows a web browser window with the URL 192.168.234.197. The interface is titled 'General Setup' and 'Additional Options'. It features a table with columns for 'Login Area', 'Description', and 'Status'. Below the table, there are sections for 'Index or File System' and 'Access'. The 'Index or File System' section has a 'Selection' dropdown set to 'hfs_SuseVM (Online database for client S)' and a 'Path' input field containing '/Users'. The 'Access' section has a text input for 'Allow access for group(s):' with 'Restorers' entered, and a 'Group name:' input field with minus and plus buttons. At the bottom right, there are 'Close' and 'Apply' buttons.

Description:

Enter the desired log-in area name here.

Status:

The status flag defines the availability of this log-in area. The possible states are:

Enabled: The log-in area is available for use by the system.

Disabled: The log-in area will not be used by the system.

Index or File System:

This field gives you the list of all the configured databases. Choose the database for which you wish to enable access.

Path:

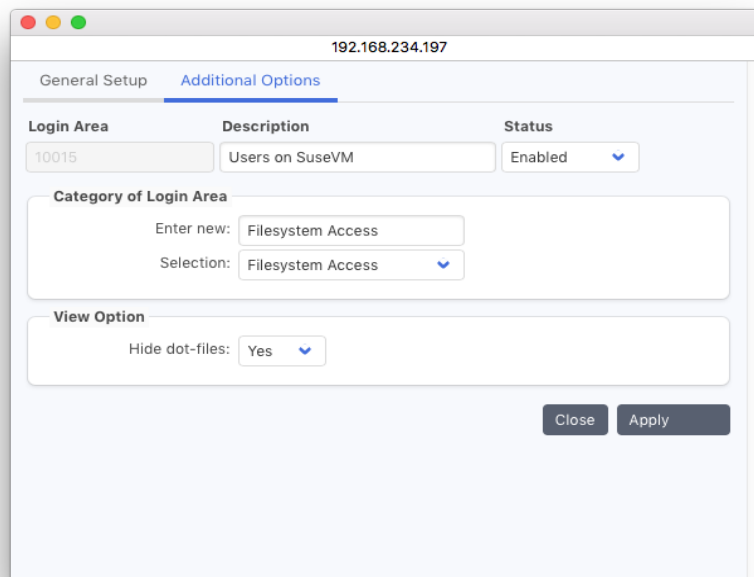
The database you choose contains a directory structure. Enter the path in this structure, i.e. the point from which you wish to enable access into it. Should this be the whole database, enter '/'.

Allow access for group(s):

Only users belonging to the listed OS groups are allowed access to this log-in area.

Group Name:

This is a utility field by which to enter the user group's. By clicking the +/- button add a group or remove it from the list. Groups are maintained in P5, see section 6.15, Group Preferences.

**Category of Log-In Area:**

Holds a list of logical categories you can choose from to simplify the navigation upon entering the data restore area.

Hide dot-files:

This option allows you to filter out all the files and directories beginning with a dot from the displayed list of objects. These files and folders are usually invisible on Unix operating systems.

yes: The filter is active.

no: The filter is not active.

6.14 User Preferences

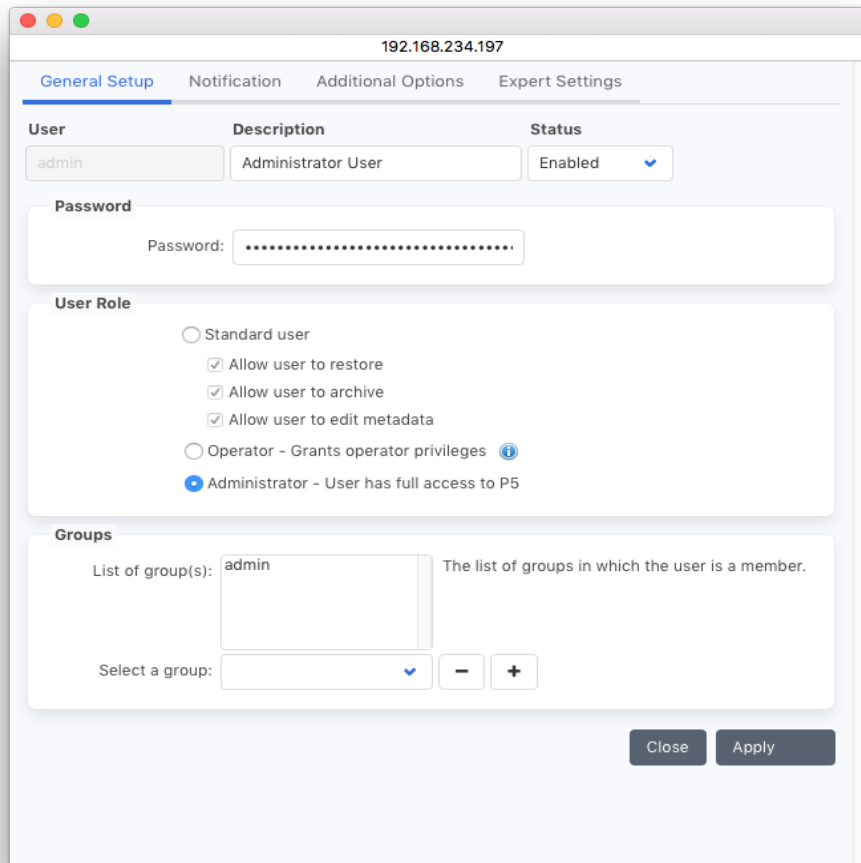
Starting with version 6, P5 manages user accounts internally in the user preferences.

The user preferences define the user account, the access password and the user's role in P5. Here, certain functions are allowed or denied to users, for example 'archiving'. Users initially inherit their preferences from the generic user.

P5 tracks and preserves all the actions the user makes while working with the application. This makes it possible to continue the user's session where he left off last.

To manage User Preferences:

1. Log into P5 using a web browser as admin user.
See section 4.1, Starting the Browser.
2. Select one of the modules from the top menu and then User Preferences from the Advanced Options sub-menu.
3. Click New in the bottom bar
4. Fill in the required fields.



User Name:

Enter the desired user name here. The users log-in name may not contain a colon or a blank.

Description:

You can describe the user here.

Password:

Set a login password here. P5 authenticates the user against this password.

Status:

The status flag defines the availability of this user's preferences:

Enabled: The user preference record is available for use by the system.

Disabled: The user preference record will not be used by the system.

User Role:

Each user can individually be granted different privileges to execute different operations. These privileges are always associated with a log-in area.

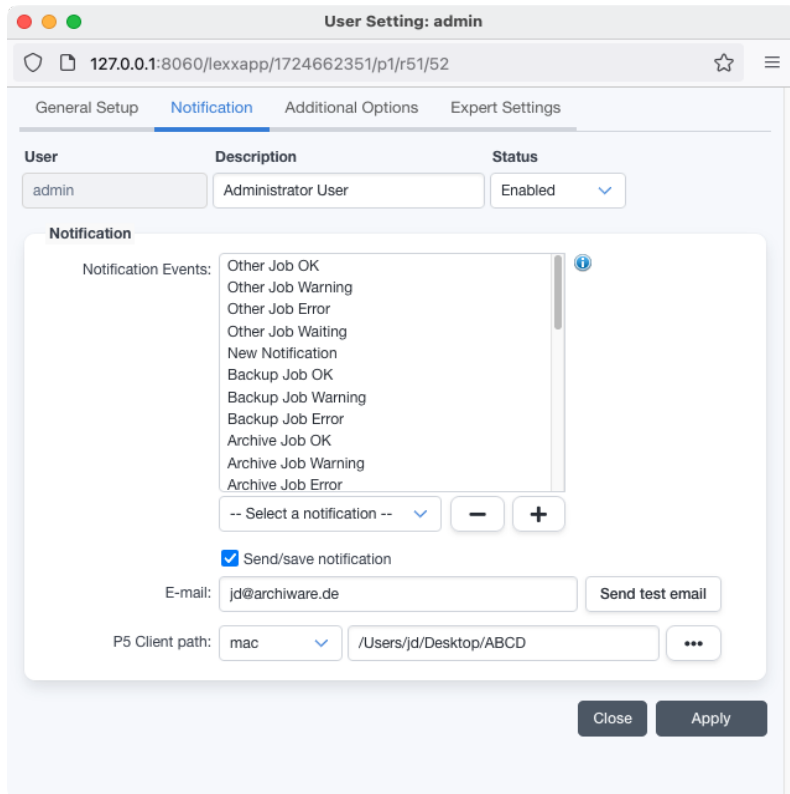
Standard User - Allow Restore:	Allows the user to restore data from the backup or the archive.
Standard User - Allow Archive:	Allows the user to archive on-line data.
Standard User - Allow Metadata Editing:	Allows the user to add or edit metadata in the archive index.
Operator:	Grants operator privileges, i.e. start and stop jobs and change tapes
Administrator:	Grants full access to the user
Allow administrator to change the folder structure of an archive index	Allows administrator to edit folder structures in the Archive Index Browser.



Caution: The settings of the *Generic* user account are used as the default template for creating a new user record. It is recommended not to set a password for that user and not check-mark the "Standard user" field

Starting with version6, P5 does no longer use or support system (OS-) accounts. Users that are not defined in the User preferences cannot login to P5. Also the user permissions in P5 are not related to permissions in the OS anymore.

Users preferences – Notifications



Notifications provide for sending E-Mails when one of the selected events in P5 occur. A notification is sent to each user whose setup matches the event. Unlike in previous versions, different notifications can be sent separately to different users.

Send /save notifications

Check this box in order to enable E-mail notification, notifications to file or both for the user.

Notification Events

Define which events are to be sent. Any event added to the table will cause an E-mail to be sent when the matching event occurs.

E-mail notifications – E-mail

The users E-mail address. Set this address in order to send email notifications

E-mail notifications – Send test E-mail

Use this button to send a test e-mail to the configured user.

File notifications – P5 client path

Specify here a client host and a folder on the client where job logs and lists are to be stored. A notification will then create a text file in the format *JobID_job.txt*, in case a job inventory is configured, an additional file *JobID_job_inventory.tsv* is created.

How notifications are sent:

Jobs executed by administrator users or by the P5 scheduler send notifications to all administrator users.

Jobs started by operator users send notifications to all operator users.

All other jobs send notifications to the user who started the job.

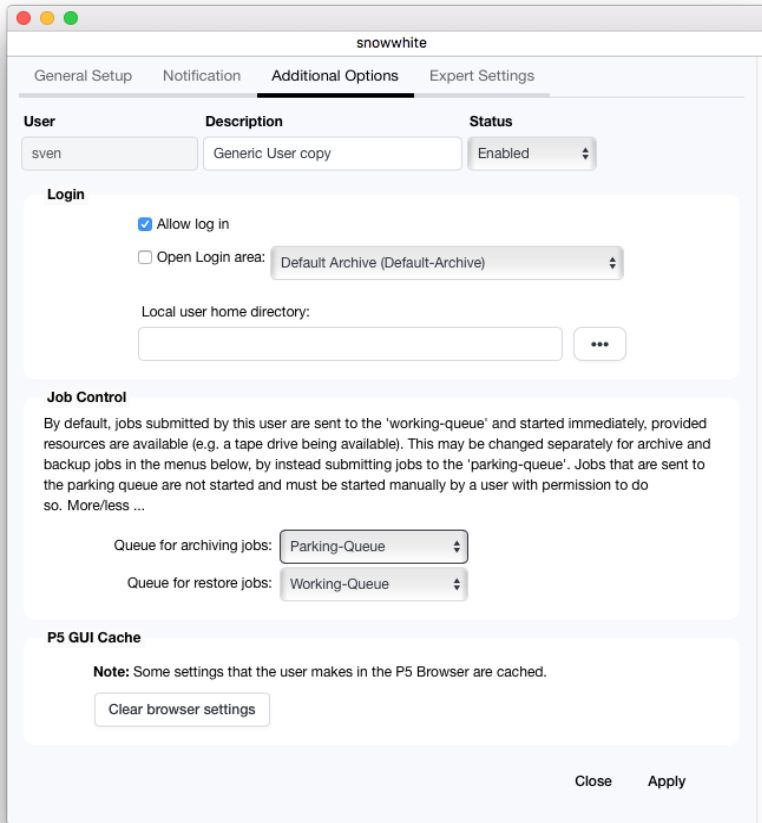
The following events are defined to send notifications:

The notification	Is sent when:
Archive/Backup/Sync Job Ok	An Archive/Backup/Sync job ends OK
Archive/Backup/Sync Job Warning	An Archive/Backup/Sync job ends with warning
Archive/Backup/Sync Job Error	An Archive/Backup/Sync job ends with error
Other Job OK	An other job (different from the above) ends OK
Other Job Warning	An other job ends with warning
Other Job Error	An other job ends with error
Other Job Waiting	A job (any job, including Backup, Synchronize and Archive) is in a wait state and cannot continue.
Index-Backup Job OK	An automated Index-Backup job ends OK
Index-Backup Job Warning	An automated Index-Backup job ends with warning
Index-Backup Job Error	An automated Index-Backup job ends with error
Label new Volumes	Additional volumes are required
Volume Recycled	A volume has been recycled
Device Needs Cleaning	A tape drive blocks and requires cleaning
Device Needs Check	A configuration problem with a device is detected
Volume Suspect	A volume state is set to suspect (due to an error)
Restore Job Ok	A Restore job ends OK
Restore Job Warning	A Restore job ends with warning
Restore Job Error	A Restore job ends with error



In order for e-mails to be sent correctly, a valid mail server must be defined in the setup of the localhost client. See section 11.3, Customizing Reporting

Users preferences – Additional Options



Date/Time Format:

Offers a list of possible date/time formats to choose from .

Allow log in:

Check this box to allow this user to log in into the P5 interface.

By disabling this check box in the generic user's preferences, new users will not be automatically allowed to log in

Login Behavior:

Here you can define what the application will do when a user logs in. By default the full interface is presented. By setting a log-in area, the user's access is restricted to that log-in area.

Local user home directory:

Set here the user's home directory on the P5 server host. This folder is then added and shown as home folder when browsing in the file system, for instance when manually archiving files.

Queue for archiving jobs:

Queue for restore jobs:

Define here how archive and restore jobs started by the affected user shall be started.

Parking queue: Jobs are enqueued as parked jobs, i.e. not started automatically.

Working queue: Jobs are started directly

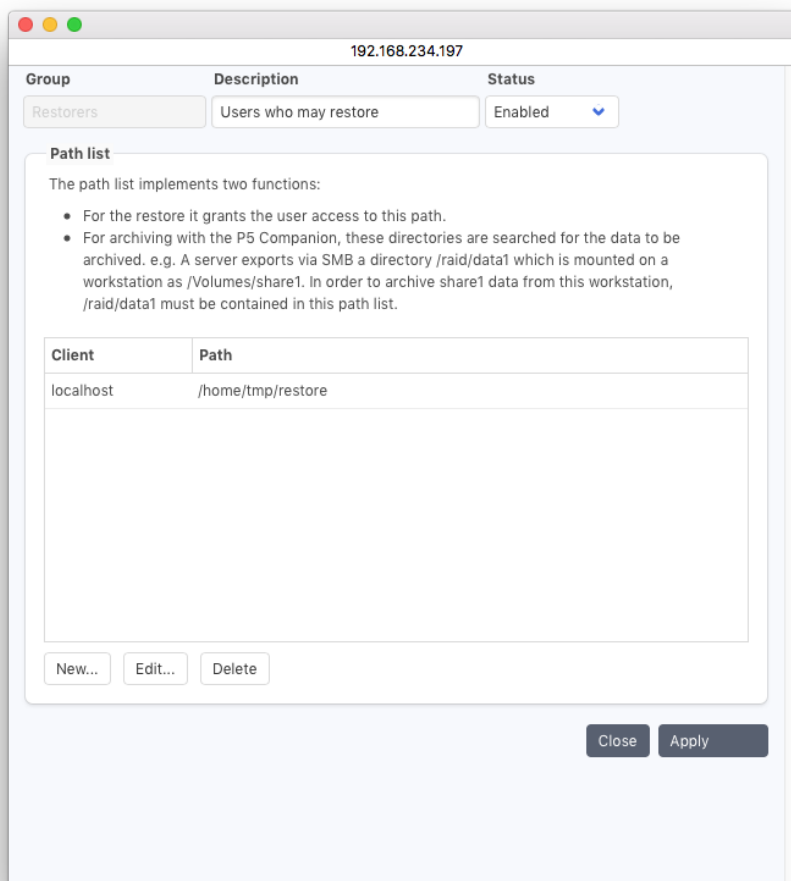
See section 6.10.1, The Job Monitor for how to handle parked jobs.

6.15 Group Preferences

User groups are used in P5 to maintain access restrictions for users. From the group preferences area, you can allow or deny access to the filesystem when restoring. Groups are also referred to when setting up Archive Plans and Index Access Control (Login Areas), so groups are also used to limit access to filesystems, indexes and archiving abilities.

To manage Group Preferences:

1. Log into P5 using a web browser as admin user.
See section 4.1, Starting the Browser.
2. Select one of the modules from the top menu and then Group Preferences from the Advanced Options sub-menu.
3. Click New in the bottom bar
4. Fill in the required fields.



Group Name:

Enter the desired group name here. The group name may not contain a colon or a blank.

Description:

You can describe the group here.

Client and Path:

Add here those paths on the client computers where this group may restore files to.

By default, no path is set in the group preferences. In case the path list is empty, no restrictions are set.

In case a user is member of multiple groups and none of the groups contains a path restriction, the user is allowed to restore to any place in the file system.

Users may belong to multiple groups, so it is possible to define restore paths and access to indexes according to multiple schemes. For instance different departments may have separate restore paths and different indexes, so different user groups are defined for both departments. Assigning a user to both groups will allow that user to restore and/or archive for both departments.



There is a special case implemented in P5: When restoring and after resolving all group memberships there is only one path available for a user to restore to, exactly that path is set automatically as restore path.

Otherwise a user may select from the paths and also specify a target path in a subfolder of the given paths.



To simplify understanding of permission setting:

- User and Group preferences define access permissions
- Restrictions in Login Areas are for browsing only and restrict read access

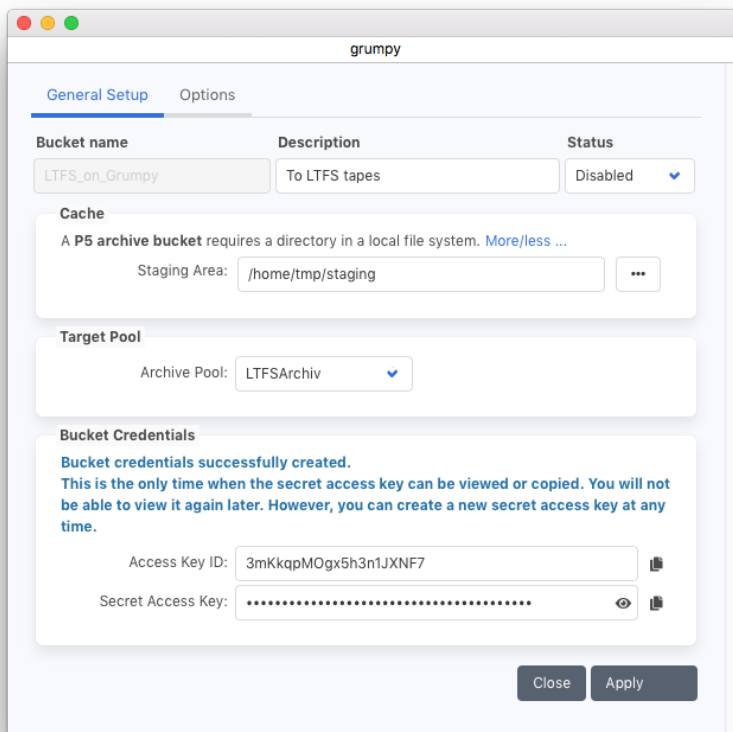
6.16 S3 Object Archive

P5 Archive provides an S3 Object Storage service, presenting an S3 end-point for other software to connect to via the http/https ports. Compatible with Amazon S3 Glacier Flexible & Deep Archive storage classes, objects received via the S3 interface are archived to LTO/LTFS, with full management of tapes, drives and libraries.

6.16.1 Creating an Archive Bucket

To create a S3 Object Archive:

1. Log into P5 using a web browser as admin user.
See section 4.1, Starting the Browser.
2. Select the Archive module from the top menu and then S3 Object Archive from the Advanced Options sub-menu.
3. Click New Archive Bucket in the bottom bar
4. Fill in the required fields.



Each Archive Bucket is represented by a bucket in the S3 API. In P5, the bucket points to an Archive Pool (see section 6.3, The Media Pool) where the data is stored.

Bucket name:

The bucket in the S3 API that must be used.

Description:

Describe the Object archive here.

Staging area:

A folder on the S3 server host used as temporary storage. Data written to or read from the bucket is stored here intermediately. Make sure there is sufficient space to allow a bigger write and a bigger read operation. The effective space requirement vary with the settings in the options page, see section 6.16.2, Archive Bucket Options.

