

ARCHIWARE P5

Reference Manual





Table of Contents

1 Preface.....	6
1.1 About This Manual.....	6
1.2 Who Should Read This Document?.....	6
1.3 Typographic Legend.....	6
2 Introduction.....	7
2.1 Backing Up.....	7
2.2 Synchronization.....	8
2.3 Archiving.....	9
2.4 Backup2Go, the Portable Backup.....	9
2.5 The Configuration.....	9
2.6 Program Description.....	12
2.7 Inside P5.....	12
2.7.1 The Data Backup.....	12
2.7.2 Recovering Data.....	13
2.8 Software Licensing.....	13
2.9 Registration.....	14
2.10 Technical Support.....	14
2.11 Documentation and Help.....	14
3 Installation.....	15
3.1 Overview.....	15
3.2 System Requirements.....	15
3.3 System Resources.....	16
3.4 Installing on Mac OS X.....	18
3.5 Installing on Windows	21
3.6 Solaris and Linux Installations.....	24
3.6.1 Installation Requirements.....	24
3.6.2 Mounting the P5 CD.....	24
3.6.3 Creating an Installation Directory.....	24
3.6.4 Copying the Installation File.....	25
3.6.5 Unpacking the Program Directory Structure.....	25
3.6.6 Starting the Installation Script.....	25
3.6.7 Special considerations when installing on 64-bit Linux System.....	26
3.6.8 Testing the Installation.....	28
3.7 Upgrading to P5.....	28
3.7.1 Hints regarding the upgrade.....	28
3.7.2 Program Update.....	29
3.7.3 Upgrade from P4 to P5.....	30
4 Signing In and Entering the License.....	33
4.1 Starting the Browser.....	33
4.2 License Activation.....	34
4.3 A General Introduction to the Server's User Interface.....	35
4.3.1 The Installation Overview.....	38
4.4 Using the P5 Agent Setup GUI.....	39
5 First Steps With P5.....	40
5.1 First Steps to Set Up the Backup.....	40



5.2 First Steps to Set Up an Archive Plan.....	41
5.3 First Steps to Set Up a Synchronization Plan.....	42
5.4 How to Set Up Backup2Go.....	43
5.4.1 Setup the Backup2Go Server	43
5.4.2 How to Set Up a Backup2Go Workstation.....	45
6 Configuring the P5 Server.....	47
6.1 Licenses.....	47
6.1.1 Creating a License	48
6.1.2 Importing a License.....	49
6.2 Storage Devices.....	50
6.2.1 Creating a Stand-Alone Tape Device.....	50
6.2.2 Creating a Removable Disk Storage.....	55
6.2.3 Devices: Base Buttons.....	58
6.3 The Media Library (Changer or Jukebox).....	59
6.3.1 Creating a Jukebox	59
6.3.2 Create a Virtual Jukebox.....	63
6.3.3 Modifying a Jukebox Setup.....	64
6.3.4 Storage Manager: Gear Menu.....	67
6.3.4.1 Label Volumes.....	67
6.3.4.2 Take Library Inventory.....	67
6.3.4.3 Load/Unload Slot.....	68
6.3.4.4 Import Volume.....	69
6.3.4.5 Reset Jukebox.....	69
6.4 The Volume Pool.....	69
6.4.1 Creating a Pool	70
6.4.2 Pool Configuration: Additional Options.....	72
6.4.3 Pool Configuration: Concurrency Options.....	73
6.5 Volumes.....	75
6.5.1 Creating a Volume (Labeling).....	75
6.5.2 Modifying a Volume	75
6.5.3 Volume Manager: Gear Menu.....	79
6.5.3.1 Label Volume.....	79
6.5.3.2 Regenerate Volume	80
6.5.3.3 Migrate Volume(s).....	81
6.5.3.4 Job History.....	82
6.6 Archive Plans.....	83
6.6.1 Creating an Archive Plan.....	83
6.6.2 Archive Overview.....	94
6.7 Backup Plans.....	95
6.7.1 Creating a Backup Plan.....	95
6.7.2 Backup Overview.....	102
6.8 Synchronize Plans.....	104
6.8.1 Creating a Sync Plan	104
6.8.2 Synchronize Overview.....	111
6.9 Filters.....	113
6.9.1 Creating a Filter	113
6.10 Cloud Storage.....	116



6.10.1 Defining a Cloud Storage.....	116
6.10.2 Supported Cloud Storage.....	117
6.11 Logs.....	118
6.11.1 The Job Monitor.....	118
6.11.2 The Log History.....	119
6.12 P5 Clients.....	120
6.12.1 Creating a Client.....	120
6.12.2 Client: Interactive Functions.....	125
6.12.2.1 Ping.....	125
6.12.2.2 Update Client.....	126
6.12.3 Setting up Encryption for a Client.....	127
6.13 The Backup2Go Server.....	128
6.13.1 Editing the Backup2Go Template	128
6.13.2 The Workstation Manager	134
6.13.3 The Workstation Record	134
6.13.4 The Workstation Overview	139
6.14 Index Databases.....	141
6.14.1 Creating a Database	141
6.14.2 Database: Functions.....	143
6.14.2.1 Fields.....	143
6.14.2.2 Check.....	144
6.15 Index Access Control (Login Areas).....	145
6.15.1 Creating a Log-In Area	145
6.16 User Preferences.....	148
6.16.1 Managing User Preferences.....	148
7 Configuring the P5 Backup2Go Workstation.....	153
7.1 The Workstation Setup Wizard	154
7.2 Manual Workstation Setup.....	155
7.3 The P5 Workstation Main Screen.....	160
7.4 Restoring Files on the Workstation.....	161
8 Restoring saved files.....	162
8.1 Restoring Data From Backup.....	162
8.1.1 Restoring Files.....	163
8.1.2 The Restore Selection.....	164
8.1.3 The Restore To Button.....	165
8.1.4 The Index Options.....	166
8.1.5 The Versions Button.....	167
8.1.6 The Search Button.....	168
8.2 Restoring data on Backup2Go Server.....	170
9 Archiving and Restore from Archive.....	171
9.1 Archiving Files.....	171
9.2 Restoring from Archives.....	173
9.2.1 The Restore To Button.....	174
9.2.2 The Restore Selection.....	176
9.2.3 The Versions Button.....	176
9.2.4 The Search Button.....	177
9.2.5 Index Options.....	179



10 Operating Procedures.....	180
10.1 Labeling Media.....	180
10.2 Mounting Volumes.....	181
10.2.1 Mounting or Un-mounting a Volume.....	181
10.3 Replacing Volumes in a Library.....	182
10.3.1 Jukebox Inventory.....	183
10.4 Handling Media Pools.....	184
11 Customizing.....	185
11.1 Customizing Index and File System Access	185
11.2 Customizing Reporting.....	186
11.3 Customized Views.....	188
12 Handling LTFS volumes with P5.....	190
12.1 Creating a LTFS volume.....	190
12.2 Importing an LTFS volume.....	190
12.3 Mounting and Unmounting an LTFS volume.....	190
12.4 Accessing an LTFS volume for reading and writing.....	191
13 Troubleshooting.....	192
13.1 Downloading Support Data.....	192
13.2 Restoring the Configuration.....	193
13.2.1 Disaster Recovery.....	194
13.3 Fixing tape drive problems.....	195
13.3.1 Tape access problems.....	195
13.3.2 Tape operation problems.....	196
13.4 Connection problems.....	197
13.4.1 The network connection cannot be established at all.....	197
13.4.2 Bad Connection / Connection Breakdown.....	198
13.4.3 Network Ports and Protocols used by P5.....	199
14 Best Practice	200
14.1 Prerequisite for Backup and Archive.....	200
14.2 Backup Scenarios	201
14.2.1 Standard backup for small business.....	201
14.2.2 Progressive Backup for huge amounts of data	201
14.2.3 Backup with Alternating Pools	202
14.2.4 Archiving Methods.....	202
14.3 Using Synchronize as Backup to Disk.....	203
14.3.1 A fast accessible backup for small businesses	203
14.3.2 Backup to disk to tape	203



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?