Archive Pool:

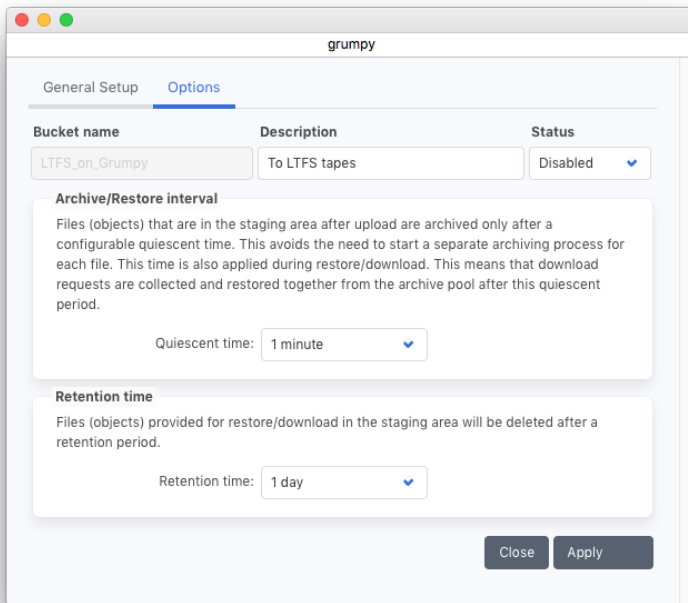
The archive pool to be used to store the data. The archive pool also defines via its type whether the data is stored on tape, disk or in the cloud. See also see section 6.3, The Media Pool.

Bucket Credentials: Access Key ID and Secret Access Key:

Access Key ID and Secret Access Key are required as authentication via the S3 API. Note that the secret access key can only be visualized after creation. Later, the saved key will not be readable anymore. It is however possible at any time to create new credentials that can then be read.

6.16.2 Archive Bucket Options

The Bucket Options allow defining the timing behavior of the underlying archive and restore mechanisms.



Archive/Restore interval Quiescent time:

This time defines how long P5 will wait for further data before the archive process is started. Smaller intervals may result in more separate archive jobs while bigger intervals may require more space in the staging area. The interval is also applied to download / restore requests.

Retention time:

The retention time defines how long data restored for download via S3 remains in the staging area before it is deleted.

6.16.3 S3 API Capabilities

The P5 Object Archive is built for compatibility with the AWS S3 API and is tested with AWS S3 for expected functionality and behavior. The P5 Object Archive provides no guarantees for other S3-compatible services, as their S3 API implementation is unknown and therefore potentially unsupported. While the P5 Object Archive commands may work as documented, any such usage is at your own risk.

For details on the S3 Server API implementations, please read more here:
https://blog.archiware.com/redoc/p5_s3_server_api/awp5s3api.html

6.17 LTO Encryption Key Management

This section in P5 manages the LTO encryption keys used in P5.

The LTO encryption is a hardware encryption performed directly in an LTO drive.

This encryption is not related to the client-specific encryption that can be setup in the localhost client definition (see 6.11, P5 Clients), but can be enabled in the media pool (see 6.3.6, Tape Based Cloud Pool: LTO Encryption).

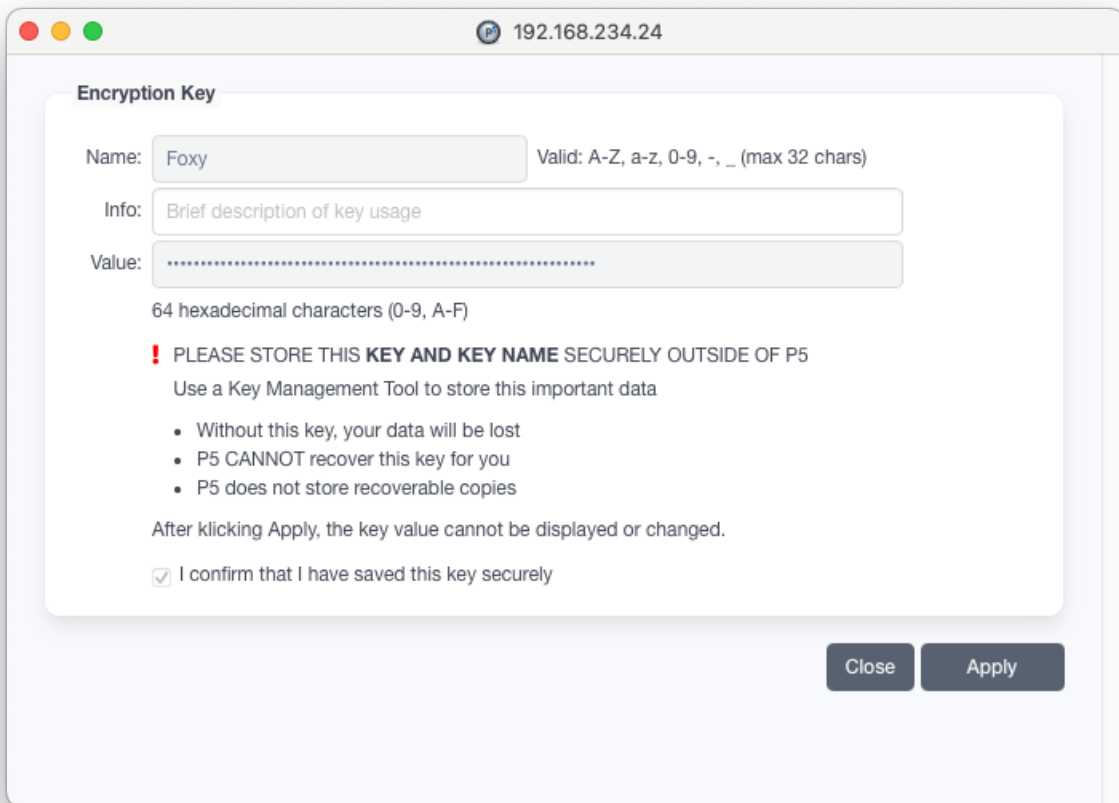
When enabled, the LTO Encryption is applied to all data written to the tape.

The hardware encryption can also be used with LTFS formatted tapes.

6.17.1 Creating an Encryption Key

To create an Encryption Key:

1. Log into P5 using a web browser as admin user.
See section 4.1, Starting the Browser.
2. Select the **Archive** module from the top menu and then **LTO Key Manager** from the **Advanced Options** sub-menu.
3. Click **New** in the bottom bar
4. Fill in the required fields.



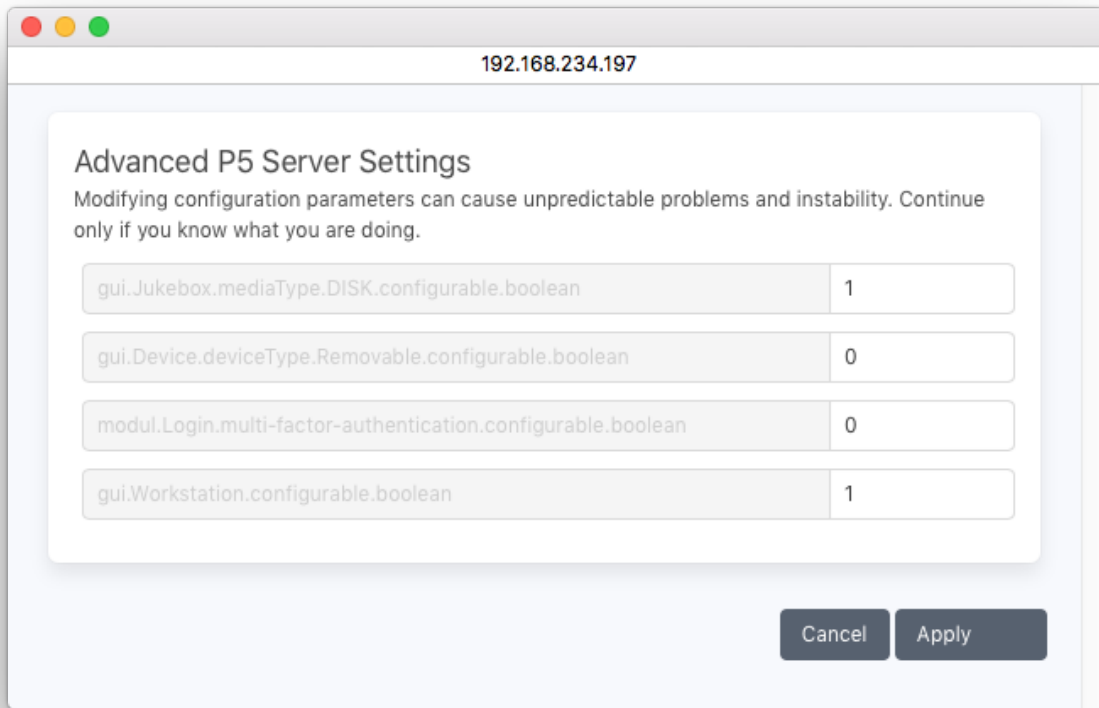
PLEASE STORE THIS KEY AND KEY NAME SECURELY OUTSIDE OF P5.

In case the key is lost, the data on tape cannot be restored.

Unlike the configuration data, this key is not automatically saved on tape.

6.18 Advanced Server Settings

In the P5 menu, the item Advanced Settings leads to server wide setting that can be configured:



Note that after changing a server wide setting, it is required to log out and login again to P5 in order to activate the change.

6.19 Lifecycle Management

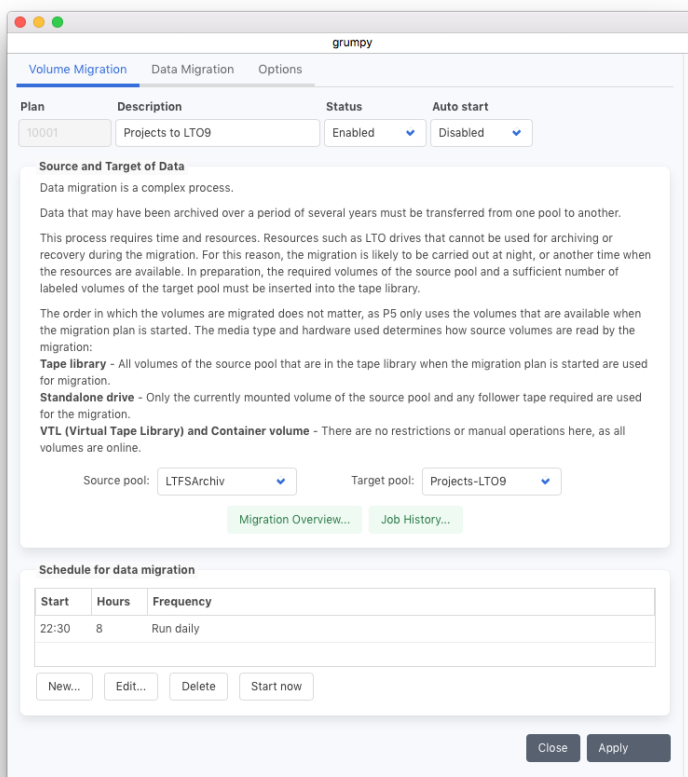
'P5 Archive DLM' (Data Lifecycle Management) is an additional module for managing the lifecycle of data previously archived with the P5 Archive module. It provides migration plans to migrate/replicate data from old to a new storage and the verify plans to check existing data on tapes is readable. The P5 Archive DLM module is a new name for P5 Data Mover, and includes the new verification feature.

6.19.1 Migration Plan

The migrations plans allow the migration of archived data from one set of media to another. For instance when a newer tape generation might be used, the data from the old tapes could be migrated, or in order to move or copy data from tape to cloud or to disk, or other way round.

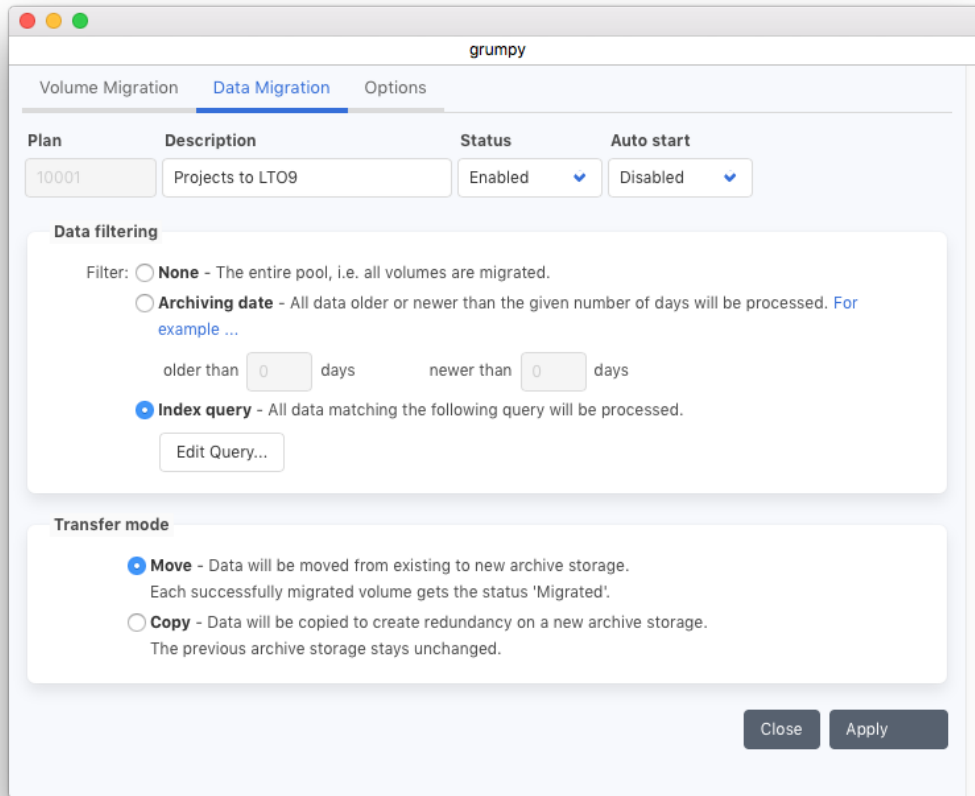
The migration is available in the left side menu of the archive module.

Click **New** to create a migration plan:



- Description:** Here it is possible to describe the migration plan
- Source Pool:** Select here the archive pool that shall be migrated
- Target pool:** The target pool defines where the data shall be migrated to

Migration Plan Data



Filter :

It is possible to activate a filter to include only new files or only old files:

- Older than X days* Include only files that have been archived more than X days ago.
- Newer than X days* Include only files that have been archived within the last X days

Index Query:

The index query allows a more detailed selection, which data shall be migrated. The query is applied to one archive index, the following queries are possible:

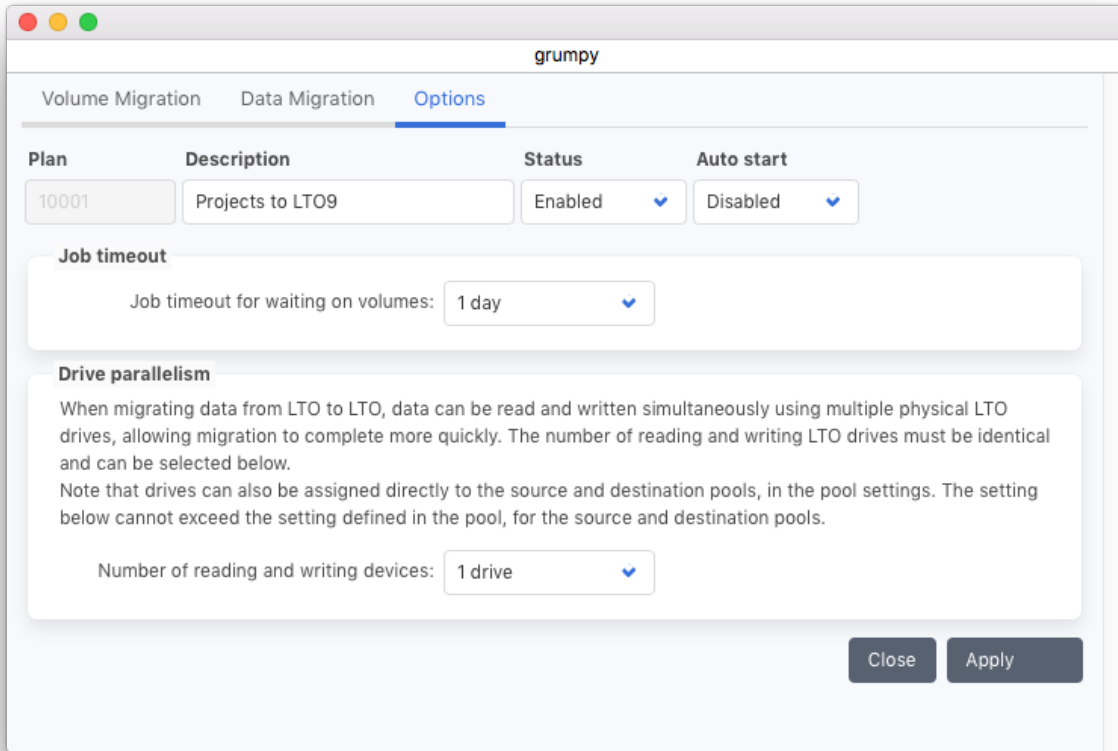
<i>Directory Path</i>	Select files and folders within one directory in the index
<i>Name</i>	Select files containing a specific word or string The string comparison matches any string contained in the file name. For details regarding the "contains word" comparison, see section 9.2.4, The Search Button.
<i>Archive Date</i>	Select files archived at, before or after a given time
<i>Modification Date</i>	Select files last modified at, before or after a given time
<i>Size</i>	Select files by their size
<i>Metadata fields</i>	Select files by their metadata

Transfer mode:

The mode defines what happens with the old data after the migration:

<i>Move</i>	The old files and folders are invalidated and may be removed with the next subsequent cleanup.
<i>Copy</i>	The old data is kept unchanged.

Migration Plan Options



Job timeout:

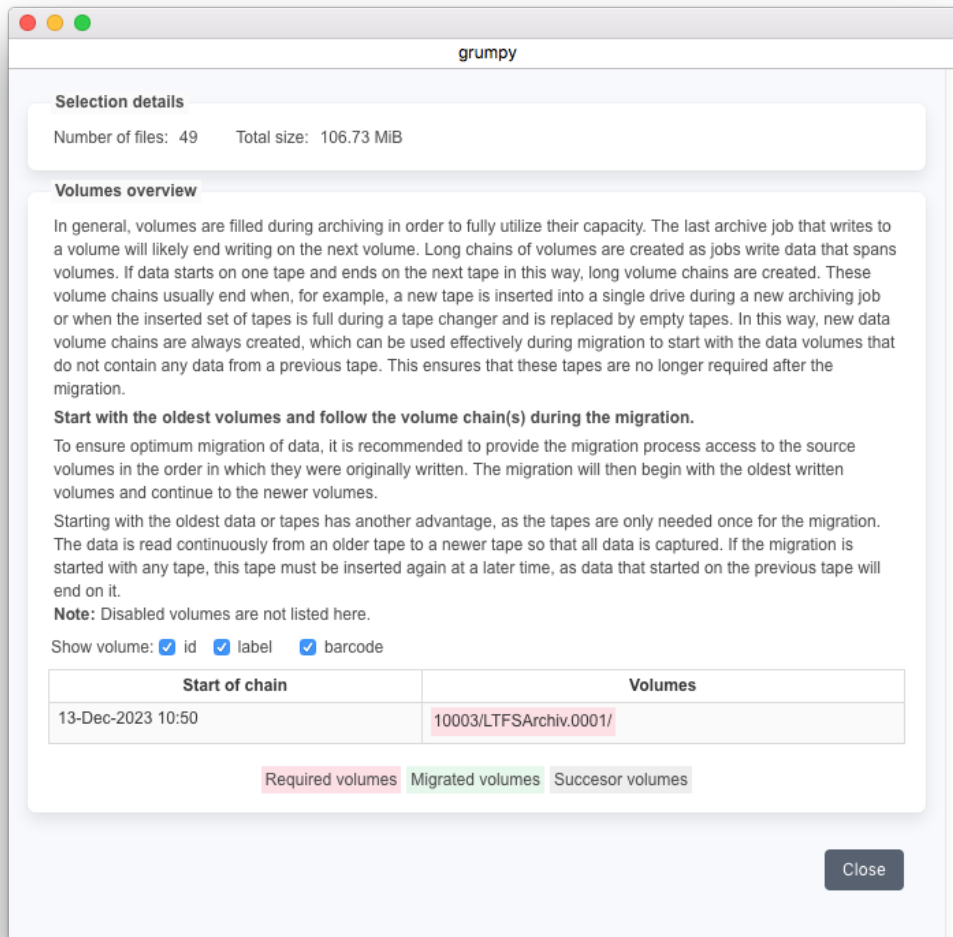
This timeout defines how long migration jobs will wait for writeable volumes of the target pool.

Drive parallelism:

When migrating data from one tape pool to another, it is possible to use multiple drives in parallel. This requires per configured parallel stream one drive for reading and one drive for writing (or two drives when cloning is active). So for a parallelism with two drives, at least two drives must be available for readign and two more for writing.

Migration Overview

The migration source windows given an overview about the selected data for the migration, its size and the required volumes.

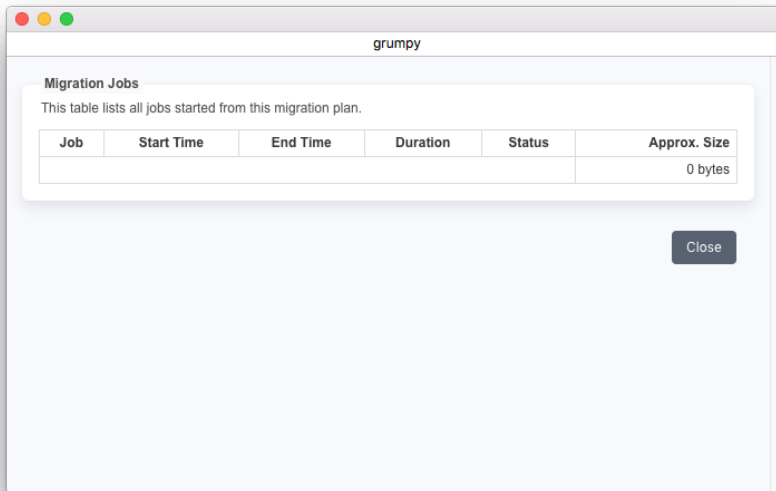


The volume list shows how the single volumes are chained together and which volumes are already migrated. Volume "chaining" appears when files are split over volume boundaries. In order to read and migrate such splitted files, both volumes are required, the first written volumes and it's follower. So when planning the migration, it makes sense to make the volumes of one or mutliple volume chains available for the migration by putting them online.

When starting the migration, each job will start with those volumes that are online. In case a volume has a follower tape, then also the follower must be online in order for the tape to be selected for migration.

Migration Jobs

The migration Job log will show the jobs of that migration plan and how much data was migrated.

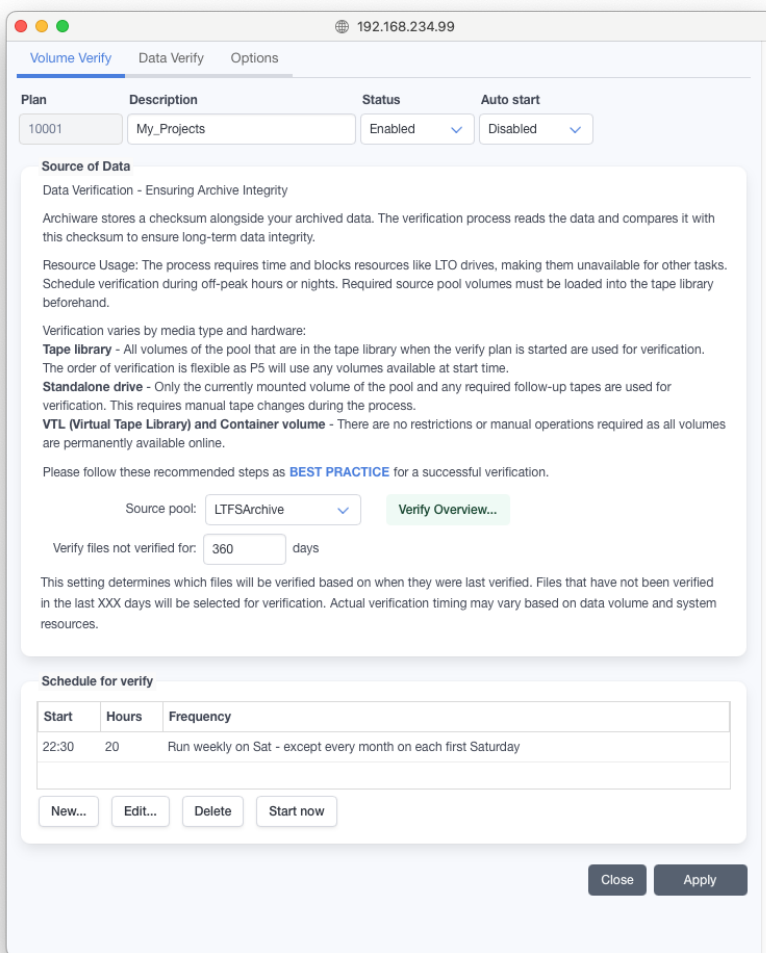


6.19.2 Verify Plan

A verify plan can be defined in order to automatically verify tapes that are online and were not verified over a given period. During verify, the tape is read and the checksum of the read data is compared to the saved value.

The verify plan is available in the left side menu of the archive module.

Click **New** to create a Verify Plan:



Description: Describe the verify plan

Source Pool: Select here the archive pool to be verified

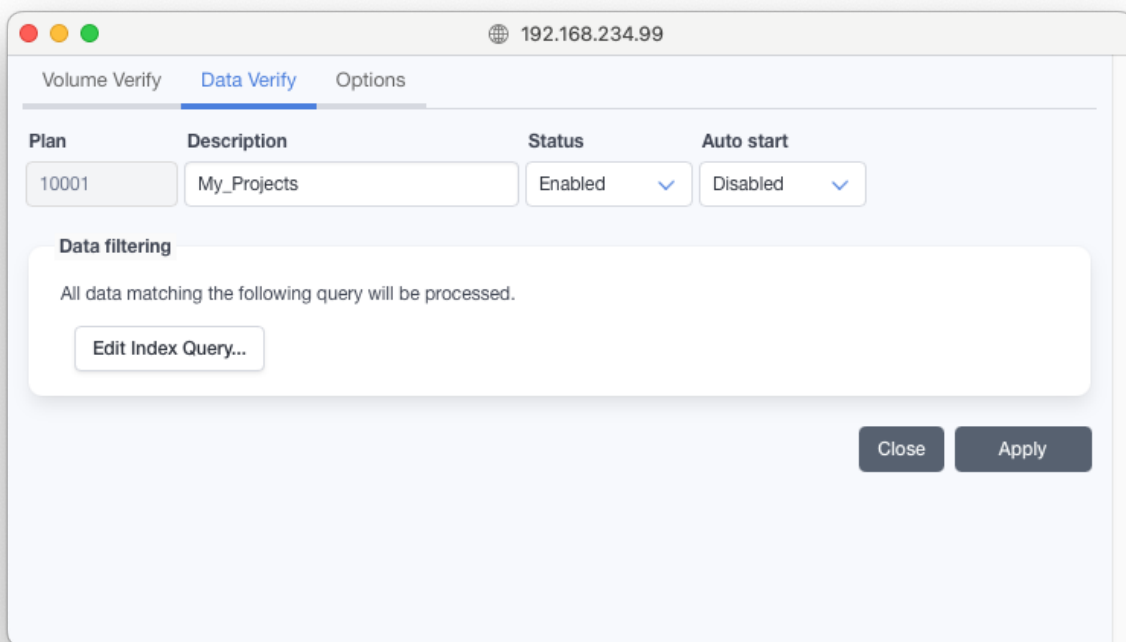
Verify files not verified for Define an interval here, how often files should be verified

The verify interval defines the duration after which a file will be verified again. As a consequence, tapes available for verify are ignored if these were checked within the given time.

For tapes, the verify requires reading the tape, usually in their entirety. That in turn reduces the lifetime of the tape, as the number of end to end passes on a tape is limited. On the other hand, the verify plan will make sure that the media is still readable. The number of verify runs on one tape must therefore be balanced between the knowledge that the tape is readable and the tape lifetime reduction caused by the verify. Verifying once per year may be a good interval, assuming the tape can be read or written approximately 200-250 times.

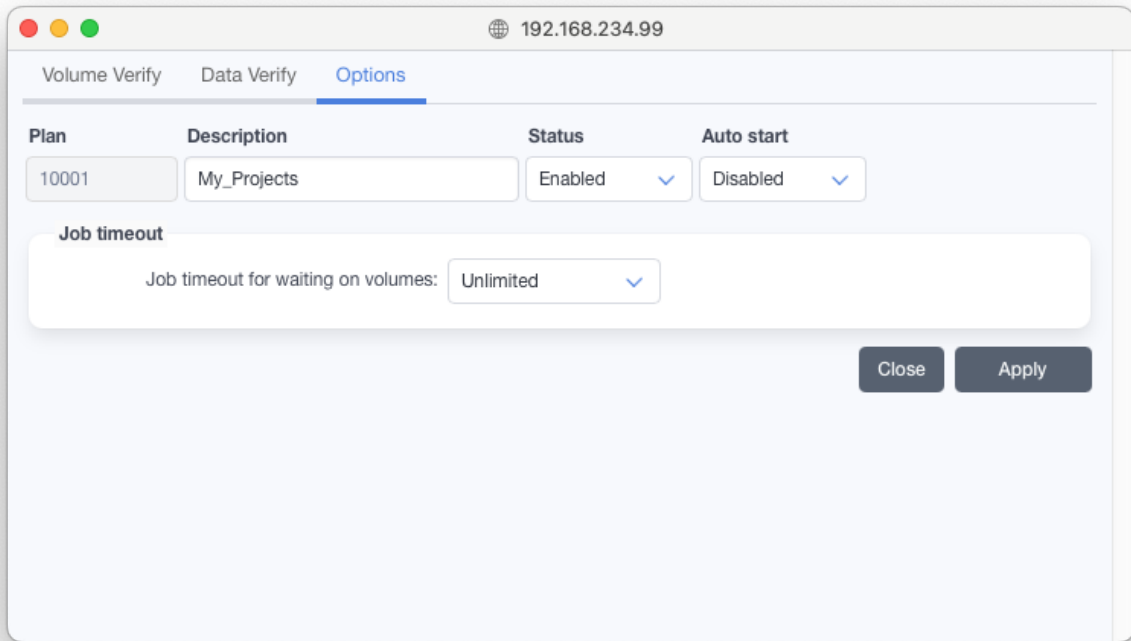
The verify job, when started with a connected tape library, will select tapes which are online available and were not checked over the defined time interval.

Verify Data



The verify plan can be restricted to the files of a specific archive index and can further be filtered, for instance for a specific folder tree or archive date only.

Verify Options



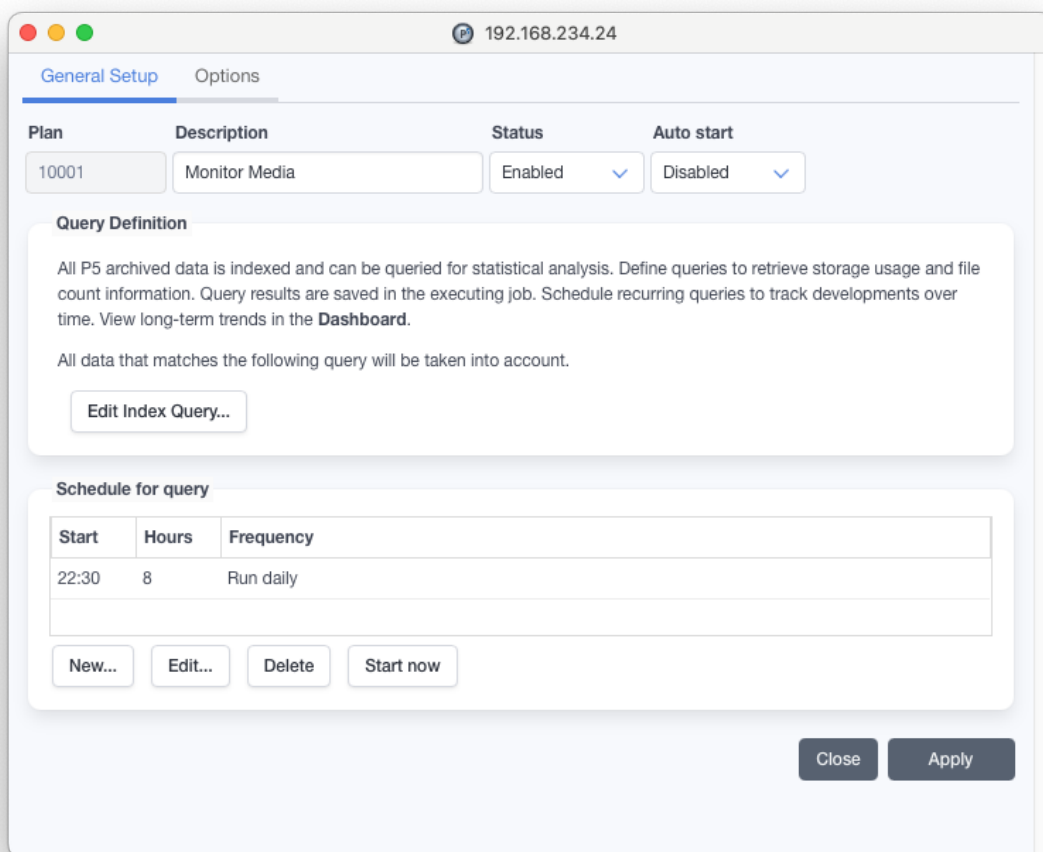
Job timeout for waiting on volumes Define here how long to wait for the next volume before the job will time out.

6.19.3 Utilization

Utilization Plans allow monitoring of sub-sections of Archive Indexes. Queries are created, and the total size of files matching the query is logged in a regular basis, e.g. daily. This allows analysis of archive storage use on a per-project basis, where a projects storage can be identified via a query on an Archive Index.

Utilization Plans only perform queries, they do not change the Archive Index or archived data.

Click **New** to create a utilization plan:



Description:

You can describe the utilization plan here.

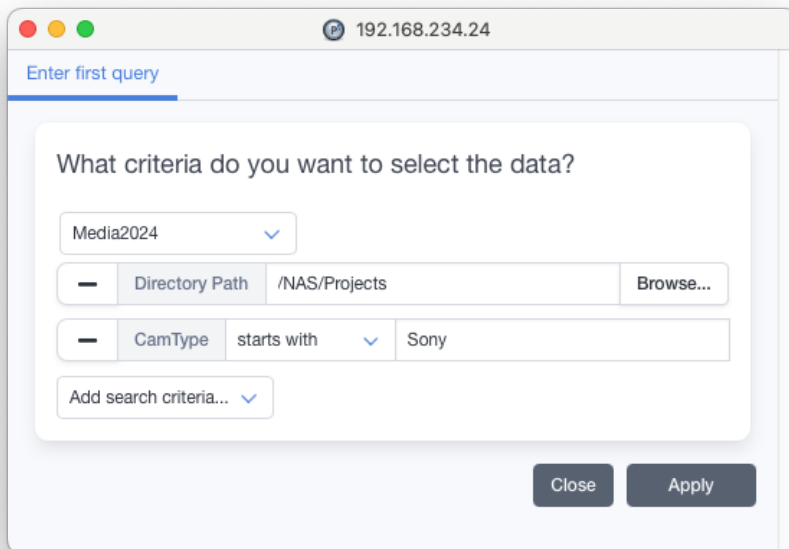
Status:

The status flag defines the availability of this plan. The possible states are:

- » Enabled: The plan is available for use by the system.
- » Disabled: The plan will not be used by the system.

Utilization Plan Query

The index query defines what exactly the utilization plan shall look up in the index.



It is possible to define multiple criteria in the index query. The found files must fulfill all given criteria.

When for instance one index criteria is *Directory path* = /NAS/Projects and there is a second criteria with a user field *CamType* contains SONY, then the query will check for files with the given metadata in the folder /Volumes/Mydisk and below.

In addition, it is possible to define multiply Queries in one utilization plan. The results of all defined queries are summed up when the utilization job is run.

Multiple queries can for instance be used to query for files in multiple indexes, or for different folders in the index.

Utilization Plan Schedule

the utilization schedule defines when the utilization shall be executed.

The screenshot shows a dialog box titled "192.168.234.24". It has three main sections:

- Start:** Time: 22:30, Date: 13-Nov-2025. Below it, "Stop after: 8 hours".
- Frequency:** Radio buttons: None, Daily, Weekly (selected), Monthly. Below: "Run query every 1 week(s) on:". Checkboxes: Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday (checked).
- Exception:** Radio buttons: None (selected), Daily, Weekly, Monthly.

Buttons: Close, Apply.

Start

Specify the date and time when to run the utilization plan. The start date is used as the calculation base for the frequency (see below).

Stop after:

This value is just here as a generic scheduler is used. Utilization plans are fast and just contain index queries that cannot be stopped.

Frequency:

You can choose between four different calculation bases:

- None** The schedule event will appear only once, at the specified start time
- Daily** The event will be triggered every n days
- Weekly** The event will be triggered every n weeks, the day(s) of the week can be specified.
- Monthly** The event will be triggered every n months, the week and the day of the week can be specified.

Exception:

An exception can be specified in the same way as the frequency. If, for example a migration must be skipped when a full backup is planned, configure an exception to cover that time.

7 LTFS Exchange Functions

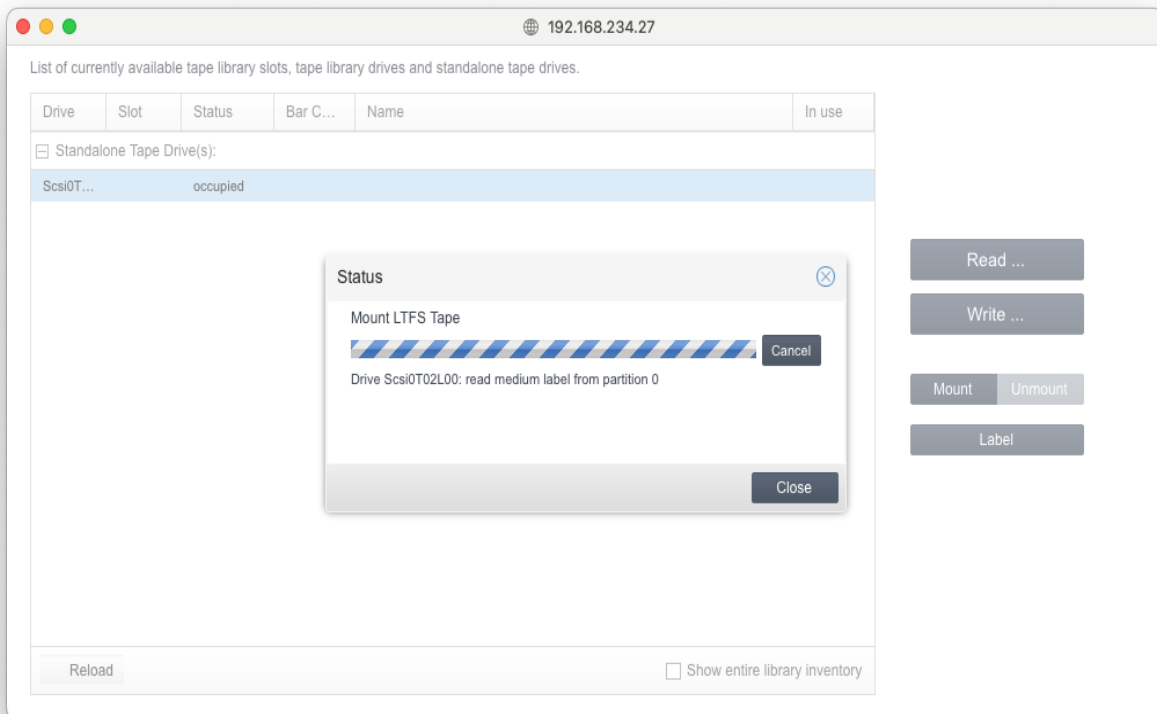
P5 comes with built-in support for LTFS, the Linear Tape File System. That means there is no separate LTFS driver required to access LTFS media.

In addition to archiving to LTFS tapes provided, it is possible to use LTFS tapes to exchange data with other systems. The LTFS exchange functions provides functions for the data exchange.

The LTFS exchange functions support formatting, mounting, reading from and writing to LTFS tapes. Tapes formatted and used in the LTFS exchange module are intended for data exchange only. These tapes are not maintained like local volumes and will not appear in the P5 Volume table, nor will their contents be indexed in the Archive Index.

To access the LTFS exchange functions:

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Open the P5 menu in the upper right edge and select LTFS Exchange.



To access LTFS tapes, an LTO-5 drive or newer and at least LTO-5 media are required.

7.1 Mount / Unmount

In order to access an LTFS tape, the tape must be mounted first. In a tape library, that means it is transported into a drive and physically mounted there. In a single tape drive the tape must be inserted into the drive manually.

In P5 this also means that the tape is made known to the program, so it can be accessed for reading and writing. Before removing the tape, click **unmount** to make sure the tape is no longer accessed.



Removing an LTFS tape from the drive without unmounting it may result in the tape being unreadable as the latest LTFS index is not written on tape.

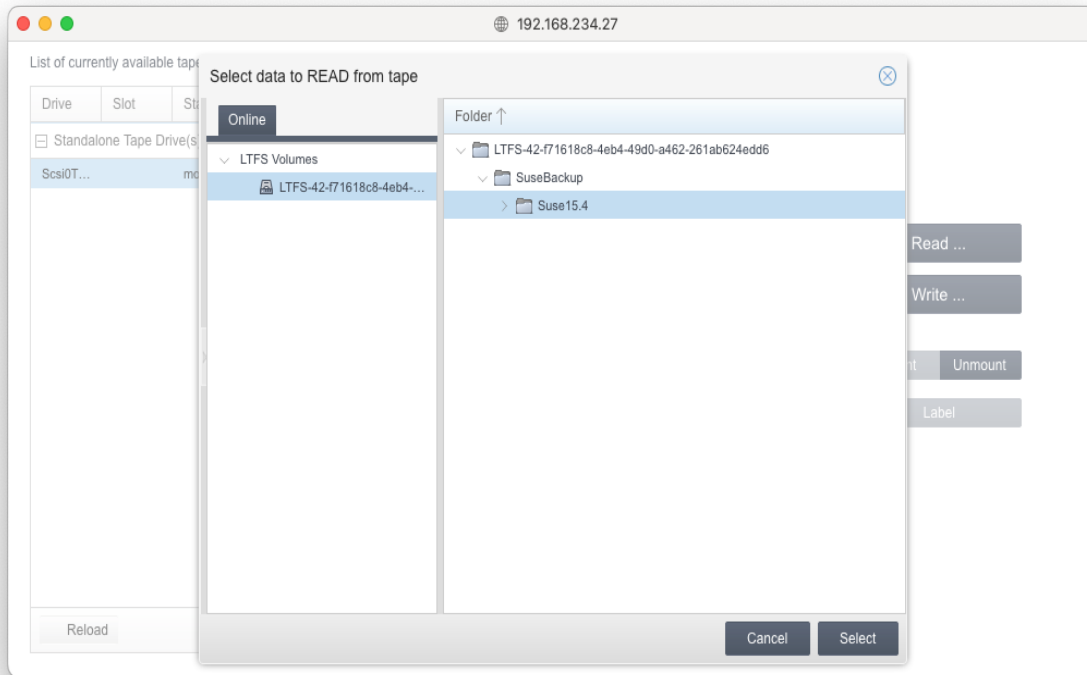
7.2 Label

Labeling a tape for LTFS means the tape is formatted and partitioned, the previous content of the tape is lost by that step.