This information is intended for the system administrators who will install and maintain P5 as well as for the user 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 in this document are used to simplify its use.

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</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</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 <code>Volume Listing</code> .
<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 <i>Login Area > Administration > Backup Management</i> .
	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 MAC 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 in the night, during which such huge data volumes have to be backed up. This must be done automatically 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 in the given and limited 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 that all the required operations in conjunction with the underlying hardware require a minimal amount of configuration, monitoring and maintenance. P5 logs all its operational steps and activities



and e-mails them on demand.

The backed up data is indexed into a file structure, which 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 index file. 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 do not change, and restore them to their correct new place in the file system. Furthermore, intentionally deleted files will also be detected and not restored, thereby maintaining the exact picture of the file system at any particular point in time. Additionally, the incremental backup takes care of continually 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 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 contains 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 employing 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.

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 lie on different machines, one running a P5 server, the other a P5 client, or vice versa.

There are several advantages of 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 discs or disc arrays to external media to better utilize the limited on-line disc 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. On 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 Backup2Go, the Portable Backup

Backup2Go is specially designed for backing up workstations. In today's corporate environments, traditional desktop PCs are increasingly being replaced by laptop computers. This fact and the necessity to provide a backup for a large number of such workstations, demand a fast, flexible and easy-to-use solution.

Backup2Go answers these needs by providing a backup to disk solution for workstations that will back up their data as long as they can reach the server and will wait while they cannot. With its highly generic approach, new workstations can be added with minimal effort with virtually no administrator intervention.

Furthermore, the users can easily restore their data themselves by simply selecting files from the backup that is presented as a file system in their web interfaces.

2.5 The Configuration

The configuration entities in P5 such as devices (drives or changers), volumes, or clients are called resources. These are entered and kept in the resource (or configuration) database. The database has a tree like structure that you can navigate through. 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 the group has a unique ID, description and several further 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 and limited resources between its diverse modules, for example to restore different file off multiple volumes, where jobs where initiated from different users at the same time.

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

The following table 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 can choose from a list of predefined and pre-configured plans, thus not having to be 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 others 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. Using the appropriate drive, the volumes are labeled, which meaning that a label is physically written on them to uniquely identify them.	Pool
Pool	Each volume (labeled media) belongs to a pool. The pools are an administrative and a control element. The job may be directed by means of the pool to use and reserve (and limit the use of) dedicated resources.	Archive Plan, Backup Plan, Device, Volume
Device	Entries in this resource group are associated with drives or other storage elements which P5 supports. The drives may be stand-alone or 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, underlie the control and are executed as jobs of the job scheduler. 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 via user preferences. Using the attributes, users may be granted use of specific operations or access to diverse areas in P5.	
Workstations	Each workstation connected to the P5 Backup2Go server is maintained in a workstation record.	Template
Backup2Go Servers	The Backup2Go workstation keeps its server connection information in a server record.	
Backup2Go Templates	An instance of the Backup2Go template is used to create a new workstation record.	Work stations

Table of Resource Groups



2.6 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 run or changing the attributes of a resource.

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

For Mac OSX:

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

Note that up to OSX 10.4, there will be a second *nsd* process same as on the Unix systems below.

For Solaris and 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 its functions. P5 consists of a series of further modules, which are run and controlled by the application server. With its integrated Web server, the entire user interface is controlled from an Internet browser, thus making further client applications superfluous.

2.7 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 discs 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 changer or jukeboxes) and automatically manages the media.

2.7.1 The 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 according to each of the backup plans at its configured time. The job manager



overtakes the supervision of the job and takes into consideration all the configured attributes. It request 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. Now the actual backup begins. If one of the volumes gets 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.

2.7.2 Recovering Data

The user initiates a data recovery. Using the P5 data browser, he/she marks a file or directory to recover. From the recover context menu, further options controlling the recover 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 exchange the volume. The operator has to exchange the volumes manually upon request in case that no changer is available. The files and folders will be recovered identically to the original.

2.8 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 further machine (PresSTORE client) you wish to back up, you will require an additional server or workstation agent license. For each further 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 up to ten 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 further machine (P5 client) from which you wish to archive files, you will require an additional server or workstation agent license. For each further 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 further 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.



2.9 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.

2.10 Technical Support

Contact information, software updates, technical documentation

- n and information about the current support programs and options may be found on Archiware's web site.

2.11 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 the documentation or selected parts of it.

This documentation can be found at the following locations:

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



3 Installation

This section describes how to install P5 on Mac OS X, Windows, Linux and Sun Solaris.



- 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.

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

3.2 System Requirements

P5 is available for the following platforms and operating systems:

P5 Server

- Mac OS X: Intel 10.5+
- Windows: Server 2008R2, Server 2012, Windows 7, Windows 8
- Linux: Kernel 2.6.4+ (x86) 32bit or 64bit with 32bit libs
For FS-Events with inotify, Kernel 2.6.13+ is required
For B2G repository on btrfs, btrfs 0.19+ is required
For B2G repository on zfs, zfs 0.62+ is required
- Sun Solaris: Sparc Solaris 2.8+, x86 Solaris 2.10+

P5 Client

- Mac OS X: Intel 10.5+
- Windows: Server 2008R2, Server 2012, Windows 7, Windows 8
- Linux: Kernel 2.6.4+ (x86) 32bit or 64bit mit 32bit libs
For FS-Events with inotify, Kernel 2.6.13+ is required



- Sun Solaris: Sparc Solaris 2.8+, x86 Solaris 2.10+



On Solaris and Linux P5 requires the Unicode UTF-8 Language package is installed.

P5 is released for the following file server products:

- Helios EtherShare und PCShare 3.1+
- Xinet Fullpress 14.0+
- Netatalk 2
- Windows Services for Macintosh

P5 supports the following browsers:

- Safari 5.1+
- FireFox 3.5+
- Internet Explorer 9+
- Chrome 4+

P5 supports the following virtualizations:

- Parallels
- VMWare
- Linux-Xen
- Hyper-V

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. Disc Space

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



It is recommended to configure the program and index areas on backed up file systems, for example a mirrored disc or a RAID system. This has the advantage that a crash of the disc 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 disc space are recommended. The installation script suggests a directory with sufficient disc 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, so that a search can be performed later. The amount of disc 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 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.