When using LTO-9 tapes, partitioning empty tapes or tapes that were not formatted as LTFS tapes before may take rather long, 30 minutes or more.

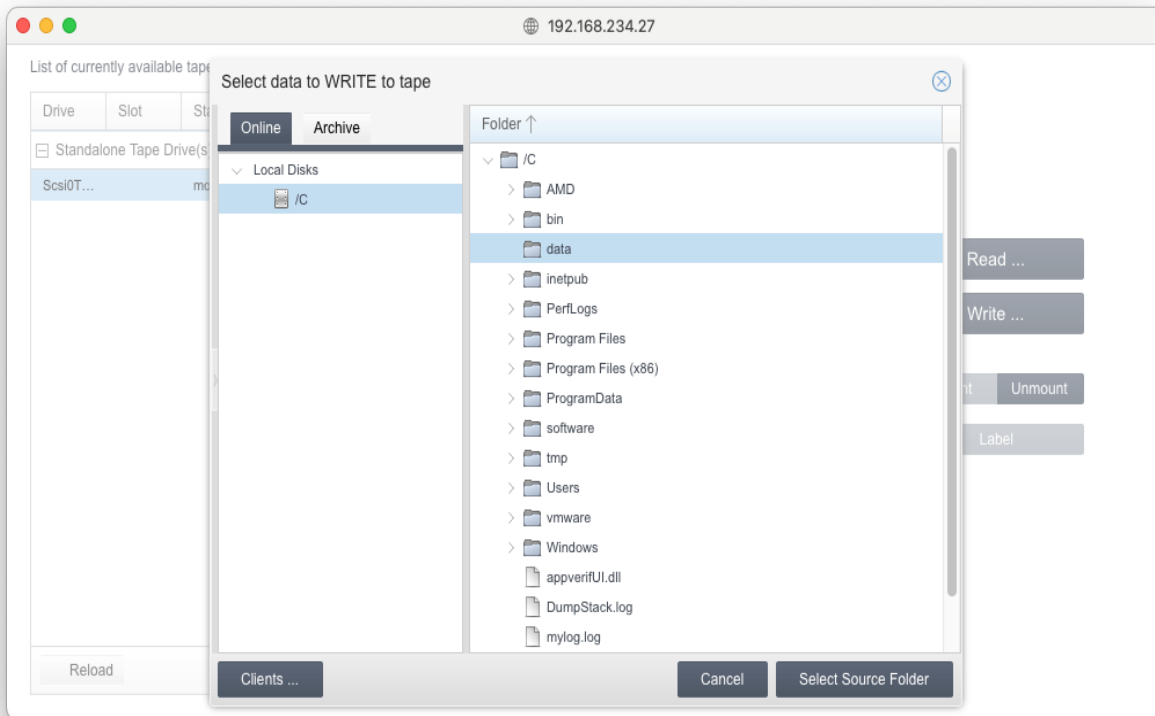
7.3 Read from LTFS



P5 does not expose the LTFS tape into the filesystem, but access to files and folders is granted through the P5 GUI. Select the files and folders on the LTFS tape for reading as first step, then select a target folder to write to.

Note that even in LTFS format tapes are sequentially accessed. It thus makes no sense to access a tape with more than one job at a time due to the required positioning. In P5, multiple read operations can be configured, but these are performed sequentially.

7.4 Write to LTFS



P5 grants access to files and folders through its GUI. Select the files and folders on the host for writing first, then select a target folder on the LTFS tape to write to.

When writing to LTFS tapes, there are two inventory files written to the root of the LTFS file system:

A file in tab separated format containing the checksum and the file path and name. This file contains all files with a checksum that fits the currently selected checksum in the volume pool (see section 6.3, The Media Pool). This file is named `[volume label].[checksum type].tsv`, for instance `volume1.crc32.tsv`.

The second file contains the checksums and file information as ASC Media Hash List (ASC MHL) defined by the Advanced Data Management Subcommittee, THE AMERICAN SOCIETY OF CINEMATOGRAPHERS, Specification v0.10 Committee Draft. This file is named `[volume label].mhl`, for instance `volume1.mhl`.

Both files will remain invisible until the tape is unmounted and will become visible when the tape is mounted the next time.

8 Restoring Data From Backup

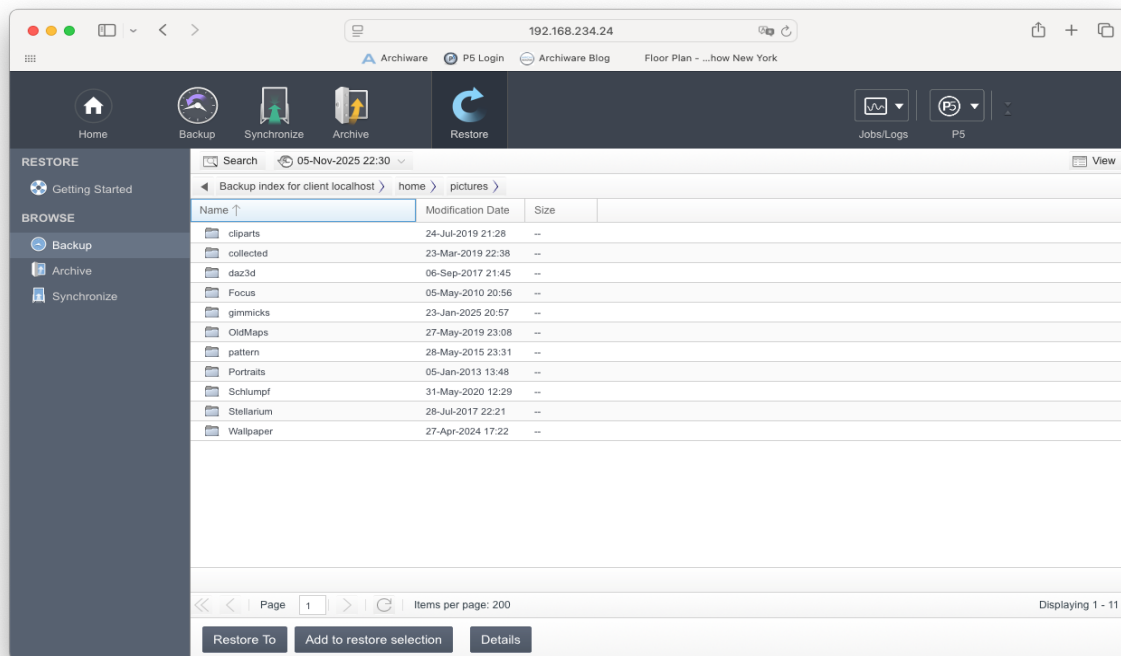
Restoring data is done through the P5 data browser. Each backed file has a corresponding entry in the index. The index tells P5 on which volume the file is located and where, and carries other meta-data describing the file. A machine (client) that was never backed up will have an empty backup index.

You can search and browse through backed up directory structures with the data browser. Data can be restored to any of the states that were backed up. How many versions of a file will be available for restore depends on your setup. This depends on the number of full backups that have been done and the number of times the file had changed in-between.

When restoring files back to their original position in the file system, a conflict between the current file on the file system and the version from the backup may occur. You can instruct P5 on how to resolve such conflicts before starting a restore. Alternatively, you can tell P5 to restore the data to another location (directory) in the file system where such conflicts will not occur.

Starting with P5, the index subsystem shows snapshots to represent the saved folder tree on disk as it was when the backup was done. Note that each backup creates a new snapshot, and each snapshot represents all saved data of the selected client. This does not imply that a new copy of each file was created when the backup job ran and the snapshot was created.

By default, P5 will always represent the most recent state of saved files in the restore screen.



8.1 Restoring Files

Follow these steps:

1. Log into P5 using a web browser. See section 4.1, Starting the Browser.
2. Navigate into the backup index where the data that you wish to restore is located. Select **Restore** in the top menu, then choose **Restore from Backup** from the left sidebar. The main screen will show the backup indexes available to restore from.



You can limit the access to backup indexes as needed. A log-in area for each index is required to grant or deny the access to that index. See section 6.13, Index Access Control (Login Areas)

3. Navigate to the folder where the files and/or folders that you wish to restore are located. Select one of the files by clicking its name. Then select multiple elements by holding down the Shift key and clicking the name of additional files.
4. Select **Restore To** from the bottom bar to restore the file to the location from where it has been saved from during backup.

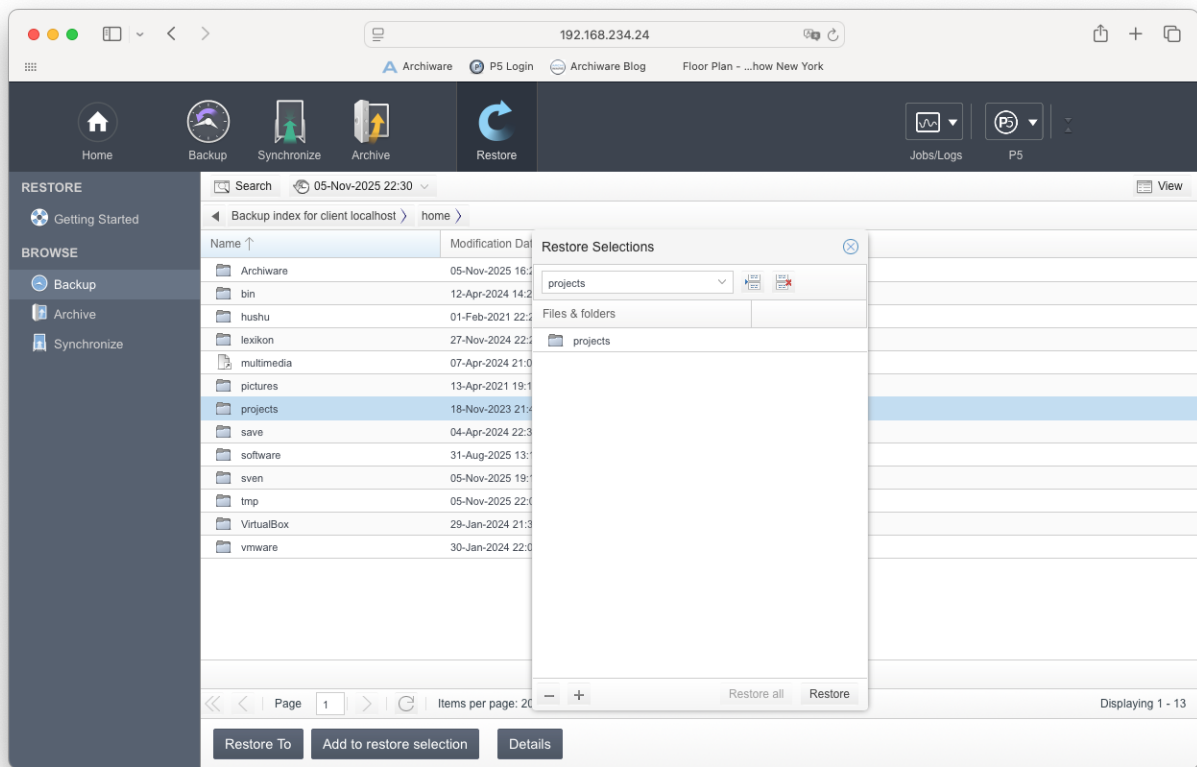
When restoring files from multiple position in the index, it may be better to first select all files and to restore them together. That way, P5 can optimize the tape access, and the restore time is sped up. To do so, Select **Add To Restore Selection** and restore the selected files later together from within the **Restore selection**.

8.2 The Restore Selection

When browsing the index tree to select files to restore, the restore selection can be opened by simply adding a file via the context menu. The restore selection allows to select multiple files and/or folders to restore with a single restore job.

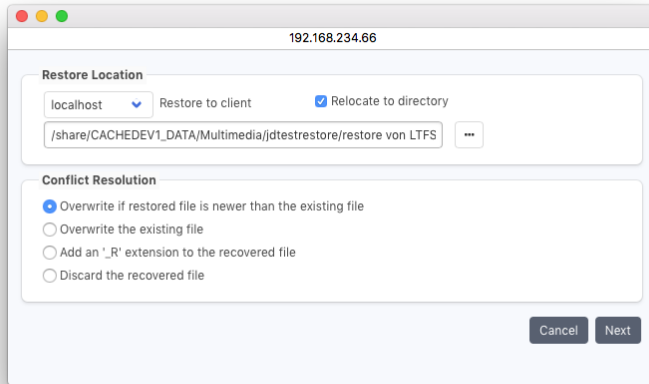
To restore the selected files, click **Restore To** in the Restore Selection window. Usually this procedure is faster than restoring files one by one as tape loading and positioning can be done in a more optimal way.

The restore selection remains active in the background until the user logs out from the P5 GUI. There is a separate Restore Selection for each index. In order to restore files from multiple clients, P5 uses a separate Restore Selection for each client, each selection must be restored with a separate restore job.



8.3 The Restore To Button

Use **Restore To** to restore with detailed control of the target and the restore options. The pop-up window allows the user to define the details of the restore process.



Restore Location:

This section defines the place to restore to.

Restore to client:

Select the P5 client the data shall be written to.

Relocate to directory:

Check this box to change the directory to restore to. Sub-directories selected for restore are generated automatically below the selected folder.

The ability to select a client or directory depends on the user's permissions defined in the user preferences (see section 6.14) and the defaults defined for the respective client (see section 6.11).

Conflict Resolution:

Files with identical names might exist in the folders where the restore should be done. You can choose how to resolve such a conflict:

Overwrite if restored file is newer than the existing file:

P5 will overwrite the existing file only if it is older than the one that it is restoring.

Overwrite the existing file:

P5 will unconditionally restore the file, regardless if a file with the same name already exists.

Add an '_R' extension to the recovered file:

If a file with the same name exists, the file which is restored will get an '_R' appended to its name.

For example, picture.tif will be renamed to picture_R.tif.

Discard the recovered file:

If a file with the same name already exists, the file will not be restored.

Before the actual restore takes place, you are informed about the volumes required to complete the restore operation and their availability. If all the required volumes can be found in a library that is online, P5 will manage volume changes autonomously, otherwise missing volumes will be requested and the restore process paused until the tape becomes available. Click **Start** to proceed.

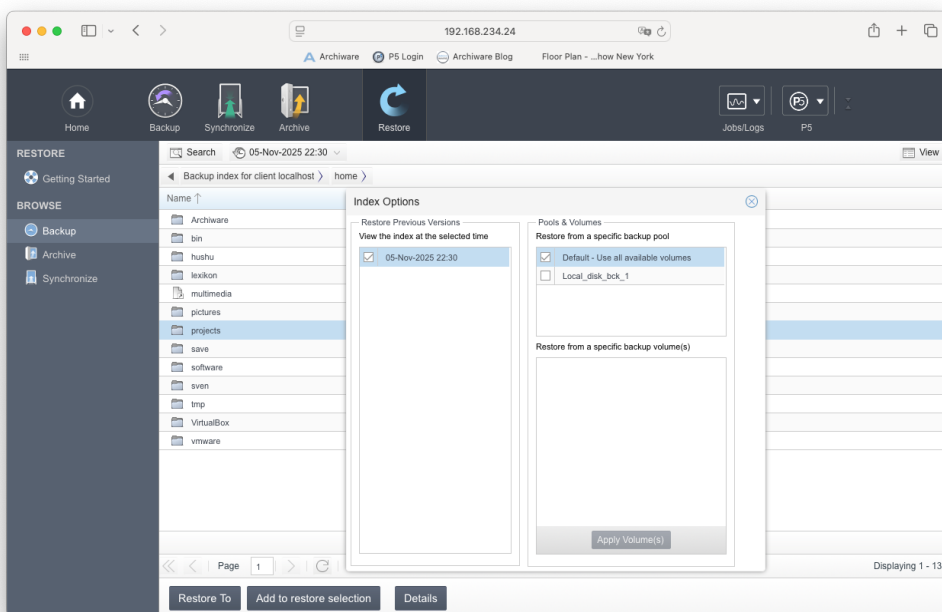
To monitor the restore process, start the **Job Monitor** from the **Jobs/Logs** section of the top menu, see

section 6.10.1, The Job Monitor.

8.4 The Index Options

P5 maintains all saved files in snapshots. Each snapshot represents the saved host as it was seen at the given time. Each running backup job creates a new snapshot by adopting all unchanged files and adding those files newly saved. By default, P5 always shows the last recent snapshot to restore from.

In order to restore files from an earlier state, simply select an earlier snapshot from the Snapshot dropdown list in the window heading (by default, Snapshot dropdown list shows the latest snapshot creation date). The item **More Options** opens a sub-window with advanced options for selecting the snapshot, backup pool and specific backup volumes.



The snapshot timeline allows to select a specific snapshot, containing only files up to that point in time.


The sub-window's **Pool Selection** area allows to select a Media pool to restore the files from. For example in case the backup planning saves the same files and folders multiple times to tape, the pool selection may be used to preset the tapes that are to be used for restore. That may also reduce the number of available snapshots presented in the restore area.

The **Volume Selection** area allows to select a Volume pool to restore the files from. That may also reduce the number of available snapshots presented in the restore area.

8.5 File Versions

When selecting a single file for restore (by double clicking corresponding row in restore selection), Backup Information window lists all the versions of that file that have been saved. Versions are shown over all snapshots that exist. The preselection of a media pool in the Index Options may reduce the number of versions shown.

This functionality is intended to enable quick selection of a specific version of a single file. In order to select a previous state of entire folder tree, please select the appropriate snapshot.

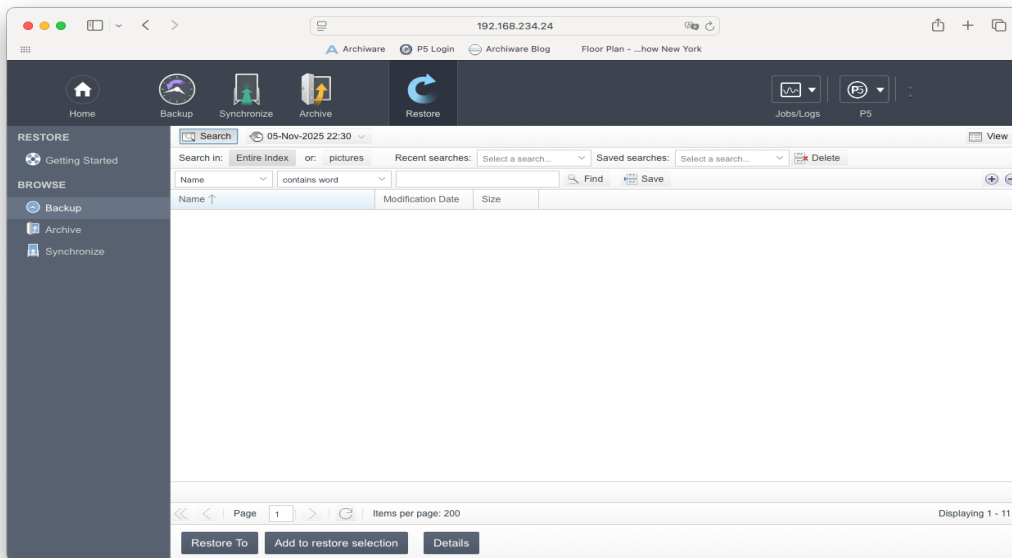


This info window allows checking whether a saved file still exists on tape, and if so on which. The browser window may also show files that were saved but the last tape containing it has been removed or recycled.

8.6 The Search Button

To simplify the search through the backup index and to provide quick access to the backed up data, P5 gives the user a comprehensive search tool. This tool enables searching through the current content of the backup. If you are looking for an older backup state, select Index Options instead.

To commence a search, navigate via Restore from the top level of the P5 browser into the backup index and into the folder that you wish to search through. Then click Search from the windows heading.



Below the Search button, select where to search. By default, the entire index is searched. Click the current folder name in the window headline to search only in that folder sub-tree.



Note that restricting the index search to a sub-tree is slow, it may take even longer than searching the complete index due to technical limitations. The same speed reduction is the case when a login area is used that restrict the view to a part of the index only.

When running a slow search, please wait until the result is presented, the index search is active in background (and will remain active until it completes).

As an optimization, the first search operation restricted to a folder tree will create a helper table in background so that the next subsequent search over the some folder subset will be sped up.

Next, select what to search by constructing a search criteria.

Constructing search criteria:

Search criteria consists of one or multiple search expressions. Each expression is defined by *Attribute*, *Operator* and *Value*. Multiple expressions can be added to search criteria by pressing the '+' button at the right side of the expression. In case of multiple expressions, search will show only the results that meet ALL search expressions. The following table shows possible search expression combinations.

Attribute	Operator	Value	Description
Name	contains word	any string	Matches file names that contain a word. One or multiple words are allowed, file names can also be entered complete. This search is complex, see the description below.
	contains string	any string	Match any part of the file name with the entered value. This search operation is slow on big indexes.
Type	is	File	Match all files.
	is	Folder	Match all folders.
Modification Date	is exactly	any date	Match entries modified exactly on a given date.
	is before	any date	Match entries modified before a given date.
	is after	any date	Match entries modified after a given date.
	within last	X Days/Weeks/Months	Match entries with modification date within last X days (or weeks or months), where X is any number.
Size	equals	X Byte/KB/MB/GB	Match files with size equal to the given number of Bytes (KB/MB/GB).
	is lesser than	X Byte/KB/MB/GB	Match files smaller than the given number of Bytes (KB/MB/GB).
	is greater than	X Byte/KB/MB/GB	Match files bigger than the given number of Bytes (KB/MB/GB).

Click Find to start the search.

Files and folders can be selected for restore directly, same as in the index browser. Double clicking a found folder takes you directly to that folder, allowing to restore files or folders from there or to start a new search in that sub folder.

Details on the "Name contains word" search operation

This is the recommended search operation to search file names. P5 uses an index based search for words. To achieve that, in addition to the file names, name fragments – *words* – are added to the index in addition to the complete file names. Technically, a word is any fragment that is separated with blanks or punctuation characters from the rest of the file name. Words are added to the index if they have three or more characters.

Examples

A filename	Is separated into words
a_bugs_life.mp4	bugs life mp4
My Project 12/08/2015	Project 2015
NotReallySingleWords.12	NotReallySingleWords

When searching with "name contains word", also the search expression is split into word and these words are searched, where only the first of the entered words uses an index based search. That implies the search is faster when the first searched word limits the search for only few files.

Examples

The search expression	Will find
a_bugs_life	a_bugs_live.mp4 bugs bunnys life 1990.wav life of antony and the bugs.txt
bug life	a_bugs_live.mp4 bugs bunnys life 1990.wav life of antony and the bugs.txt
NotReally	NotReallyAll.2015 NotReallySingleWords.12

One can optimize the search by putting a longer, more significant word as first word in the list. In the above examples, the first expression, the name a_bugs_life will find the file, but the second search finds the same and is significantly faster. However, when entering a complete file name, that file name will always be among the hits.

The search tends to deliver more hits than expected, but by adding further words the result may be reduced. Also, search expressions can be combined. When doing so, the search by "name contains word" should be part of the combined search in order to speed up the search.

Saving search criteria

After constructing the search criteria, you can save it for future use by pressing Save button. Enter a name for your search criteria in the popup window and press OK button to finish saving.

Saved search criteria can be accessed through Saved searches drop down menu. Pressing the Delete button will delete selected search criteria.

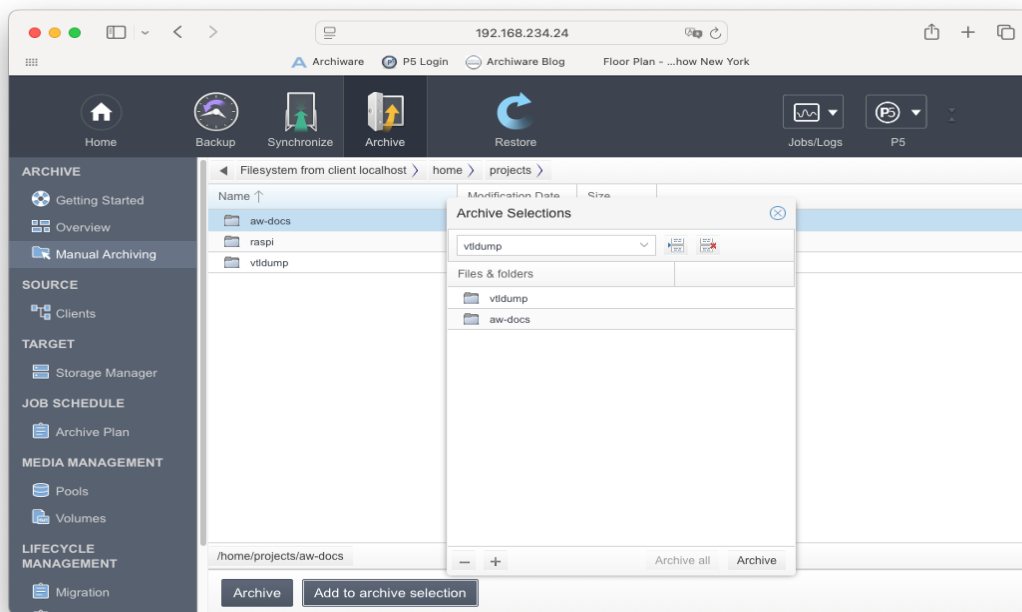
9 Archiving and Restoring from Archive

P5 supports both automatic and manual archiving. An archive plan must exist before you can start archiving. The plan defines parameters to simplify and speed up the submission of a file or the collection of files to be archived. Among other parameters, the archive plan defines the target pool of volumes to store the data to and the index to hold the list of files that have been archived. See section 6.5, Archive Plans.

9.1 Archiving Files

Follow these steps:

1. Log into P5 using a web browser.
See section 4.1, Starting the Browser.
2. Select Archive from the top menu, then select Manual Archiving from the left sidebar.
3. The main window shows the available clients and their file systems.



You can create log-in areas into the file system as need be. They enable access to a specific folder structure within the file system. See section 6.13, Index Access Control (Login Areas)

4. Navigate to the folder that you wish to archive. Select the folder by clicking its name. Click Archive in the bottom bar.
5. Select an archive plan from the list. The archive plan defines the media pool and the index to use for archiving.

You can start archiving immediately. If not checked, archiving will be started as configured in the chosen

plan.

Specify a location within the index:

The file can be placed in the index either at the same location as in the file system or relocated to a new location in the index. If you define a new location, the path to the folder where the file or folder you are archiving is located will be replaced in the index by the newly chosen path.

You may define additional fields (meta-data) in each of the archive databases (indexes) to describe and systematize your data in more detail. This makes the task of finding and selecting the data much easier later on. The entries you make here will be associated with each of the files and directories archived in this run.

Click Start to submit the archiving job.

To monitor the restore process, start the Job Monitor from the Jobs/Logs section of the top menu, see section 6.10.1, The Job Monitor.

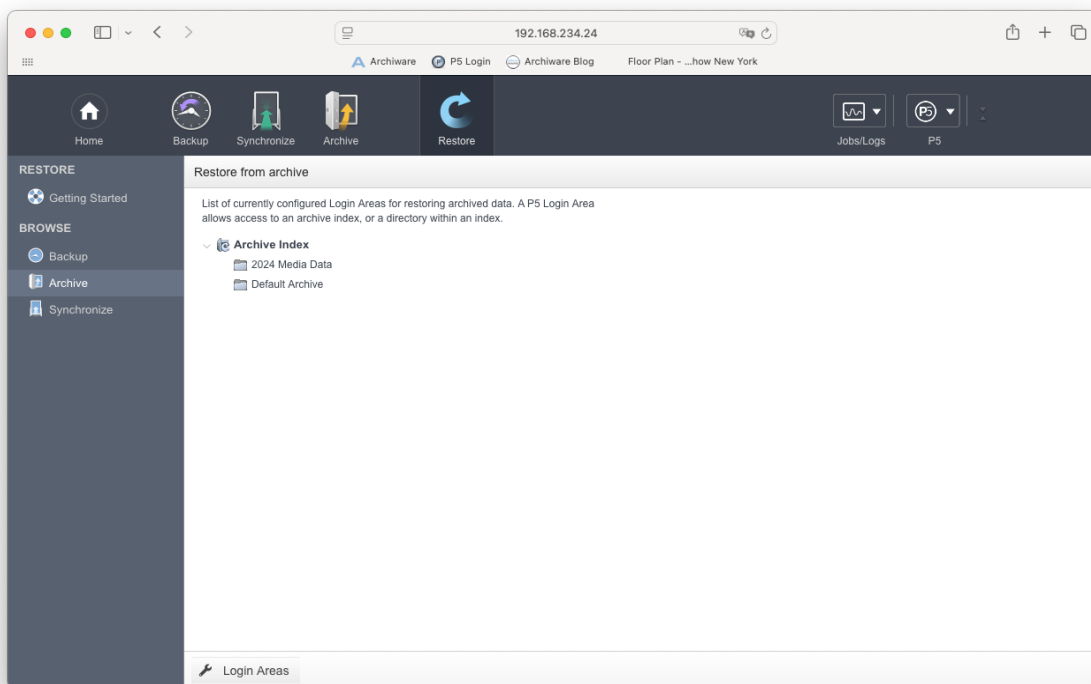
As an alternative to directly Archiving files, it is possible to the files to an Archive Selection and archive the selected files together in one archive job. That is usually faster than starting multiple archive jobs as it avoids the job specific overhead like loading and positioning the tape.

Multiple Archive Selections can be created manually, these are named by the first contained file and will remain active in background until either the user logs out or an archive job for that selection is started.

9.2 Restoring from Archives

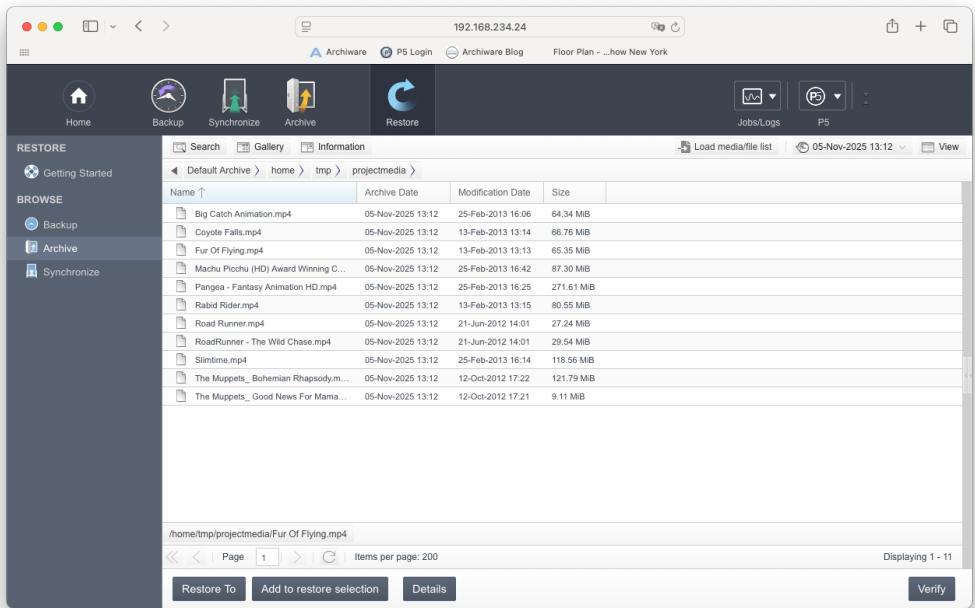
Follow these steps:

1. Log into P5 using a web browser.
See section 4.1, Starting the Browser.
2. Navigate into the archive index where the data that you wish to restore is located. Select Restore in the top menu, then choose Restore from Archive from the left sidebar.
The main screen will show the archive indexes available to restore from.



Note that in order for user defined archive indexes are made available here, a login area is required that defines who is permitted to access the index and whether the indexes appear grouped. The button *Login Areas* at the bottom of the screen directly links to the login area definitions. See section 6.13, Index Access Control (Login Areas)

After selecting an archive index, the browser view shows the files and folders saved.



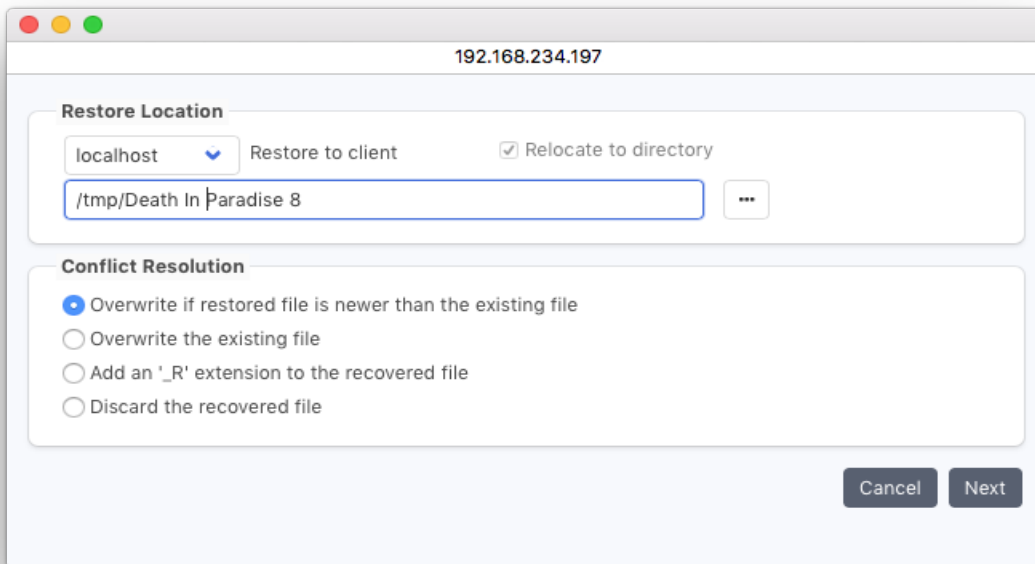
Navigate to the folder where the files and/or folders that you wish to restore are located. Select one of the files by clicking its name. Then select a second file by holding down the Shift key and clicking its name. You can select multiple elements this way.

Select **Restore To** from the bottom bar to restore the file.

To monitor the restore process, start the **Job Monitor** from the **Jobs/Logs** section of the top menu, see section 6.10.1, **The Job Monitor**.

9.2.1 The Restore-To Button

Use **Restore To** in order to restore with detailed control of the target and the restore options. The pop-up window allows to define the details of the restore process.



Restore Location

This section defines the place to restore to.

Restore to client:

Select the P5 client that data will be written to.

Relocate to directory:

This box is checked by default. P5 then restores to the selected target folder. Sub-directories selected for restore are generated automatically below the selected folder. When the relocation is selected, the used paths are the index paths as shows in the index browser.

When deselecting the relocate function, P5 restores to the original paths, these paths are not the index paths but the paths where the files have been archived from.

Note that the relocation affects all files and folders that are restored, even when the files were selected from different locations in the index, for instance via restore selection window.

Restoring to the original location is possible only if there are no restrictions defined for the restore path, i.e. the group preferences for the logged-in user must not have a path set, in order to use this function. See also section 6.15, Group Preferences.

The ability to select a client or directory depends on the user's permissions defined in the user preferences (see section 6.14) and the defaults defined for the respective client (see section 6.11).

Conflict Resolution:

Same named files might exist in the folders where the restore should be done. You can choose how to resolve such a conflict:

Overwrite if restored file is newer than the existing file:

P5 will overwrite the existing file only if it is older than the one that it is restoring.

Overwrite the existing file:

P5 will unconditionally restore the file, regardless if a file with the same name already exists.

Add an '_R' extension to the recovered file:

If a file with the same name exists, the file which is restored will get an '_R' appended to its name.

For example, picture.tif will be renamed to picture_R.tif.

Discard the recovered file:

If a file with the same name already exists, the file will not be restored.

Before the actual restore takes place, you are informed about the volumes required to complete the restore operation and their availability. If all the required volumes can be found in a library that is on line, P5 will manage volume changes autonomously, otherwise missing volumes will be requested and the restore process paused until the tape becomes available. Click **Start** to proceed.

To monitor the restore process, start the **Job Monitor** from the **Jobs/Logs** section of the top menu, see section 6.10.1, **The Job Monitor**.

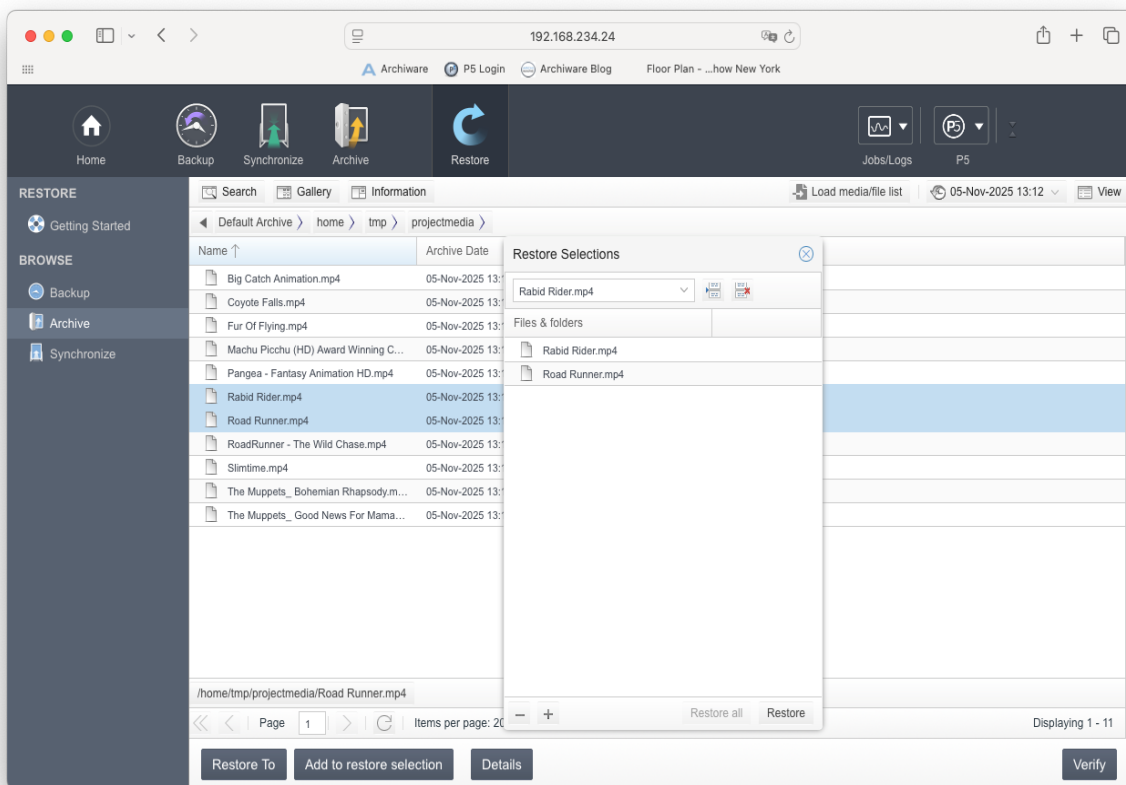
9.2.2 The Restore Selection

When browsing the index tree to select files to restore, the restore selection can be opened by simply adding a file via the context menu. The restore selection allows to select multiple files and or folders to restore in a single restore job.

To restore the selected files, click **Restore To** in the Restore Selection window. Usually that procedure is faster than restoring files one by one as tape loading and positioning can be done in a more optimal way.

There is a separate Restore Selection for each index automatically. Different Restore Selections for one index can be created manually, these are named by the first contained file. All restore selection remains active in the background until the user logs out from the P5 GUI or the restore operation is started.

In order to restore files from multiple clients, P5 uses a separate Restore Selection for each client. Each selection must be restored with a separate restore job.



9.2.3 File Versions

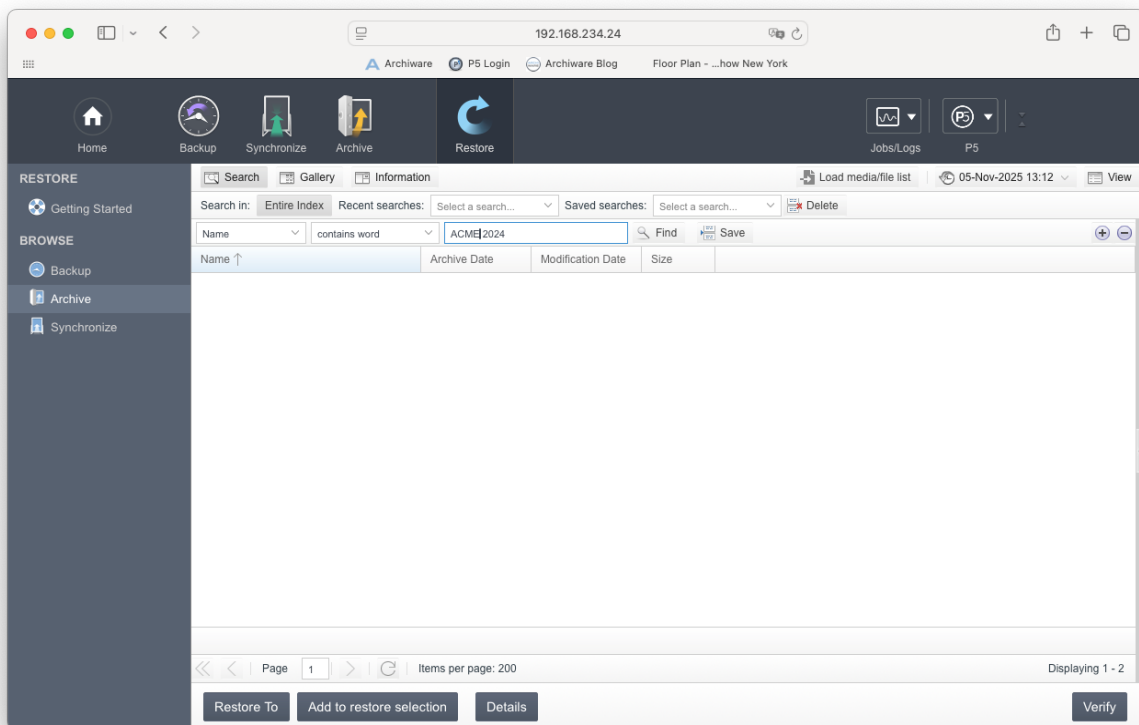
When selecting a single file for restore (by double clicking the corresponding row in the restore selection), the Archive Information window lists all the versions of that file that have been saved. Versions are shown over all snapshots that exist. The preselection of a media pool in the Index Options may reduce the number of versions shown.

This functionality is intended to enable a quick selection of a specific version of a single file. In order to select a previous state of an entire folder tree, please select the appropriate snapshot.

9.2.4 The Search Button

To simplify searching through an archive index and to provide quick access to the archived data, P5 gives the user a comprehensive search tool. This tool enables searching through the current content of the archive.