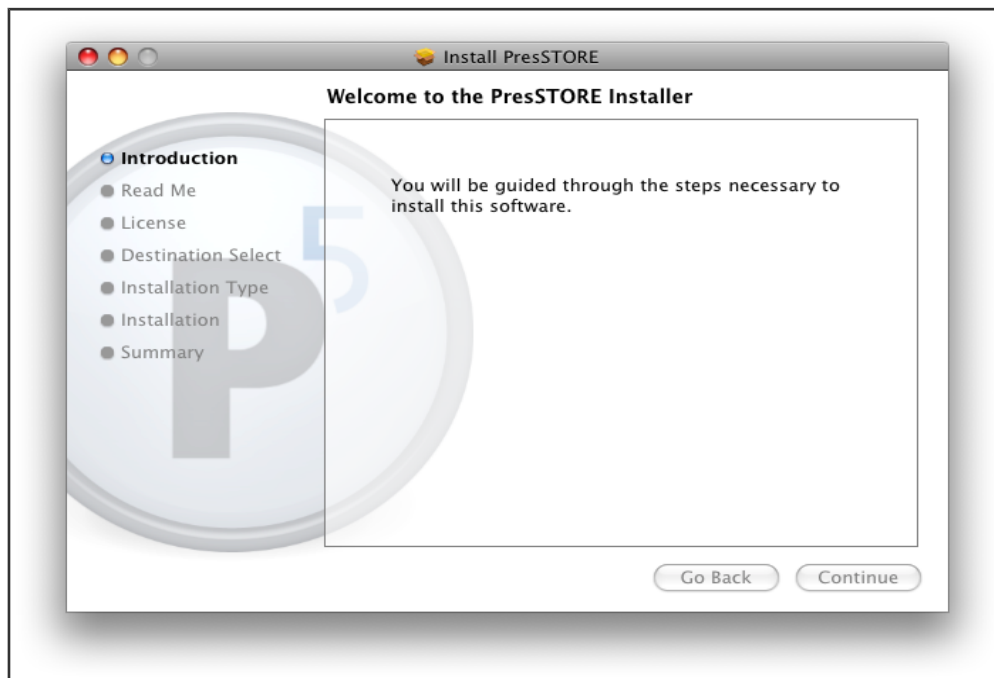
3.4 Installing on Mac OS X

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

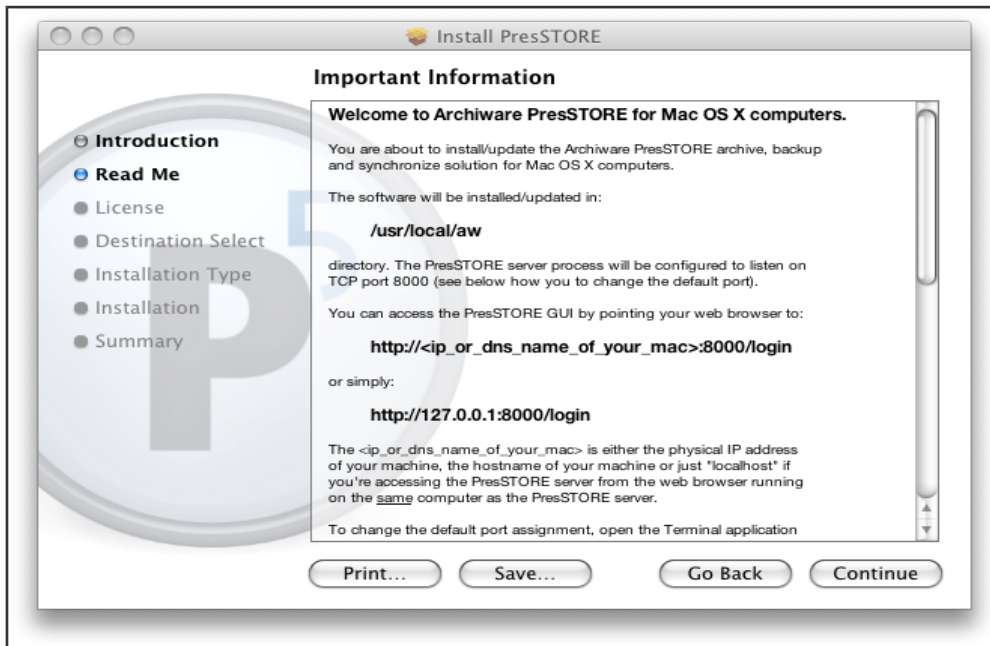
In case you are upgrading from an earlier version of PresSTORE, please see section 3.7, Upgrading to P5.

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

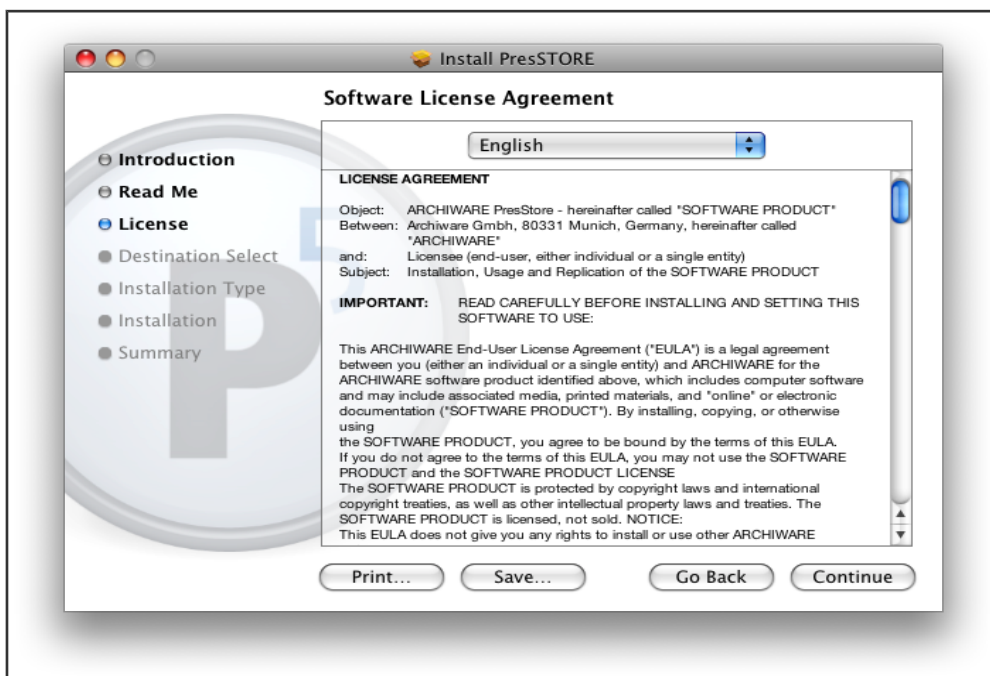
1. Begin installing on Mac OS X. You will be asked for your user name and password. Please enter and confirm with 'OK'.



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



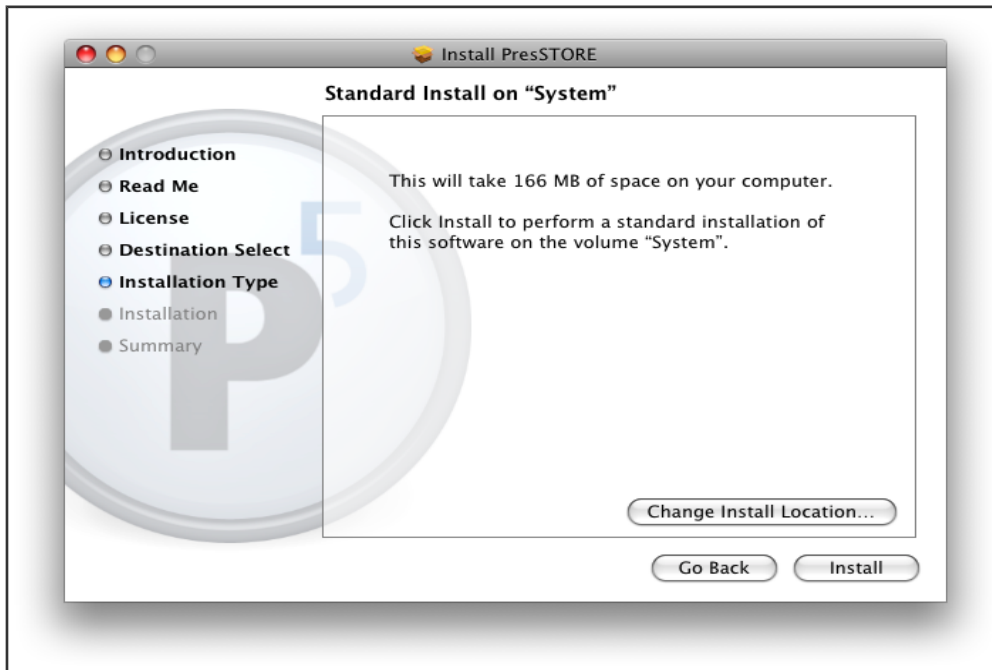
3. Please read through the License Agreement and familiarize yourself with the terms of use and of the application of P5. If you own a full or a time limited P5 license, click 'Continue' followed by 'Accept'.