To commence a search, navigate via Restore from the top level of the P5 browser into the archive index that you desire to search through. Then click Search from the window headline, which will open the following window for you:



Below the Search button, select where to search. By default, the entire index is searched. Click the current folder name in the window headline to search only in that folder sub-tree.



Note that restricting the index search to a sub-tree is slow, it may take even longer than searching the complete index due to technical limitations. The same speed reduction is the case when a login area is used that restrict the view to a part of the index only.

When running a slow search, please wait until the result is presented, the index search is active in background (and will remain active until it completes).

As an optimization, the first search operation restricted to a folder tree will create a helper table in background so that the next subsequent search over the some folder subset will be sped up.

Next, select what to search by constructing a search criteria.

Constructing search criteria:

Search criteria consist of one or multiple search expressions. Each expression is defined by *Attribute*, *Operator* and *Value*. Multiple expressions can be added to search criteria by pressing the '+' button at the right side of the expression. In case of multiple expressions, search will show only the results that meet ALL the search expressions. The following table shows possible search expression combinations:

Attribute	Operator	Value	Description
Name	contains word	any string	Matches file names that contain a word. One or multiple words are allowed. Entire file names can also be entered. This search is complex, see the description below.
	contains string	any string	Match any part of the file name with the entered value. This search operation is slow on big indexes.
Type	is	File	Match all files.
	is	Folder	Match all folders.
Archive Date	is exactly	any date	Match entries archived exactly on a given date.
	is before	any date	Match entries archived before a given date.
	is after	any date	Match entries archived after a given date.
	within last	X Days/Weeks/Months	Match entries with archival date within last X days (or weeks or months), where X is any number.

<i>Modification Date</i>	<i>is exactly</i>	any date	<i>Match entries modified exactly on a given date.</i>
	<i>is before</i>	any date	<i>Match entries modified before a given date.</i>
	<i>is after</i>	any date	<i>Match entries modified after a given date.</i>
	<i>within last</i>	<i>X Days/Weeks/Months</i>	<i>Match entries with modification date within last X days (or weeks or months), where X is any number.</i>
<i>Size</i>	<i>equals</i>	<i>X Byte/KB/MB/GB</i>	<i>Match files with size equal to the given number of Bytes (KB/MB/GB).</i>
	<i>is lesser than</i>	<i>X Byte/KB/MB/GB</i>	<i>Match files smaller than the given number of Bytes (KB/MB/GB).</i>
	<i>is greater than</i>	<i>X Byte/KB/MB/GB</i>	<i>Match files bigger than the given number of Bytes (KB/MB/GB).</i>
<i>Description, further metadata fields</i>	<i>starts with</i>	any string	Match the entries whose description begins with the entered value.
	<i>contains string</i>	any string	Match any part of the description with the entered value.

Click **F**ind to start the search.

Files and folders can be selected for restore directly from the search result, same as in the index browser. Double clicking a found folder takes you directly to that folder, allowing to restore files or folders from there or to start a new search in that sub folder.

Details on the "Name contains word" search operation

This is the recommended search operation to search file names. P5 uses an index based search for words. To achieve that, name fragments – *words* – are added to the index in addition to the complete file names. Technically, a word is any fragment that is separated with blanks or punctuation characters from the rest of the file name. Words are added to the index if they have three or more characters.

Examples

A filename	Is separated into words
<code>a_bugs_life.mp4</code>	<code>bugs life mp4</code>
<code>My Project 12/08/2015</code>	<code>Project 2015</code>
<code>NotReallySingleWords.12</code>	<code>NotReallySingleWords</code>

When searching with "name contains word", the search expression is split into words too. These words are searched, where only the first of the entered words uses an index based search. That implies the search is faster when the first searched word limits the search to only a few files.

Examples

The search expression	Will find
<code>a_bugs_life</code>	<code>a_bugs_live.mp4</code>

The search expression	Will find
bug life	bugs bunnys life 1990.wav life of antony and the bugs.txt a_bugs_live.mp4 bugs bunnys life 1990.wav life of antony and the bugs.txt
NotReally	NotReallyAll.2015 NotReallySingleWords.12

One can optimize the search by putting a longer, more significant word as first word in the list. In the above examples, the first expression, the name `a_bugs_live` will find the file, but the second search finds the same and is significantly faster. However, when entering a complete file name, that file name will always be among the hits.

The search tends to deliver more hits than expected, but by adding further words the result may be reduced. Also, search expressions can be combined. When doing so, the search by "name contains word" should be part of the combined search in order to speed up the search.

Saving search criteria:

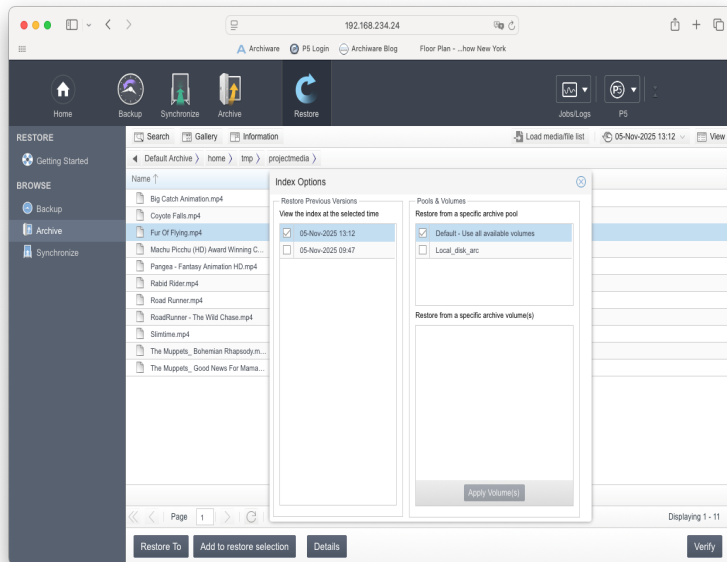
After constructing the search criteria, you can save it for future use by pressing the `Save` button. Enter a name for your search criteria in the popup window and press the `OK` button to finish saving.

Saved search criteria can be accessed through the `Saved searches` drop down menu. Pressing the `Delete` button will delete selected search criteria.

9.2.5 Index Options

Each run of an archive job creates a new snapshot. By default, P5 always shows the most recent snapshot to restore from, including all older files.

In order to restore files from an earlier state, simply open the Index section in the window heading. This opens the index option sub window:



The snapshot timeline allows to select a specific snapshot, containing only files up to that point in time.

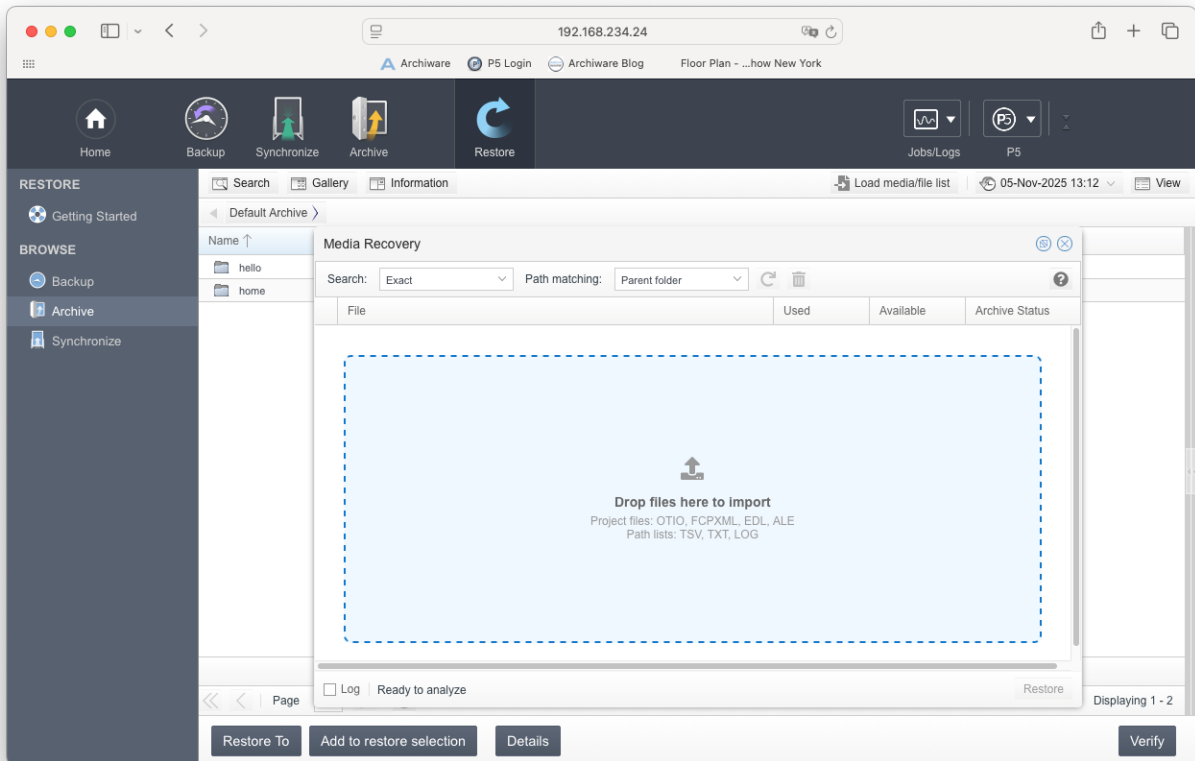
The Pool Selection area allows to select a Media pool to restore the files from. That may also reduce the number of available snapshots presented in the restore area.

The Volume Selection area allows to select a Volume pool to restore the files from. That may also reduce the number of available snapshots presented in the restore area.

9.2.6 Restore from a File List

Starting in version 8, P5 supports restoring files referenced by a file list. To do so, click **Load Media File List** in the heading bar of the restore section.

P5 then offers loading a file list:



Drag a supported media file list into the marked area in the media recovery window. P5 will show the found files with a button to restore them.

The following file formats are supported:

TSV, TXT, LOG Simple text files with one file per line

- OTIO OpenTimelineIO-Exchange file
- FCPXML Final Cut Pro project file
- EDL Edit decision list
- ALE Avid log exchange

9.2.7 Editing the Archive Index

Right clicking on an entry in the Archive Index brings up a context menu with additional options:

Add to restore selection	adds the entry to the existing restore selection, same as clicking on the "Add to restore selection" button
Restore to	restores the selected entry, same as clicking on the "Restore to" button
Description	brings up the description window for the selected Archive entry. The description window shows available versions, a preview for the supported document types and metadata information, where you can change the item's description and UUID.
Element...Rename	Rename the selected folder
Element...Move	Move the selected folder / subtree
Element...New Folder	Create a new subfolder
Element...Delete	Delete the selected folder / subtree



Note that access to the index editing functions is disabled by default. It has to be enabled for a user in the user preferences before it can be used, see 6.14 User Preferences.

When deleting a folder in the index, it is removed from the index. In case the element is saved in a Container storage, it is also removed on the storage. Files on discrete media like tapes cannot be deleted from the storage.

10 Operating Procedures

This section describes the tasks necessary for an optimal and smooth day-to-day operation of the system. This includes:

- checking the backup.
- managing the backup devices and volumes.

10.1 Labeling Media

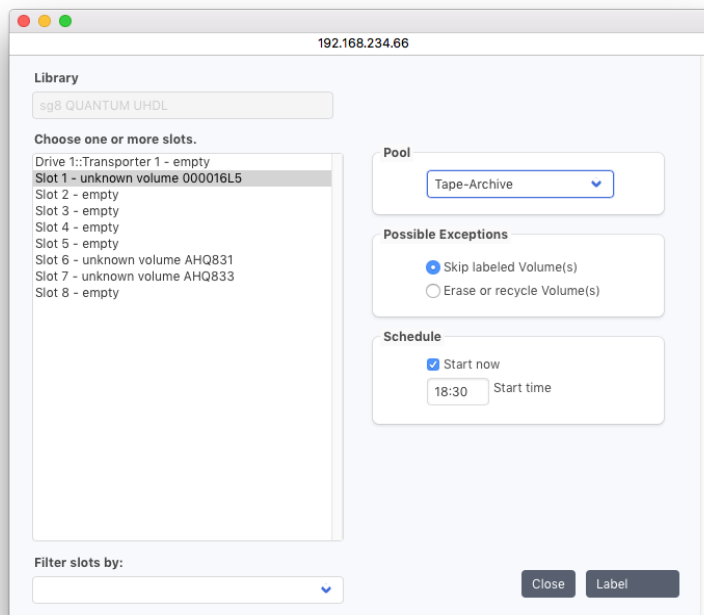
Backup is usually done on mass media such as tapes. P5 manages the media as volumes. Before a media can be used in P5, it must be labeled. Labeling physically writes a label on the media and adds a volume into the configuration database. The label and the volume carry diverse configuration information. A name and a sequence number uniquely identify each volume, the name being that of the pool it belongs to.

Follow these steps to label volumes:

1. Log into P5 using a web browser.
See section 4.1, Starting the Browser.
2. Depending on the hardware available to you, a single drive or a media changer, you will submit the labeling job from the corresponding management area.

Select the Archive or Backup module from the top menu and then from the left sidebar: Storage Manager if you have a media changer, or Standalone Tape Drive if you have a single tape drive.

3. From the bottom menu, choose **Label**



Slot Range:

This defines the first to last slot, for example a range of media in the jukebox that you wish to label.

Select a Pool:

Choose the pool to which the volumes will be assigned.

The check box `Skip labeled Volume(s)` is selected by default to avoid erasing media containing data. Such volumes are skipped by default. Chose `Erase or recycle Volume(s)` instead, to erase any data already on the media you wish to label.

10.2 Mounting Volumes

Before P5 can start to read or write a volume, it has to be mounted. Mounting reserves the correct volume(s) for backup, archive or restore. They request certain volumes or pools from the media manager. Where a media library is available, the media manager will automatically mount the requested volume itself. If there is no library available for the operation or if the location of the media is not known to the system, the operator needs to mount the volume manually.


If a volume has been stored outside of the system (the library), it has to be put into the system (the library) again. Then an inventory of the library has to be run to make the volume position and availability known to P5 again. See section 6.2.6.2, Take Inventory...

When the volume is not needed any more, it can be 'unmounted' and removed from the device. With a library, this happens automatically.

10.2.1 Mounting or Unmounting a Volume

Follow these steps:

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select Archive or Backup from the top menu
3. Select the Standalone drive from the left sidebar
4. Select the device and click Mount/Unmount from the bottom bar



Click Mount to mount the volume. Click Unmount to unmount the volume.

10.3 Replacing Volumes in a Library

P5 keeps a list of labeled volumes in the configuration database. Volumes located in a library have the corresponding jukebox and position attributes set. If you remove volumes from the library or introduce old volumes into it, the configuration database has to be updated by running an inventory. The volumes will either be mounted into drives and their labels read, or where available, the bar code will be read and compared to the configuration database. The bar code inventory is considerably faster.

P5 automatically runs a bar code inventory of a library when it detects discrepancies between its internally stored inventory and the actual inventory of the changer, when the volumes it detects differ from the ones it expected to find.

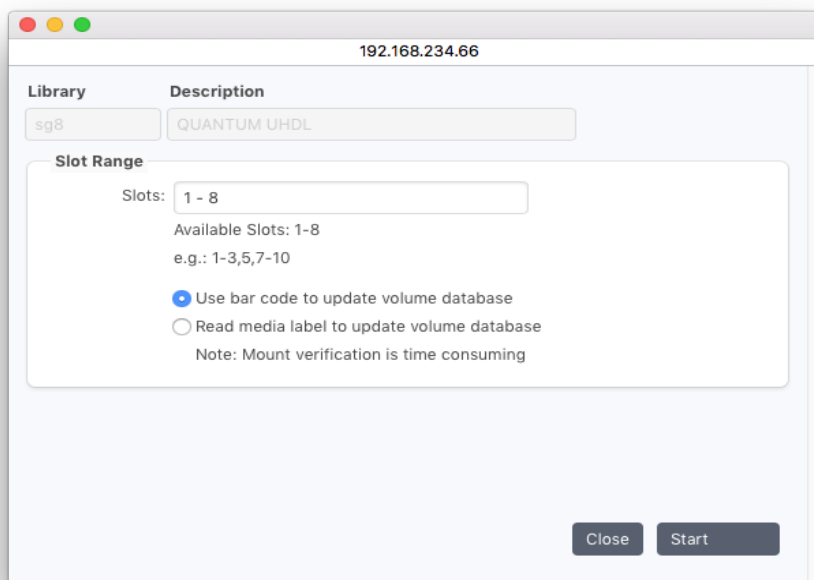


A bar code inventory can only be done if the library is equipped with a bar code reader.

10.3.1 Jukebox Inventory

Follow these steps:

1. Log into P5 using a web browser as an admin user. See section 4.1, Starting the Browser.
2. Select Archive or Backup from the top menu.
3. Select the Storage Manager from the left sidebar.
4. Select the jukebox and click Inventory from the base menu.



Slot Range:

This defines the first and the last slot, for example a range of media in the library that you wish to run through the inventory.

Use Bar code to update volume database.

This tells P5 to compare the bar codes with the configuration database to determine which volumes are available.



The bar code inventory can only be done if the library is equipped with a bar code reader.

Read media label to update volume database.

This tells P5 to compare the label written on each volume with the configuration database to determine which volumes are available.



Since each volume has to be mounted, this can take some time.

10.4 Handling Media Pools

In P5, media pools are used to group sets of tapes and are addressed in every backup or archive job. P5 does not address single tapes when writing. To assign tapes to media pools, the tapes must be labeled, see section 10.1 Labeling Media above for how to label a tape.

Media pools define an organizational unit that can be used in several ways to fine-tune the behavior of P5 when writing:

- Different pools are always used to write backup and archive data onto different tapes. It is also possible to differentiate the data into different tape sets, for instance to write different archive to different pools, or to define different backup sets that are to be completed with backups.
- Different pools are also used to write to tape and to disk (a "virtual tape"). In addition, it is possible to use media pools to differentiate which device is used for writing. For instance, the tapes in one tape library can be labeled for one media pool while the tapes in another library belong to another pool. That way, the media pool can be used to address the library.
- The same mechanism can be used to address the magazines within one tape library. P5 keeps a strict assignment of the tape and the slot where it was taken from, so P5 does not mix up tapes among different magazines.
- Media pools can be limited to address specific drives. In that way, it is possible to separate for instance an LTO-6 pool and an LTO-5 pool within the same changer so that LTO-6 tapes are only used in an LTO-6 drive while LTO-5 tapes can be bound to the LTO-5 drives (the latter is technically not required).
- The parallelization options and the cloning option in the media pool allow to define different schemes for writing to tape. For instance an archive job may write to two (equal) drives with cloning, thus creating two identical copies of each tape, while a backup pool may address the same two drives to write different data in parallel.
- The parallelization options in the media pools may be limited to a predefined number of jobs that may write in parallel to one drive. In addition, it is easy to enforce a rule where only one pool's tape can be in one drive at a time. This enables restricting the total system load and regulating the data flow within the system.

Note that empty tapes must be labeled for a specific media pool before they can be written. In general it makes sense to provide a sufficient number of tapes before starting to write.

10.5 Handling Network Shares on Windows

It is possible on Windows to use network shares, i.e. drives that are visible from other computers. In some situations this is the only way to access a remote computer, for instance if P5 is not available as a client for that host.

On Windows, P5 runs as a system service. That implies that it runs in the context of the local system user account. Since on Windows network shares are mounted per user, it is thus impossible to access the network shares that a Windows user connected to, for instance via Windows explorer. Instead, the network share must be mounted in the context of the local system account. That can be achieved with the Windows / Sysinternals utility `psexec`. This tool can be downloaded here:

<https://docs.microsoft.com/en-us/sysinternals/downloads/psexec>

10.5.1 Mount a Network Share and Assign it to a Drive Letter

To mount a network share and assign it to a drive letter:

1. Open a CMD Window as Administrator
2. Enter the `net use` command to mount the network share. Use `psexec` to do so in the context of the local system account:
`psexec -s -i net use X: \\192.168.234.31\Data /persistent:yes`
This command would mount the network share and assign it to drive letter X. The parameter `-i` allows to specify user and password for this mount interactively.

P5 will then show the network drive letter same as local drive letters and allow to read from or write to that network share.

10.5.2 Mount a Network Share and Use the URL Directly

Starting with version 7.1, you can mount a network share without assigning it to a drive letter:

3. Open a CMD Window as Administrator
4. Enter the `net use` command to mount the network share. Use `psexec` to do so in the context of the local system account:
`psexec -s -i net use \\192.168.234.31\Data /persistent:yes`
This command would mount the network share without assigning it to a drive letter. The parameter `-i` allows to specify user and password for this mount interactively.

P5 will then allow to access the network share for reading or writing. Note that the share must be used in a P5 specific syntax. Please specify `/@server/share/path` instead of `\\server\share\path`.

However, the share is not listed in the GUI for browsing. In order to browse the network share, it is required to create a login area for this share. Enter the share, for instance `/@192.168.234.31/Data` as the path in the login area. To create a login area, see section 6.13, Index Access Control (Login Areas).

11 Customizing

This section describes the customization possibilities in P5, which allow each user to quickly and easily complete his P5 related tasks.

11.1 Users and Groups

Users

Starting with version 6, P5 maintains user accounts itself and does no longer use the operating system users. So for each user who shall be able to access P5, a user account must be created. User accounts are maintained in the User Preferences, see section 6.14. The user preferences define the role of a user, administrator, operator or normal users, the functions a user can reach and the group(s) a user belongs to.

User Groups

User groups define the file system access permissions during restore. For each group, one or multiple paths can be defined where data can be restored to. Groups are described in section 6.15, Group Preferences.

11.2 Index and File System Access

In addition, the user groups can be associated with individual log-in areas, which allows the given group access to a defined log-in area and may disallow access to everybody else. If no group is associated with a log-in area, this area can then be accessed by everyone.

This mechanism allows to define access permissions to indexes and file systems by group, but also to grant access to multiple log-in areas while allowing restrictions.

In order to restrict access to log-in areas to defined groups, add these groups in the log-in area record, see section 6.13, Index Access Control (Login Areas). Log-in areas with no group entered are not restricted.

P5 organizes and maintains databases in a tree structure comparable to a file system on the disk. Access to the databases is enabled through the log-in areas. This access control affects:

- the indexes visible to a user in the P5 Restore section of the top menu, and
- the file systems accessible by a user when manually archiving files.

Customizing the Home Log-In Area

P5 allows you to individually direct each user to a predefined place within the application after the user signs in. For end users who seldom need to employ P5 and always tend to perform only a limited variety of tasks, this simplifies finding the correct functionality and navigating through the application.

To configure user preferences, see section 6.14, User Preferences.

To restrict each individual user to a particular log-in area, set the corresponding log-in area as `Login Behavior` in that user's preferences.



In order to grant log-in area access to a non administrator user, make sure the log-in area is not restricted to administrators only, which is set by default.

11.3 Customizing Reporting

Following each job that P5 executes (for example a backup or an archive job), a protocol of the run is saved. This protocol carries all the information relevant to the execution and completion of the job and especially of exceptional conditions such as a damaged media.

The protocols can be retrieved by selecting **Jobs/Logs** from the top menu. See more about logs in section 6.10.2, *The Log History*.

In order for the P5 administrator to maintain an overview of the jobs and systems and to be able to react promptly to problems, two mechanisms are implemented to send mails and/or notifications to selected users. Notifications for all automatically scheduled jobs are being sent to all users who chose to be notified of such events. Jobs that are started manually by a logged in users are reported by email to that user exclusively.

To configure the email delivery, navigate to the user preferences configuration area and define for each user the receiver mail address and which notifications should generate emails. The setup is described in detail in section 6.14, *User Preferences*.

To configure which mail server to use for mail sending and how to connect to the mail server, open the localhost client configuration. See section 6.11.1, *Creating a Client* for how to find the localhost client.

The screenshot shows a configuration window titled "192.168.234.197" with two tabs: "General Setup" and "Additional Options". The "Additional Options" tab is active. It contains a table with columns "Client", "Description", "Identifier", and "Status". The "localhost" client is listed with "localhost" in the Description and "Enabled" in the Status. Below the table is the "Data Encryption" section, where "Encrypt data:" is set to "No". An "Important Note" states that all data is encrypted and provides the key file path: `/usr/local/aw/config/localhost.key`. The "E-Mail Configuration" section includes fields for "SMTP mail server" (localhost), "Mail sender" (root@localhost), "Port" (25 selected), "User name", "Password", and "Test email address". There are "Close" and "Apply" buttons at the bottom right.

Enter as *SMTP mail server* the name or IP address of your mail server. In case a name is entered, please make

sure the name can be resolved to an IP address. Enter as Mail sender the E-mail account used as sender address in the email.

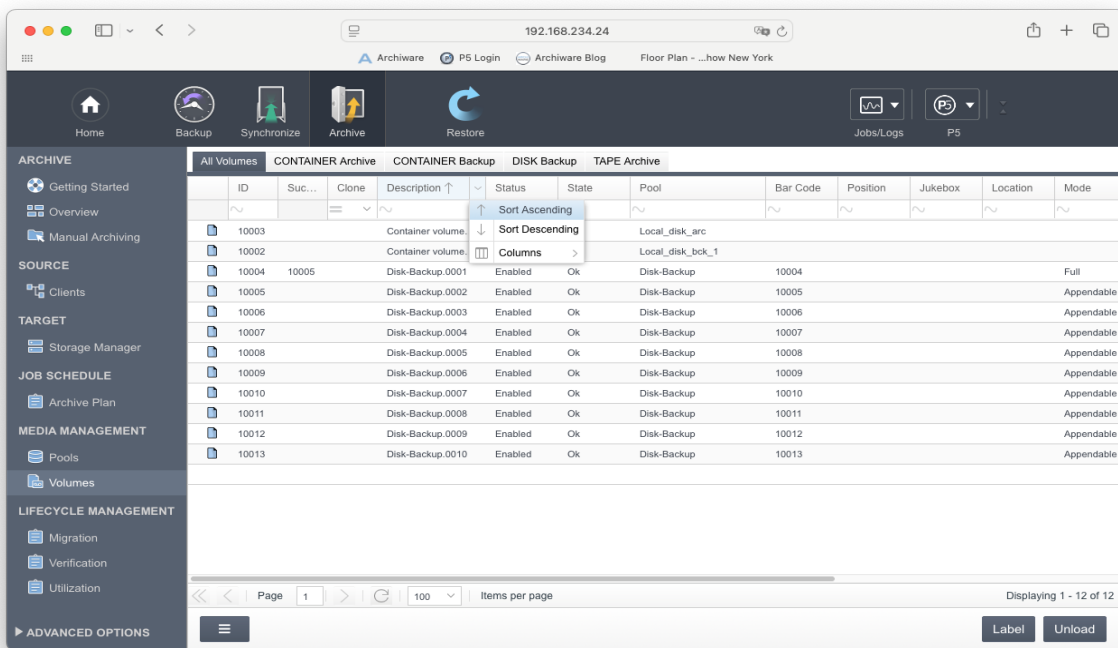
P5 supports both unencrypted authentication on SMTP port 25 and SSL/TLS encrypted authentication on port 465. Please select required authentication method and enter your SMTP server credentials.

11.4 Customized Views

As of Version 7, P5 automatically creates separate views. The number of views depends on the number of pools and volumes. When starting with P5, there is one view for each volume pool.

The table views can be customized. When clicking on the triangle next to a column headline, columns can be disabled or sorted. Useful for huge lists when there are hundreds of volumes filling multiple pages in the browser, additional filters may be defined to filter only a specific subset.

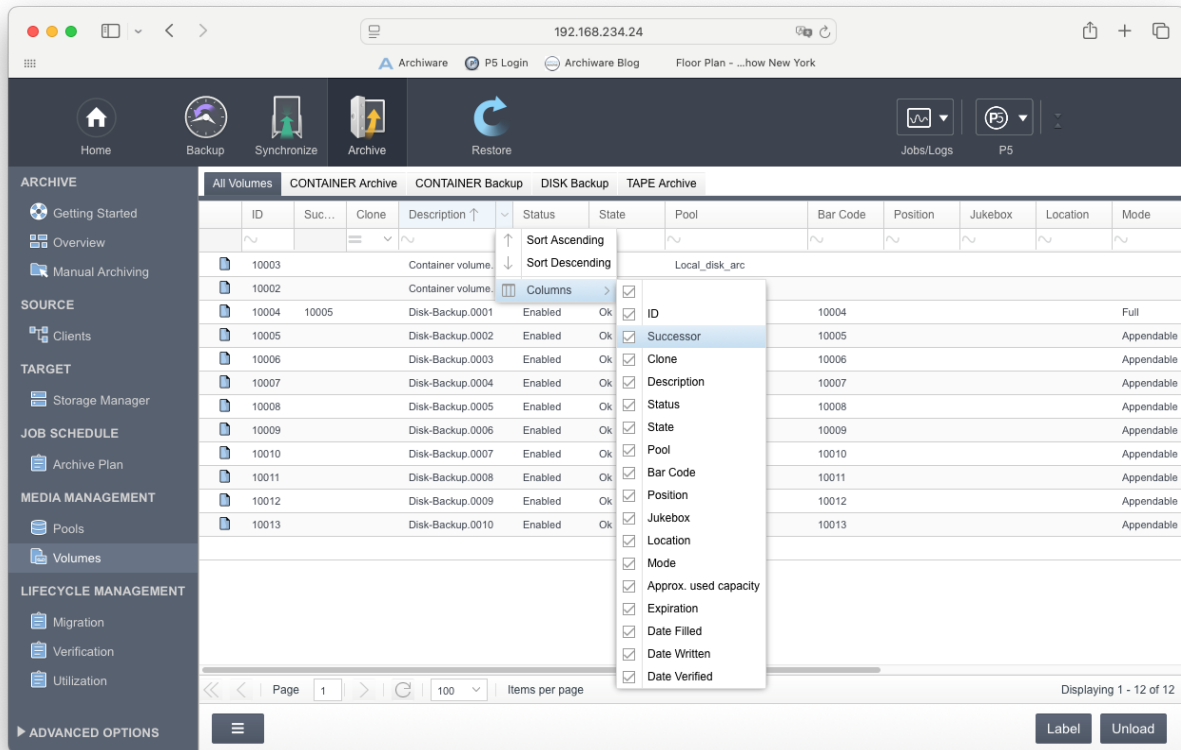
Views and applied filters are saved in the browser, not in P5 itself. So views cannot be made available to other users that make use of another browser instance and configuration.



For instance, in case the volumes are named using their bar code, both the name and the bar code columns contain the same information, so one can be omitted to enhance readability.

Filters can also be applied, for instance to focus on a specific set of tapes.

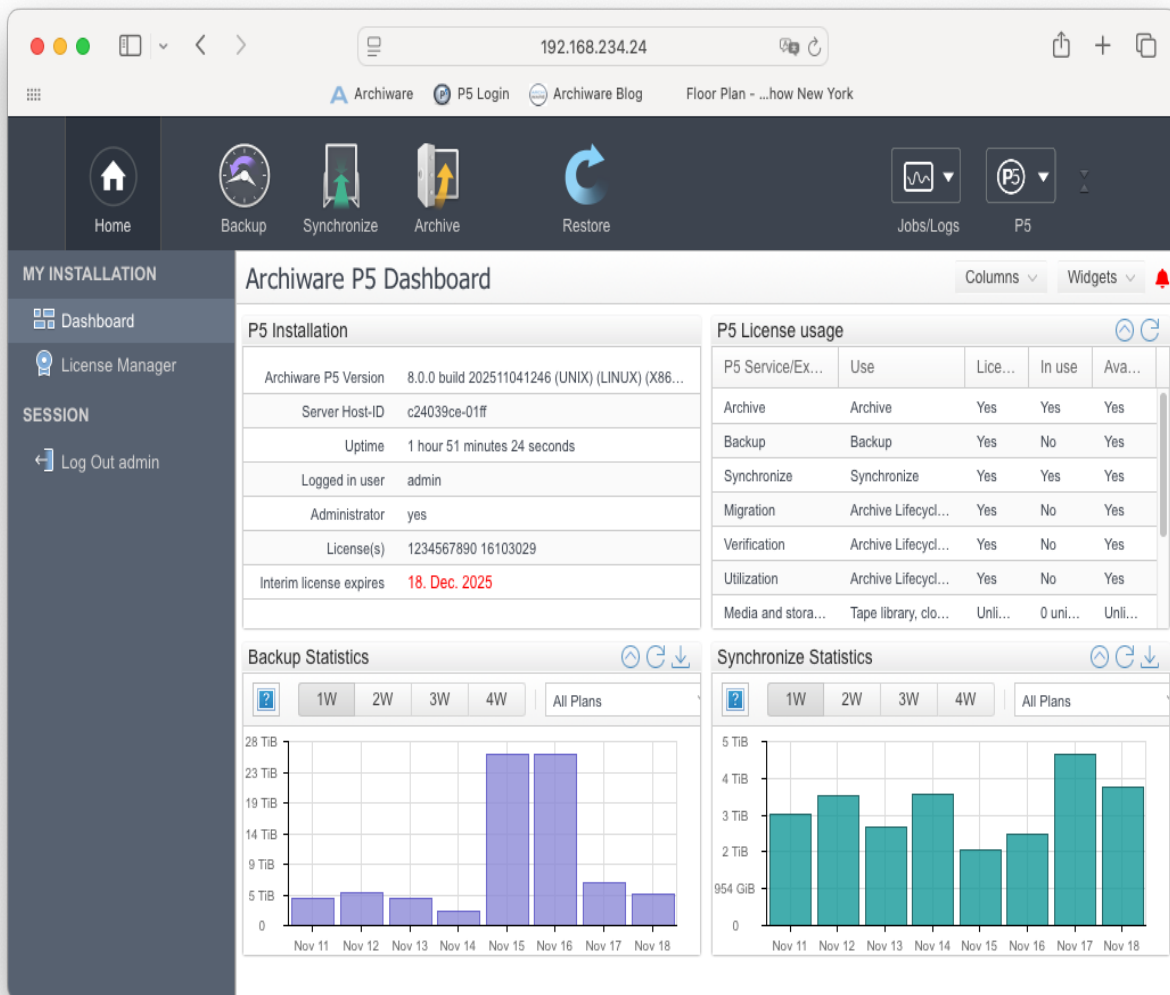
Active filters are visible in the filter field below the column header. Filters are not stored persistent and will be removed in a new browser session.



Filters are applied to the column they are defined in. A filter makes P5 show only those entries that contain the filter expression.

11.5 The Dashboard

Starting with version 7.4, P5 provides a customizable dashboard as entry screen:



Widgets can be added with the widget menu in the upper right edge and are automatically aligned according to the selected number of columns.

The following widgets can be selected:

Installation

P5 Installation Info	An overview of the installation. This Widget is always present
License Usage	The installed license(s) and their usage
Notification	Installation notifications
Welcome Widget	An entry widget that leads to the Archive, Backup and Synchronize Overview Pages

Jobs

Backup Statistics	Graphical overview for backup jobs of the last 1-4 weeks
Synchronize Statistics	Graphical overview for synchronize jobs of the last 1-4 weeks
Archive Statistics	Graphical overview for archive jobs of the last 1-4 weeks
Restore Statistics	Graphical overview for Restore jobs of the last 1-4 weeks
Job Results	An overview of jobs with their result state over the last weeks.

Storage

Tape Pool Usage	The tape pools and their fill state
Cloud/Disk Pool Usage	Disk and cloud pools with their sizes
Tape Performance	A visualization of the tape drive speed over the time

Note that the widget cannot auto-update, click on the refresh button on top of the widget to actualize the values.

11.6 Setting up a HTTPS Certificate

Starting with version 6, P5 supports https as communication protocol for browser connections.

So instead of the http URL <http://localhost:8000/login> it is also possible to access P5 via https using <https://localhost:8443/login>. The https protocol uses a SSL certificate that the browser will need to trust in when connecting. To achieve that, the certificate are intended to be signed by a trustworthy authority, and the browser checks that sign and the signing authority.

SSL certificates need the host to have a dedicated IP address and are sold by certifying companies.

To simplify initial access via https, P5 includes a so called self signed certificate. That means https access is possible, but the browser will when connecting to P5 the first time emit a warning that this access is not secured, as the certificate used by the web server, i.e. P5, is not signed by a known authority.

To install a file with a signed SSL certificate, please write that file as config/lexsrv.pem in the P5 base folder. By default that is

- On macOS or Linux hosts:
/usr/local/aw/config/lexsrv.pem
- on Windows hosts
C:\Program Files\ARCHIWARE\Data_Lifecycle_Management_Suite\config\lexsrv.pem

11.7 Setting up Login with Two Factor Authentication

By default, the login procedure in P5 is done by authenticating the entered username and password versus the credentials saved in the user preferences in P5.

With activated two factor authentication, P5 will send in addition a numeric code to the user's email account and request the code during login.

Before activating the two factor authentication, please check and ensure the following items:

- The mail interface in P5 must be set up and must be working. See section 11.3, Customizing Reporting for how to setup the mail interface.
- For each user who shall be able to login, a correct mail address must be set up in the user preferences so the code can be sent correctly. See section 6.14, User Preferences for how to edit the user preferences.
- Make sure that a login with the **admin** account is possible. In case the mail interface becomes unavailable (for instance because the mail server is unreachable), only the admin account has the possibility to bypass the two factor authentication.

To enable the two factor authentication, select the item **Advanced Settings** from the P5 Menu and set the value `modul.Login.multi-factor-authentication.configurable.boolean` to `1`.

In case the login should not be possible, there is a possibility to login for the **admin** user in the following way:

- Enter `admin` as user name
- Enter the password for the admin accounts
- After entering the password, check the file `config/mfa_admin.key` in the P5 install folder. That text file contains the numeric code for the login.

Note that in order to read the generated file, `root / administrator` permissions are required on the P5 server host.

12 Troubleshooting

You decided to use P5 in order to obtain the highest possible degree of safety and availability of your data. It is Archiware's topmost priority to offer you quick and flexible support in order to reduce downtime and provide optimal and smooth P5 operation. Additional features are also included to enable you to independently take actions in case of a P5 problem. This section describes the actions you may need to take to recover P5 operation in case of problems or after a crash of the hosting system.

These include:

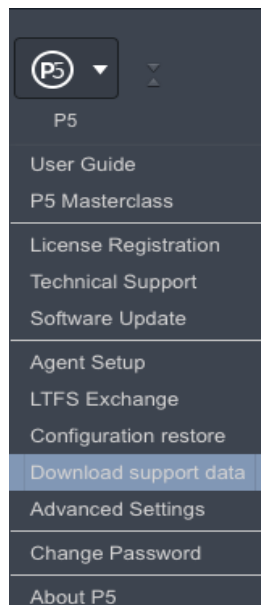
- downloading support data.
- restoring the configuration.
- disaster recovery.

12.1 Downloading Support Data

In most cases, it is sufficient to get in touch with P5 support to get your system running again. Archiware's support staff reacts promptly to your queries and is able, thanks to their extensive training and years of experience, to quickly locate and solve your problem.

To contact the support, please submit a ticket on the Archiware support page support.archiware.com/support. Please describe your problem there.

P5 logs each executed job. These logs are available to you both in the browser as well as in email form (see section 6.10, Logs) and can be used to monitor the application. All the actions performed by P5 are logged in enough detail to diagnose a problem and make a recovery possible. Furthermore, a feature is implemented in P5 that you can use to provide support staff with necessary information.

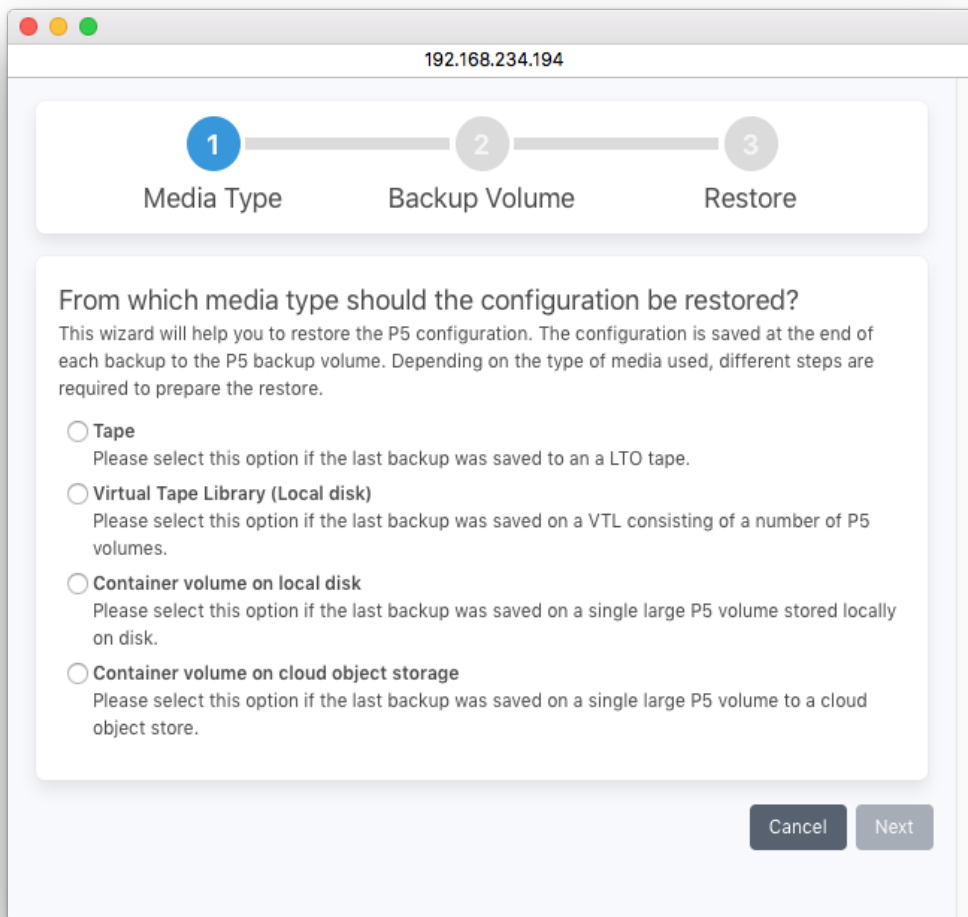


From the top level in the P5 browser, in the P5 menu you can use the `Download support data` feature. This will collect the P5 logs, system log and your P5 setup parameters file, pack and compress them and will download the resulting file and place it on your desktop or in 'Downloads' folder. Please send that log file if asked by the support staff.

12.2 Restoring the Configuration

The entire P5 configuration that is located on the hosting system is saved automatically with each backup job. This includes both the general setup (clients, databases, licenses, log-in areas and user preferences) and also the job and storage management (archive, backup, synchronize plans, as well as filters, the hardware configuration, drives, changers, volumes). This ensures that your system is recoverable at any time and that you can access older configurations and backups. Following a crash of the hosting system, the complete P5 configuration can be easily and quickly recovered. There is no need to set up P5 manually from scratch again.