4. Please choose the target volume where to install P5 and click *Continue*. Please note: A P5 server requires additional disc 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 just installed client to the client resources in the P5 server.

After completing the installation, a P5 icon is shown in the preference pane indicating that P5 is running:





3.5 Installing on Windows

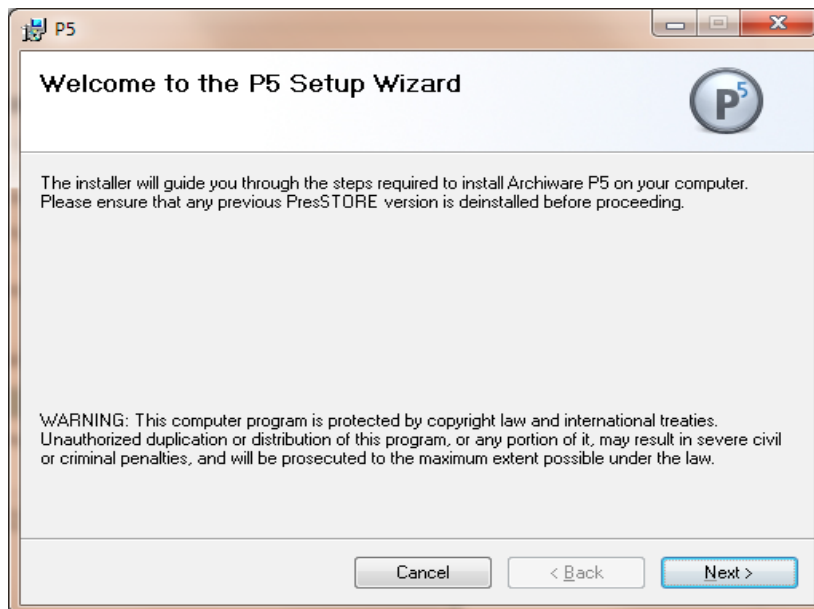
Archiware makes P5 available both on CD-ROM and online. The installation file is packed and has the format 'awpstxyz.msi'. The letters 'xyz' stand for the given PresSTORE version. The installation is done by the Windows Installer by simply double-clicking on the installation file.

In case you are upgrading from an earlier version of PresSTORE, please see section 3.7, Upgrading to P5.

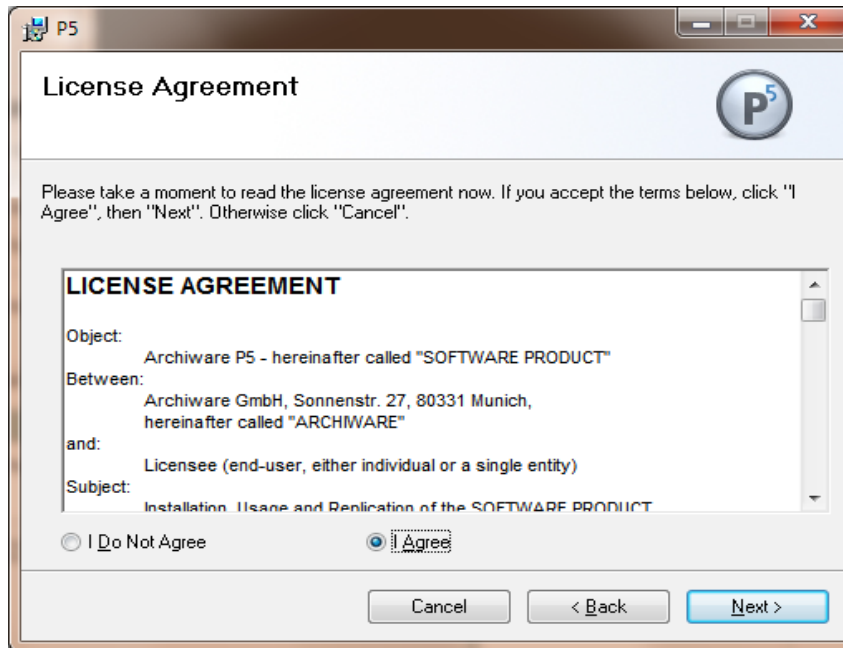


Administrator access rights are required to install P5.


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. 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.

5. 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 Solaris and Linux Installations

This section describes:

- installation requirements
- mounting the P5 CD
- 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 PresSTORE, please see section 3.7, Upgrading to P5.

3.6.1 Installation Requirements

P5 is available both on CD-ROM and online. The installation file is packed and has the format 'awpstxyz.taz' or '.tgz'. The letters 'xyz' stand for the given PresSTORE 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 Mounting the P5 CD

Normally the auto-mount process on both Solaris and Linux enable you to directly access the P5 CD. Should your CD for some reason not be mounted automatically, you can issue the following commands to mount it.

Linux:

```
Linux # mount /dev/cdrom
```

Solaris:

```
Solaris # mount -r -F hsfs /dev/dsk/c0t6d0s2 /cdrom
```

3.6.3 Creating an Installation Directory

Please create a directory for P5 taking into consideration the required disc 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
```

3.6.4 Copying the Installation File

Having created the P5 home directory, you can copy the installation file from the CD, 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 /cdrom/cdrom0/server/awpstxyz.tgz /usr/local/aw
```

Correct the source directory according to the mount point of your CD.

3.6.5 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 -
```

3.6.6 Starting the Installation Script

Now you can begin the installation of the P5 software. Start the installation script in the P5 home directory.

```
# ./install.sh
```

Please execute in sequence all the menu points and then exit the installation with point 99.



3.6.7 Special considerations when installing on 64-bit Linux System

P5 is a 32-bit application and requires 32-bit run time libraries. On 64 bit Linux distribution, these are included but often not automatically installed. Unfortunately the installation procedure depends on the Linux distribution. This section describes how to install the 32-bit libraries on some known Linux distributions:

SUSE Linux and OpenSUSE

SUSE Linux distributions come with YaST (Yet another Setup Tool). Please use yast to install the packages `libacl-32bit` and `pam-32bit`. YaST will automatically add all other required libraries.

Ubuntu

On Ubuntu, execute in a terminal window the following commands to call the Advanced Packaging Tool apt:

```
# apt-get install ia32-libs
# apt-get install lsb-core
```

RedHat Linux, CentOS and Fedora

RedHat, CentOS and Fedora come with the yum installer (Yellowdog Updater, Modified). Please call yum in a terminal to install these packages:

```
# yum install bzip2-libs.i686
# yum install elfutils-libelf.i686
# yum install glibc.i686
# yum install nss-softokn-freebl.i686
# yum install xz-libs.i686
# yum install zlib.i686
# yum install elfutils-libs.i686
# yum install libgcc.i686
# yum install libattr.i686
# yum install libacl.i686
# yum install audit-libs.i686
# yum install cracklib.i686
# yum install db4.i686
# yum install libseline.i686
# yum install pam.i686
```



Debian /GNU Linux

On Debian, execute in a terminal window the following command to call the Advanced Packaging Tool apt:

```
# apt-get install ia32-libs
```

Unfortunately, Debian Linux needs additional attention to adopt the pluggable authentication module PAM for use with 32 bit applications.

To do so, please follow these steps:

1. Check on <http://ftp.us.debian.org/debian/pool/main/p/pam> for the latest `libpam-modules*i386.deb` package.

2. Download that package to a temporary space and unpack it:

```
# wget http://ftp.us.debian.org/debian/pool/main/p/pam/libpam-modules\_1.1.3-7.1\_i386.deb
# mkdir pam
# dpkg -X libpam-modules_1.1.3-7.1_i386.deb pam
```

3. move the required libraries into the `/lib32` folder:

```
# mv pam/lib/i386-linux-gnu/security/* /lib32/security/.
# rm -r pam
```

4. You will need to adopt the installed pam configuration of P5. Please edit the file `/etc/pam.d/nsd` and change these lines:

```
auth required      pam_unix.so
account required   pam_unix.so
```

to

```
auth required      /lib32/security/pam_unix.so
account required   /lib32/security/pam_unix.so
```

That's all, after that step login must be possible in PresSTORE, too.



3.6.8 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)
```

3.7 Upgrading to P5

This section describes how to upgrade the P5 Server and Client on a Mac OS X, Windows, Linux or Sun Solaris.

3.7.1 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.
- PresSTORE versions prior to PresSTORE 4.4 cannot be upgraded. Please either upgrade to version 4.4 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.
- A backup of the system is recommended before starting the installation. This backup should include both the PresSTORE 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.
- A downgrade from P5 to a previous version is not possible without losing the data converted to or written with P5
- When Upgrading from P4, please read also section 3.7.3 Upgrade from P4 to P5 below.
- When using Backup2GO, please update the workstations directly. The remote update mechanism of these workstations cannot handle remote updates. This does not affect P5 clients as used in Archive Backup or Synchronize.



3.7.2 Program Update

a. Mac OS X:

Upgrading a Mac OS X system requires the same procedure as a new installation. See section 3.4, Installing on Mac OS X. The installation procedure retains the existing configuration and index information.

b. Windows:

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

If PresSTORE version 3.x was previously installed, it has to be uninstalled first. Please read further instructions in the README file of the corresponding PresSTORE version. Uninstalling does not remove the configuration nor the index files. It is simply required, because PresSTORE 4 uses a new installation program.

c. Linux/Solaris:

Please read the README file in the ARCHIWARE home directory for an upgrade on a Solaris or Linux system.

You have to enter the appropriate PresSTORE 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.

After upgrading, P5 will 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.



3.7.3 Upgrade from P4 to P5

A major improvement in P5 is the new indexing system. The new, more powerful indexes are based on SQLite and can scale up to accommodate even for the largest of today's systems.

The new indexes come with a few changes that affect updates from P4:

Backup indexes

After updating to P5, backup jobs will create new empty indexes leaving the old indexes intact and browsable in the *Restore* section of the GUI. New and empty indexes imply that the first run of any backup job after the update will always be a full backup. Files saved by previous backups will remain available in the P4 index structure to be restored until they fall out of the retention time. P4 indexes will not be used for further indexing by any P5 job. Furthermore, P4 indexes can only be browsed, whereas P5 can additionally also be searched.

A major change in the backup module affects how the backup job determines what to save for the next increment. In P4, the job referred strictly to the latest backup of the backup plan that was being executed. In P5, the job refers to the latest files saved by any previous job to the referenced media-pool. More specifically, each incremental backup updates the state of the saved files and folders as saved by the latest backup that used that specific pool of volumes.

This implies that:

- when two or more backup plans use the same media pool(s), the jobs will complement each other to actualize the state of that pool. Unlike in P4, incremental backups of different plans are not independent of each other. Such setups must be checked after the upgrade. If two plans are to independently run increments, each must have a separate pool. In case a backup plan exclusively uses certain pools and no other plan uses the same, there is no change in the behavior as compared to P4.
- when a backup plan uses alternating media pools in a single incremental event, a previous full backup in one pool will no longer enforce a full backup in the other pools when they get addressed.
- unlike in P4, it is now possible to create a full backup in one plan and an incremental backup in another plan that both use and refer to the same media pool.

In case backup must be deferred in order to adjust the plans to accommodate the changes, the corresponding plans can be disabled.

Archive indexes

Same as with the backup indexes, P4 archive indexes remain accessible in the *Restore* section of the GUI and can be used for the immediate restore of data. However, in order to be able to run new archive jobs that add new files to the archives, the indexes must be first converted to the new P5 database format. This conversion is not done automatically, but has to be triggered manually.



In order to convert the index, navigate to the index databases as described in section 6.14, Index Databases, select the archive index and click the **Convert** button.

During conversion the index is read from the old location in the P4 configuration subfolder `conf/index/archive` and written to the new location at `config/index/archive`. In case the index subfolder was moved to some other location and referred to it by a link in the `conf` subfolder, it is recommended to create a similar link in the `config` subfolder.

For instance, if the existing link is

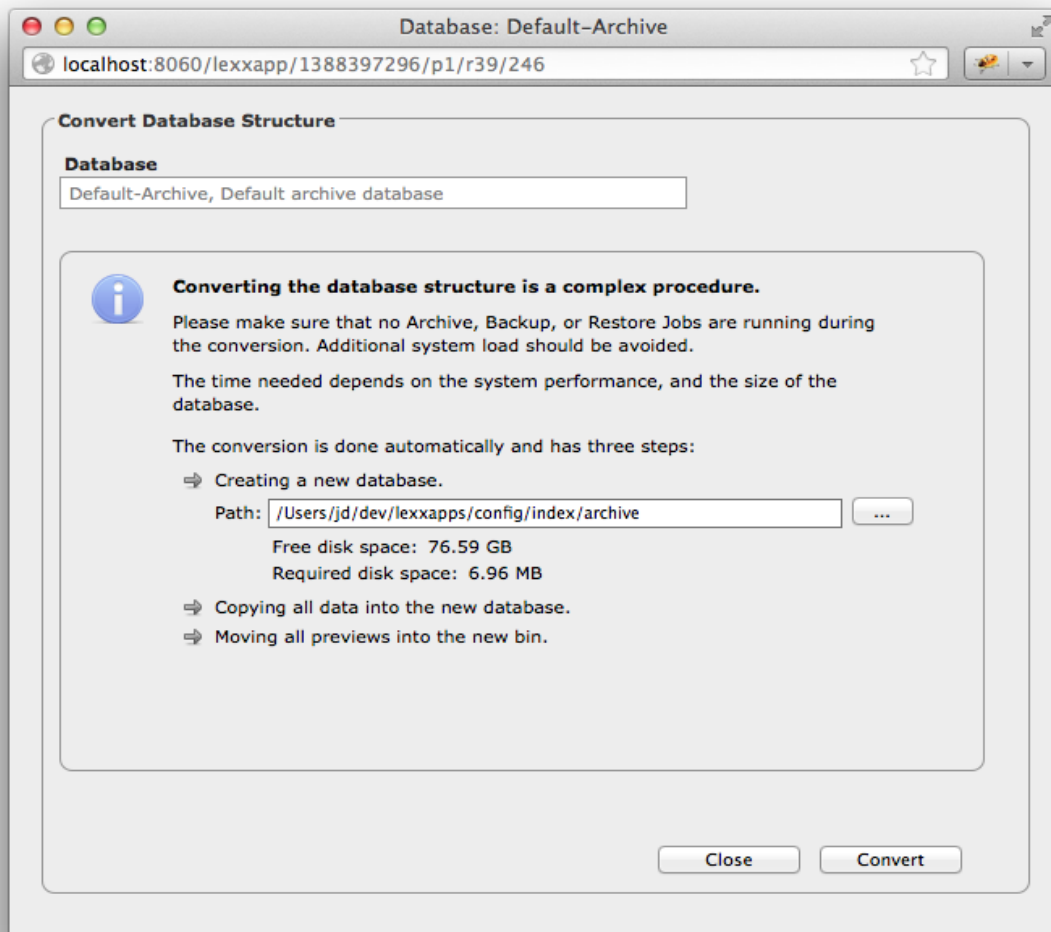
```
conf/index → /Volumes/bigdisk/aw/index
```

then a second link should be created onto the same disk but into another folder, for example:

```
config/index → Volumes/bigdisk/aw/p5index
```

Do not link both of the folders to the same target folder. It would cause the index conversion to fail.