In the P5 section of the top menu, you will find the function **Configuration restore**.



Following a crash of the hosting system and a new P5 installation, the device paths to your hardware will not yet be configured in P5. You will therefore need to enter the type of the storage and the path to the drive or container holding the volume from which to recover the configuration.

The following table shows some examples for the drive's OS Path:

Operating System	Path
macOS	/usr/local/aw/conf/dev/awst0 (This file is generated by P5)
Linux	/dev/sg0
Windows	{\\.\Tape0} (Use this syntax when the device is enabled in the Windows device manager) {\\.\Scsi1T2L0} (Use this Syntax when the device is disabled in the Windows device manager, where Scsi1 stands for the first adapter, T2 stands for SCSI ID 2 and L0 means LUN 0.)
Virtual juke box	/media/mydisk/PresSTORE_Volumes/10084 /C/VirtualJukebox/10084 (Please enter the path to the disk file)

By default, the most recent configuration will be recovered. Should you desire to recover an older configuration, you will need to enter the block number where it is located on the volume. You can find this information in the log of each backup job.

In case further info is required to restore the configuration, like a directory for a container store, an Object Store for recovering the configuration from the cloud or a decryption key for hardware-encrypted tapes, these can be set in the configuration restore dialogue.



Please note that restoring the configuration will overwrite the existing one.

12.2.1 Handling Encryption during Configuration Recovery

Encryption is used in P5 to store passwords, the encryption key is stored in the file ***localhost.key***. Ideally that file must be brought back in place before restoring the configuration, for instance from an external password manager.

The localhost.key file is not saved automatically by P5 for security reasons.

In case that is not possible, all passwords stored in the configuration cannot be decrypted anymore after recovering the configuration.

If so, reset the admin password to the default ***admin*** by placing a file ***resetadminpassword*** in the subfolder ***config/customerconfig***.

The administrator then can login with the default credentials admin/admin and needs to reset all passwords.

This affects:

- All passwords in the User Preferences in order for other users to login
- The passwords in the Clients section to connect to clients
- The passwords/access keys in all Cloud Services in the Storage manager
- Hardware encryption keys (if set) into the P5 LTO key manager



In case the P5 storage encryption was turned active in the localhost client entry of the P5 server, the localhost.key file is required in place to recover the configuration, the recovery will otherwise fail.

In case the LTO hardware encryption was active when the configuration was saved, the decryption key must be entered in the configuration restore dialogue.

12.2.2 Disaster Recovery

Basically there are two situations where a disaster recovery becomes necessary: when either a client or a P5 server crash. The data of any backed up client can easily be recovered on any other client or machine and also on a newly installed one. This can be done by Restore (See section 8, Restoring Data From Backup).

To recover a crashed system on which the P5 server was located or to install a new system in its place, several steps have to be taken before routine P5 operation can be resumed:

1. Install a new system, including a new P5 instance
2. Recover the P5 configuration (See section 12.2, Restoring the Configuration).
3. Recover the backup and archive indexes.

Because backup and archive indexes will be missing in the newly installed P5 server, they have to be recovered out of the last backup. Since the backup indexes are backed up automatically with each backup, they can be recovered with the **Check Database** function.

To do so, follow the steps described in section 6.12.5.2, Check Database.

In case of a disaster recovery in which the backup indexes are not available, phase 1 where an automatic index repair is attempted will not be executed and phase 2 will immediately start. In this phase, the volumes which are necessary to recover the index will be identified and listed so that you can make them available to the P5 media management. Next, the backup indexes will be read from the backup media and recovered into the P5 server.



Contrary to the backup index, if an archive index could not be corrected by the check index function, it cannot be recovered automatically since it does not get saved automatically. It is therefore recommended to back up the archive indexes with a backup plan, so that in the event of a corruption you can recover them from the backup.

In case the archive indexes have been saved by a backup plan, they also can be recovered with Restore (See section 8, Restoring Data From Backup).

4. After completing this step, the new P5 server will be completely recovered and ready to start routine operations again.

12.3 Fixing Tape Drive Problems

This section lists some standard procedures for the most common problems that may arise with tape drives and tape libraries. Both are, however, rather complex technical devices and it is therefore not possible to resolve every error that may come up.

There are two main groups of problems: Tape access problems when external tape device is not found or it cannot be set up and tape operation problems when errors occur during operation for devices that were operable before.

12.3.1 Tape Access Problems

Read this section in case you cannot configure the device or in case you get error messages like No such device or address or Open Error. These errors indicate that the device is either not accessible or not accessible any more.

Please check whether the device is visible to the operating system:

On Windows, check whether you see the device in the Windows Device Manager

On macOS, click the Apple menu → About this Mac → System Report

On Linux, call in a terminal
`cat /proc/scsi/scsi`

to list the connected devices

On Solaris, unload and load the generic SCSI driver and check for attached devices by calling

```
rem_drv sgen  
add_drv -v sgen
```

P5 relies on the generic SCSI driver of the Operating System to access external devices. In case the devices are not visible to the Operating system, there is no way to access them in P5. In addition to each tape drive (usually identified by its serial number and device vendor), the media changer (if present) must be visible as a separate device.

In case some devices are not visible, please perform the following checks:

Reboot or try to rescan the SCSI bus:

On Windows, call Scan for new devices in the Windows Device Manager

On macOS, rebooting is the only option

On Linux, call in a terminal
`modprobe -r sg`
`modprobe sg`

to reload the generic SCSI driver

In case the device does not become visible, the following checks may be helpful:

1. Check whether all cables are connected correctly, for parallel SCSI devices check also the terminators. Then power cycle the device, reboot and recheck.
2. If the error persists, check whether the host bus adapter is installed correctly and try replacing the cables. Then power cycle the device, reboot and recheck.
3. If possible, try checking with another tape drive and check the hardware manual.

When the devices are visible to the Operating System, but are inaccessible in P5, delete the device in the Storage Manager (for media libraries) or the Standalone Device section and add it anew. In case there are multiple devices configured, do so with all tape devices.

12.3.2 Tape Operation Problems

You should read this section in case the device was once operating correctly and a tape or changer related error appears. Here are some common error classes and some standard procedures how to handle them:

In case you get one of the following errors:

```
SCSI Bus reset
SCSI Device reset
I/O Error
Mechanical positioning error
transfer element full
internal target failure
invalid element type
incompatible format
```

These errors may indicate a configuration problem of the library itself or a problem of the configuration in P5. Delete the library configuration in P5 and add it anew, power cycle the device and reboot.

If the problem persists, check whether there is a mechanical problem, maybe a tape is stuck somewhere in the changer and cannot be moved.

In case a media changer shows errors when transporting, loading or unloading tapes and the error is one of the following:

```
Destination element full
Source element empty
```

This error may have one of the following reasons:

- P5 keeps a strict assignment between a physical tape and the slot in the library where it has been found. Make sure you do not insert a new cartridge into the slot of the tape that is currently in the drive. It is good practice to empty the drives before changing tapes. In case this appears, reset the changer from the P5 Storage Manager. Repeat the reset once if that did not help.
- Make sure there are no more tapes inserted into the changer than there are slots. Even if it is possible to get 21 tapes into a 20-slot-changer with one drive, the changer will then be unable to move any of the tapes (as any possible destination slot is occupied). Remove at least one tape.

12.4 Network Connection Problems

This section is about troubleshooting problems concerning the connection between P5 server and client. Before starting, please check whether it is possible to connect with a browser that is running on the same host by entering the URL <http://127.0.0.1:8000/login>. In case the browser cannot connect but shows an error, the problem is probably local to that machine, for example if P5 is not running or fails to start correctly.

12.4.1 The Network Connection Cannot be Established at All

Start here after having checked that P5 is accessible locally on both machines, but one machine cannot connect to the other one.

At first call the ping command in a terminal on one host to see if the remote host is visible. You should get a reply from the other side like in the following example:

```
# suse:~ ping 192.168.234.24
PING 192.168.234.24 (192.168.234.24) 56(84) bytes of data.
64 bytes from 192.168.234.24: icmp_seq=1 ttl=128 time=0.194 ms
64 bytes from 192.168.234.24: icmp_seq=2 ttl=128 time=0.187 ms
64 bytes from 192.168.234.24: icmp_seq=3 ttl=128 time=0.189 ms
```

In case you get a message like 'destination host unreachable', check the following items:

- Are both hosts connected to the network
- In case you addressed a name instead of an IP address:
check whether the name resolution results in the correct address
- Is there a firewall between the two machines that blocks the communication
- Is the IP routing set up correctly

Please check the documentation of your operation system for how to change these settings.

In case the ping gets a reply as expected, please try to connect with a browser to the remote machine, you will then need to pass the remote address to the browser. If e.g. the remote host is at the IP address 192.168.0.12, please enter in the browser <http://192.168.0.12:8000/login>

Do you get a login mask from the remote P5 instance? If so, also connecting that host as a client should work flawlessly. In case you do not get a login mask from remote (but you get it when accessing P5 locally on the remote machine), recheck the items above.

12.4.2 Bad Connection / Connection Breakdown

In case the connection between two P5 instances can be established but the connection breaks down, please check the following:

- Is there an internet security suite or an http proxy in use between the two hosts? P5 uses a http protocol with modifications. During pure data transfer, the protocol is not fully browser compliant, some security suites may consider that as a hidden attack. In that case, please define an exception in the security suite.
- The second probable reason is that there are IP components in between that do active buffering, namely where networks with different speeds are involved (e.g. when transferring over a VPN or an internet). In these setups, please edit the client settings on the P5 server and increase for that client the "Limit bandwidth:" value. The trick here is that any value except 100% introduces speed measurement and delayed transmission so that buffering on intermittent components will be limited.
- In some rare cases, a breakdown occurs when a problem arises on the client, e.g. a power down or a crash. In that case, the logs must be checked to find out why the client stopped.

12.4.3 Network Ports and Protocols used by P5

P5 always uses a simple HTTP protocol in both directions. In data connection such as between a P5 server and a client or workstation, the protocol is slightly extended with variable length packages.

By default, P5 uses TCP port 8000 for all communication. For firewall exception rules please note that in Backup, Synchronize and Archive, the P5 server connects to the P5 client's port 8000. (as opposed to the traditional server/client scheme used in TCP scenarios).

13 Best Practice

You decided to use P5 in order to obtain the highest possible degree of safety and availability of your data. To simplify your first steps with P5 and to assist you in your routine tasks, this documentation not only contains a description of the application, but also instructions for using (section 8), operating (section 10), customizing (section 11) and trouble shooting (section 12) the program.

This section describes, by way of example, some typical scenarios often used as solutions for common backup and archive approaches. By following these examples, even a new P5 user has a tool to quickly and efficiently employ P5.

13.1 Prerequisite for Backup and Archive

One frequently asked question is what the differences are between backup and archive. The following hints help to decide which module to use for saving the data to tape:

Backup in general is a safety mechanism to create copies of data available on disk onto a second media. To keep the backup copy updated, data is usually rewritten to the backup media in regular intervals. For that purpose, the backup is assigned a retention time, after which the data expires. It is good practice that the data retention time (available in the backup plan) is set after the data has been written to one or multiple other media.

Use the Backup module when:

- The data should be kept on disk and the backup is intended as a copy.
- A protection against server crashes or data loss should be created by writing a copy that is kept in the office, or
- A protection against theft or environmental damage need to be created by a copy that is stored off site.

Archive is a different module that is not intended as a safety copy but as an offline data storage for data that is no longer required online. Use the archive module when:

- Data should be archived for long term storing
- The data should be removed from the online storage

In order for backup or archive to be able to write to tapes, the media to write to must be prepared with the following steps:

1. Configure the device that should be used to write the data. That may be a standalone single tape drive, see section 6.2.2, Single Tape Drives, a media changer or library, see section 6.2.1, Tape Libraries / Media Changers, or container storage for disk or cloud, see section 6.2.3, Container Storage.
2. Define a set of media to write to. In P5, this is done by defining a media pool, see section 6.3.1, Creating a Pool . Multiple pools can be defined for different purposes, for example when creating different archives or when creating multiple backups that should be kept separate.
3. Make the volumes available for P5 by labeling the media, i.e. by assigning the physical media to a media pool, see section 7.5.3.1, Label Volume.

When these steps are done, a backup or archive can be defined by creating an archive plan (see section 6.5.1) or backup plan (see section 6.6.1).

13.2 Backup Scenarios

This section illustrates some standard scenarios to create backups. In order to extend these scenarios with a backup to disk to tape functionality or as an alternative setup for small business, see also section 13.2.5, Using Synchronize as Backup to Disk.

13.2.1 Standard Backup for Small Business

A simple and easy to understand traditional setup for the backup on discrete media like tapes is the grandfather – father – son paradigm:

- Create a backup plan and schedule a full backup in regular intervals on the weekend, for example weekly.
- Add a second schedule for incremental backups that start on weekdays after business hours.

P5 will then write a full backup on a weekly basis and extend that state with incremental backups over the week. Combined, these two plans make up a backup cycle.

The next cycle, starting with a new full backup, should then be written to a new set of tapes, the previous set may be stored in a safe place or, if there is enough space available, kept in the media changer.

When the third backup cycle is again written to a new tape set, the three "generations" are complete: cycle1 as grandfather, cycle2 as father and cycle 3 as son. The oldest tapes can then be overwritten by the next backup cycle. In this particular scenario, each tape is recycled (and used again) after three weeks.

It is possible to plan more than three cycles to keep older backups, or only two cycles in case simple alternation of the tapes is sufficient.

To maintain only a single cycle is not permissible. This would imply that the old backup data is removed before a new backup is written. This backup method is potentially insecure.

A single, separate media pool is recommended to achieve the G-F-S backup plan described above.

The data retention time in the backup plan should be set to a value that is a few days shorter than the backup cycle time multiplied with the number of cycles in order to allow P5 to automatically recycle the tapes. In our example with three cycles of one week each, a data retention time of 18 days is a good value.

The "Volume usage" field in the backup plan should be set to "Use new Media", to ensure that P5 starts each cycle on new tapes.

13.2.2 Progressive Backup for Huge Amounts of Data

P5 has a special backup mode intended for huge datasets where the time window available to perform a full backup is insufficient. Progressive backup can be activated in the Backup Plan's Advanced Options.

P5 will, since all backups are run in the same backup cycle, ensure that before a tape is recycled, the data stored on that tape is saved again. In other words, each saved file always remains saved on at least one tape until the data retention time is over. This implies that in the worst case the number of tapes required is twice the amount of tapes that would be required for a full backup plus the tapes required for incremental updates during the data retain time.

In progressive mode, the time to run a backup must be limited. Once the retain time is reached, P5 will start to backup expired files again, together with new or changed files, as far as possible in the time window. After some time, the files, as they were not all saved at the same time, will expire at different times, so only the expired portion of files is included in the backup job. The time window must be sufficient to save all new and changed files and allow saving a part of the expired files.

In case the progressive backup runs until the time limit is reached, there are still files left to save, these will then be saved with the next job. In case the backup is complete, it should end before its time limit.

In case the progressive backup does never end within its time limit, either the time window is too short or the retain time must be extended so that less data expires at the same time.

Note that the progressive mode should only be used with tape libraries or virtual disk libraries.

13.2.3 Backup with Alternating Pools

The backup module allows to define alternating backup pools. Setting up alternating backup pools defines a backup plan that writes up to three independent backup data sets. P5 will then automatically select the first of the pools that it finds available at backup time and adds all data that is missing in that pool. So a plan can be defined that dynamically adjusts to the tapes that are made available and saves the data missing in that particular pool.

To setup a plan with alternating pools just define a backup schedule where as many pools are defined as backup sets shall be created. Up to three pools can be added. The definition of the backup pool is done in the Backup Event, see section 6.6.1, Creating a Backup Plan. In case multiple backup events shall be defined, you may want to define the same pools in all backup events.

This mechanism can be used in conjunction with the progressive setup with automatic recycling as well as with a normal full / incremental backup schedule.

Note that the use of alternating pools in one Backup Event is different to defining multiple Backup Events with a separate media pools each. The later will not create independent backup sets but write one complete set of all data across multiple pools.

Note that alternating Pools will not work with Container Backups, as there are regarded as non-removable, i.e. P5 would only make use of the first pool as it is present (if the storage is not present, that may result in access errors).

13.2.4 Archiving Methods

There are two general methods to organize archiving with P5:

- When archiving **manually**, the archive plan can be considered as a parameter-set used to define the volume pool, cloning and further parameters. To archive manually, the files and folders must be ready and the drive and the volumes must be available. This method is usually preferred in case data from an external source needs to be archived or when archiving is done by a single person. See section 9.1, Archiving Files for how to archive manually.
- When archiving **automatically**, the archive plan is set up to start on its own to gather the files and folders to archive from a **drop folder**. The archive is usually started overnight and deletes the files from the drop folder after archiving.

This method is preferred when the data must be collected from different sources to be prepared for the archive, or when multiple users need to be provided with an easy method to archive data. Often, this method is used to archive data from bigger online storage servers. The automatic setup can be easily configured in the archive plan, see section 6.5.1, Creating an Archive Plan.

- A **variant** of the **automatic** archiving from a drop folder is to archive the whole server volume, but with specifying a filter to archive only folders with a specific name, possibly on multiple volumes. In order to archive folders with this method, it is no longer required to move them to a drop folder, but instead rename the folder to fit the filter requirement. For example, adding the extension ".ToArchive" could be such a filter expression, so only folder with that extensions are archived.

When archiving, data is often removed from the online storage. Doing so may make the copy on tape the only existing version of the data. Because of that, a backup copy of the archived data is recommended. P5 supports that by cloning the archived tapes. Cloning means that two tape drives are used to write two archive tapes simultaneously. In case an archive tape becomes damaged or no longer readable, this allows to rebuild the tape from its clone without data loss. Cloning requires a second tape drive and can be set when creating the media pool, see section 6.3.1, Creating a Pool . The recovery of a tape from its clone is described in section Error: Reference source not found, Error: Reference source not found

13.2.5 Using Synchronize as Backup to Disk

When writing to disks, the Backup module that emulates a media changer on disk cannot make use of the physical advantages the file system on disk can provide. As an alternative, the Synchronize module can be used to create a backup copy of the data on a backup disk.

With Synchronize, a one to one copy of all the files is created at a separate location. This includes all the file attributes, so that this location holds an exact copy of the source. The separate location can be a separate folder, another drive or a remote machine. Like the Backup module, *Synchronize* can be restricted to specific directories, so one can create a snapshot of specific folders or of the whole volume.

The Synchronize module has two features that allow creating a complete backup with multiple copies, cycles and versions:

- The **cycles** allow to save multiple snapshots of the source at different stages of the file system: When Synchronize cycles are set in the Synchronize plan, P5 will maintain a directory `__CYCLES` on the destination and will keep a complete set of data each time a mirror mode synchronize is executed. In other words, it will create a new snapshot. This feature is comparable to the backup cycles configured in backup plans. You can choose the number of such cycles that you wish to keep.
- The **versions** feature allows an even more exact configuration. After creating a snapshot at the target location, Synchronize can keep the target up to date by performing incremental updates in short regular intervals, for example every 30 minutes. If doing so while a file is being edited on the source, the destination file will be updated and the old version would normally be lost. With the versions enabled, Synchronize will keep previous versions of each file in a separate `__VERSIONS` directory.

Together with cycles and versions, the Synchronize module used for backup has multiple advantages:

- Files are directly accessible without management overhead
- Previous versions can be kept to go back a few hours in the working process
- Cycles allow to save the state of one or two weeks ago, just like a grandfather-father-son backup would do.

Find more about how to setup synchronize in section 6.7.1.

13.2.6 Backup to Disk to Tape

Backup-to-disk-to-tape is a commonly used scenario that combines the advantage of a quickly accessible direct restore and the faster possibility to update the backup as provided by Synchronize with the backup to tape in order to create transportable backups that can be stored off line.

To implement a backup-to-disk-to-tape scenario, set up the backup to disk as the first backup stage as described in section 13.2.5. For the second stage, use the backup module to save the synchronized data to tape.

The backup module will automatically filter the `__CYCLES` and `__VERSIONS` folders and thus save the latest synchronized state. Therefore, a backup of the synchronize target will save the same data as if the original data would have been saved directly without overhead of additional copies on the Synchronize target disk.

Another benefit of this method is that Synchronize performs delta updates with lower data transfer load than a full backup would require. This method can therefore be used to increase the time available for backup. To achieve that:

- Schedule Synchronize to perform a backup to disk, for example once a day when the generated network load and traffic on the server disks does not impact daily business.
- Throughout the day, a backup can be written to tapes from the synchronized data. In case the synchronize target is locally connected to a dedicated backup server, this operation can be performed locally on the backup server without load to the production system.

13.2.7 Backup to Removable Disk Media

Under most conditions, when running backup to disk (or cloud), the container backup is a more ideal solution because it can selectively reuse single containers and thus needs to transfer less data. But there is one exception: The container does not handle removable media; so when running a backup to RDX cartridges or other removable disks like USB connected disk drives, setting up these as Virtual Tape Library (VTL, for big drives with multiple virtual disk volumes) or removable drives (for single virtual disk volumes) make more sense.

P5 will then treat the disk volumes like tapes. The single volumes are then written strict sequentially, and in order to overwrite a volume, it must be recycled (which clears the complete volume) and can then be overwritten, same as a recycled tape.

In order to enable the use of Virtual Tape Libraries, select the item `Advanced Settings` from the P5 Menu and set the value `gui.Jukebox.mediaType.DISK.configurable.boolean` to `1`.

In order to enable the use of removable single disk drives, select the item `Advanced Settings` from the P5 Menu and set the value `gui.Device.deviceType.Removable.configurable.boolean` to `1`.