The conversion is then started:





Depending on the index size, the conversion can take between a few minutes and multiple hours. Since the conversion creates high CPU and disk I/O load, please run only one index conversion at a time.

Also make sure that no other jobs access the same index during conversion. The progress of the conversion can be monitored in the P5 job monitor.



4 Signing In and Entering the License

This section describes the basic elements of P5 and of the 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 4.0+ or Internet Explorer 9.0+.

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

<http://localhost:8000/login>

Please use *root* or *administrator* for your first session.

Once you have signed in 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 **Getting Started > License the software** in the main window.
3. Start entering a new license by clicking 'new' 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 [Archiware portal](#).

The screenshot shows a 'License Configuration' dialog box with the following fields and values:

- License Type:** PresSTORE 4 and 3
- Serial number:** 16090198
- Status:** Enabled
- Scope of License:**
 - Selection:** AWx511 - Workgroup Edition 24
 - Expires:** 19-Nov-2011 (Interim license)
 - Maintenance expires:** 28-Feb-2012
- System Identification:**
 - Host ID:** c8a16796-0401
 - License key:** tpfe-khrt-edbs-wthd

Buttons at the bottom: Close, Apply

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



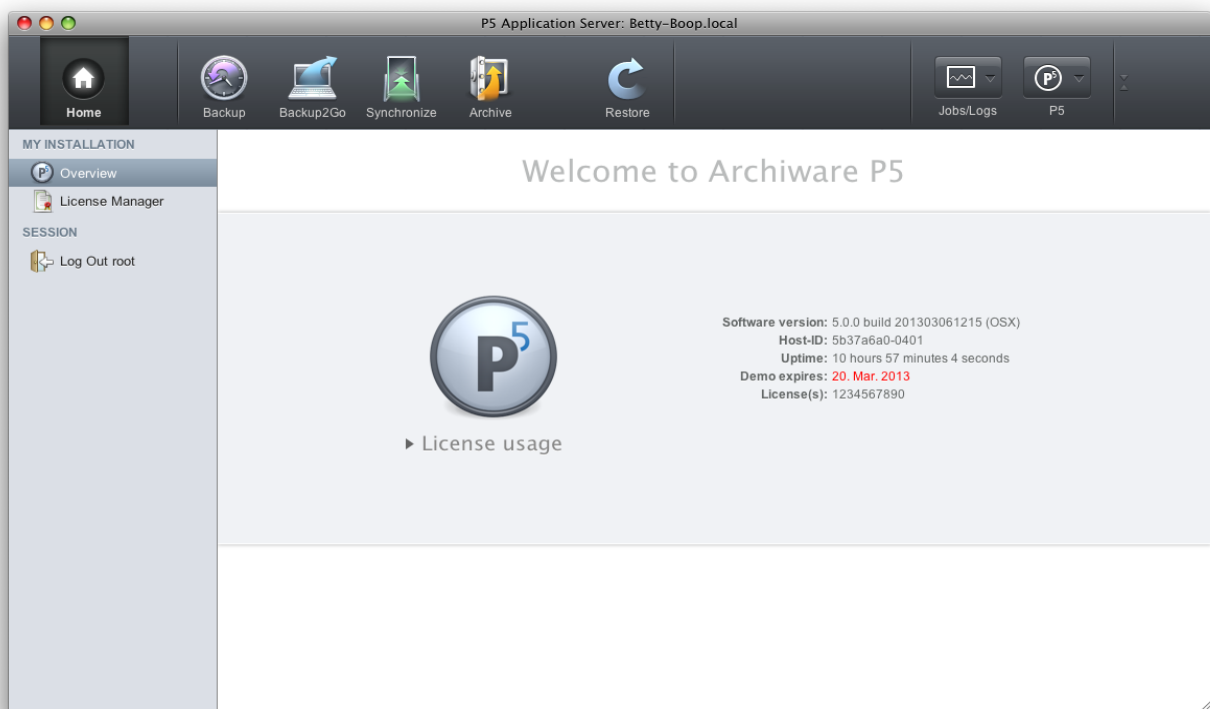
4.3 A General Introduction to the Server's User Interface

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.

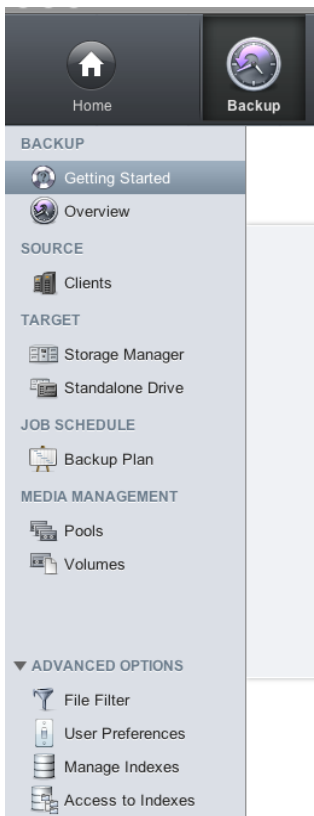


P5 – leads to help, support and further program information.

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, which can otherwise be accessed via the tree structure from the home screen.



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 rarely used configuration like the index database or the log-in areas.

Each module stores its current selection within the side bar. This allows to switch between modules and yet return to where you were positioned in that 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.

Software version: 5.0.0 build 201303061215 (OSX)
 Host-ID: 5b37a6a0-0401
 Uptime: 10 hours 47 minutes 50 seconds
 Demo expires: 20. Mar. 2013
 License(s): 1234567890

P5 Service/Expansion	Use	Licensed	In use	Available
Backup	Backup	Yes	Yes	Yes
Backup2Go	Workstation Backup	Yes	No	Yes
Synchronize	Synchronize	Yes	Yes	Yes
Archive	Archive	Yes	Yes	Yes
Media and storage management	Tape library, disk storage	Unlimited	25	Unlimited
Media drive	Standalone tape, removable disk drive	Unlimited	2	Unlimited
Server agent	Server client	Unlimited	1	Unlimited
Virtual server agent	Virtual server client	Unlimited	0	Unlimited
Workstation agent	Workstation client	Unlimited	0	Unlimited


The display shows the installed software version and build, the host ID identifying the host where P5 is installed on 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 setup for instance encryption, see section 6.12, P5 Clients,
- to **User Preferences** to setup login permissions, see section 6.16, User Preferences, and
- on Backup2Go Workstations to setup the Backup2Go Server, see section 7, Configuring the P5 Backup2Go Workstation

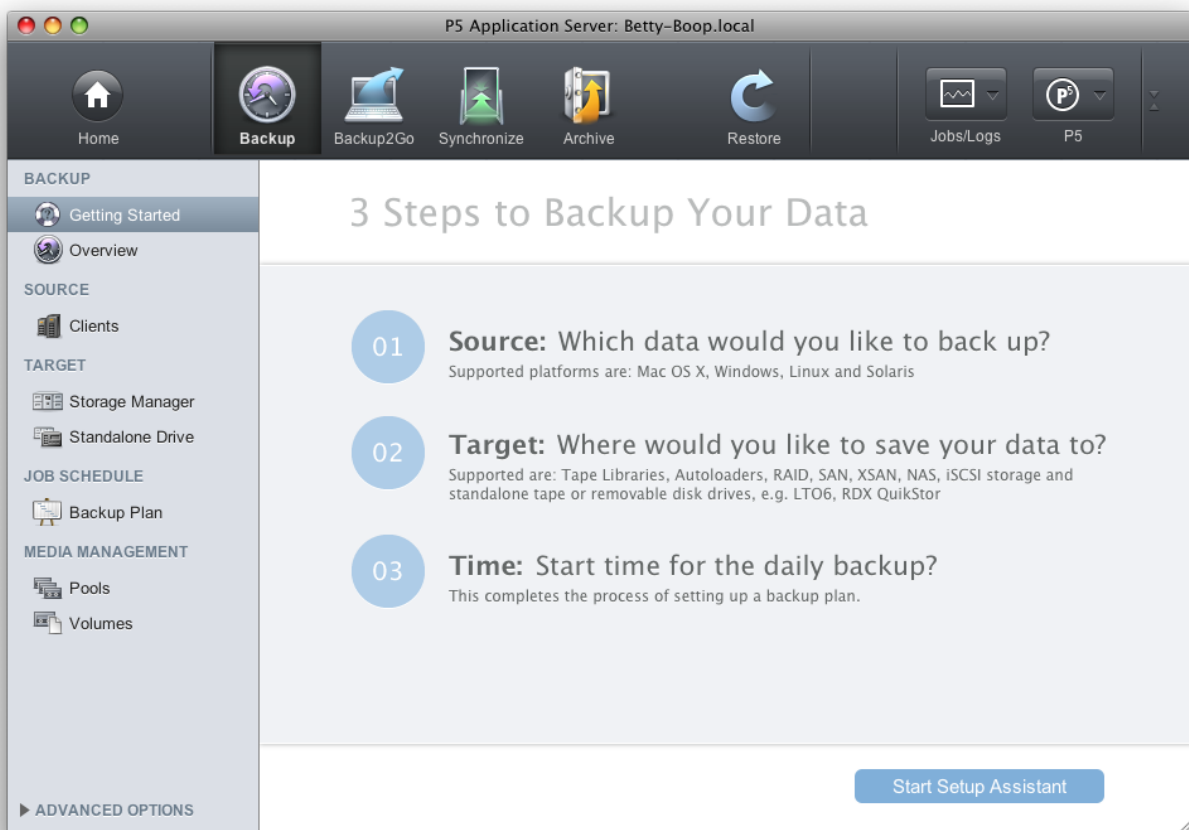


5 First Steps With P5

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.

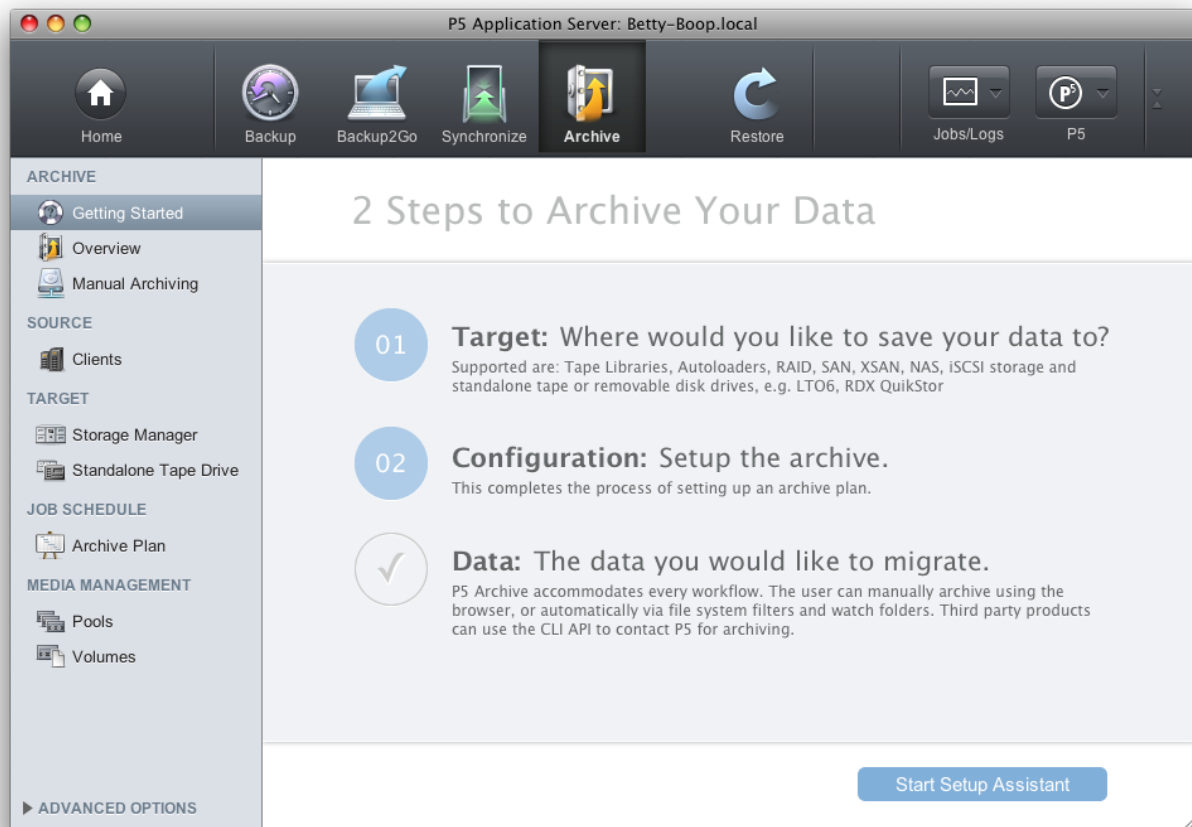
The Backup Assistant will define a Backup Plan (see more in section 6.7, Backup Plans) that holds the info of what to save and the time schedule. Furthermore, a virtual or physical media changer or single drive is configured automatically and a media pool, either Disk-Backup or Tape-Backup, is created and Volumes can be labeled in preparation of writing to them. See more about target storage in sections 6.2 to 6.4.

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 6.12.1, 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 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 6.4.

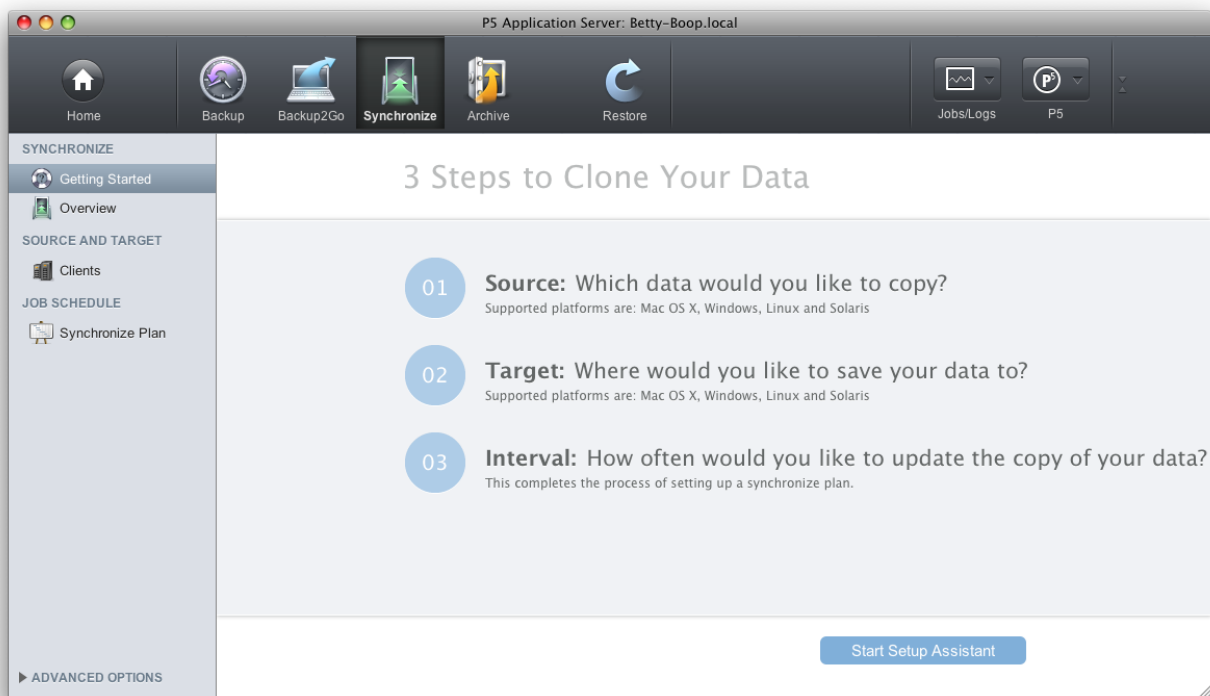
The Assistant also allows the user to define a client in order to archive data from that client instead of the local machine. This configuration requires a P5 installation on that client computer. For further information about P5 clients please see section 6.12.1, Creating a Client.

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 6.6, 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 6.8, Synchronize 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 6.12.1, Creating a Client.



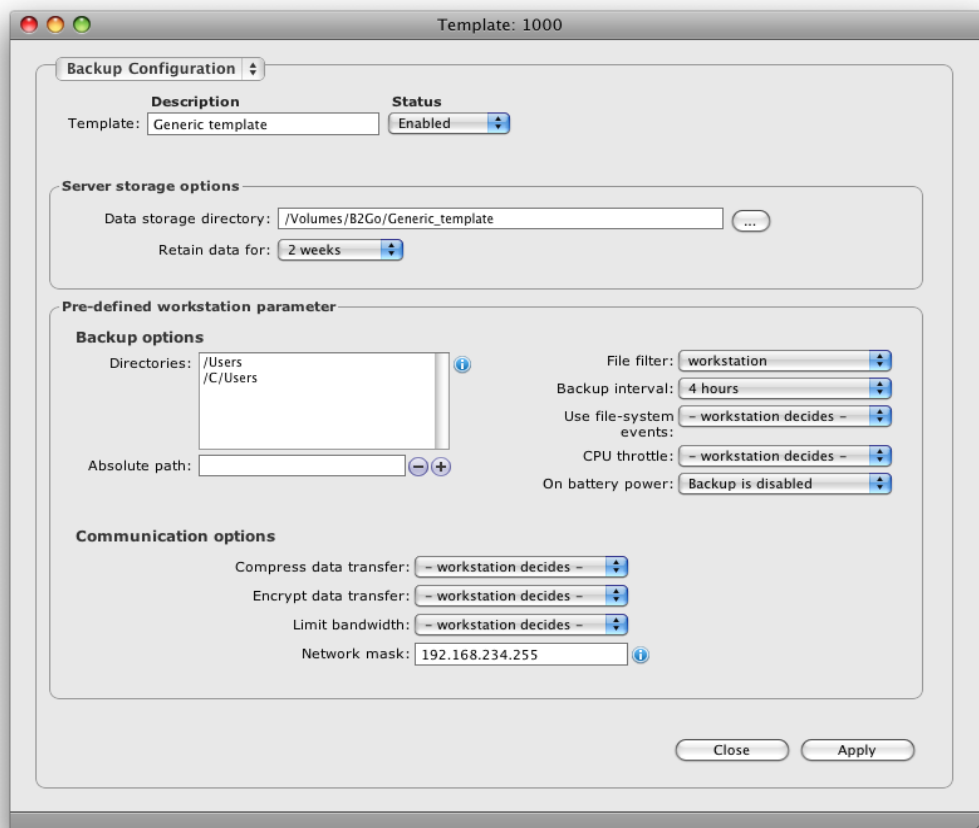
5.4 How to Set Up Backup2Go

Backup2Go is an easy to use backup to disk module designed to back up workstations that are not permanently available for backup.

5.4.1 Setup the Backup2Go Server

Select Backup2Go from the top menu.

To prepare the Backup2Go server, choose *Workstation Groups* from the left sidebar. Then double click *Generic template ...* in the main window.



For each new workstation that registers into the server, a new resource is created and linked to this template. The template defines the default parameters that will be used for backing up the workstation. If you wish to set parameters that differ from those defined in the template, you can set those individually in the workstation's resource on the server or on the workstation itself. Any subsequent changes to the template will be reflected on the workstations linked to it.



The purpose of the *Predefined workstation parameters* section is to define the parameters of the backup, which the administrator of the backup does not wish the user to be able to change. Entries that are left blank can be defined directly in the workstation's setup. The ones defined here will override and prevent setting the parameters on the workstation. When the directories field is left empty, as by default, the home directory of the user that signed in to configure the local settings on the workstation is used. The *Server storage options* section defines where and for how long to store the data of a previous snapshot. The current state of the file system is indefinitely stored. These entries are mandatory. If not set in the template, these parameters must be set later on the workstation.



P5 supports ZFS (Zetabyte File System) on Solaris. On other file systems like Journaled HFS Plus on Mac OS X 10.5+ or ext3/ext4 on Linux, soft- and hardlinks are used to create snapshots. Please note that normal filesystem tools like the OSX finder may report incorrect values for the used disk space.

This step completes the Backup2Go server setup.

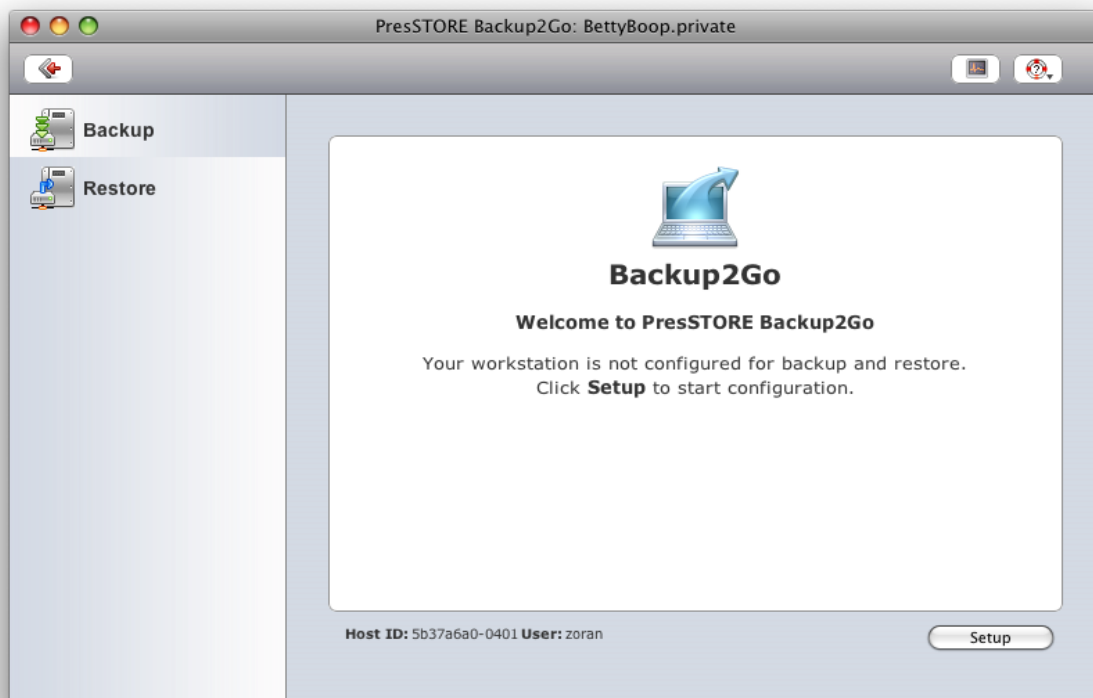


5.4.2 How to Set Up a Backup2Go Workstation

Install P5 on each workstation to be backed up with Backup2Go. The procedure is the same as for the server, except that on the workstation no license and no further configuration is required, see more in section 3, Installation.

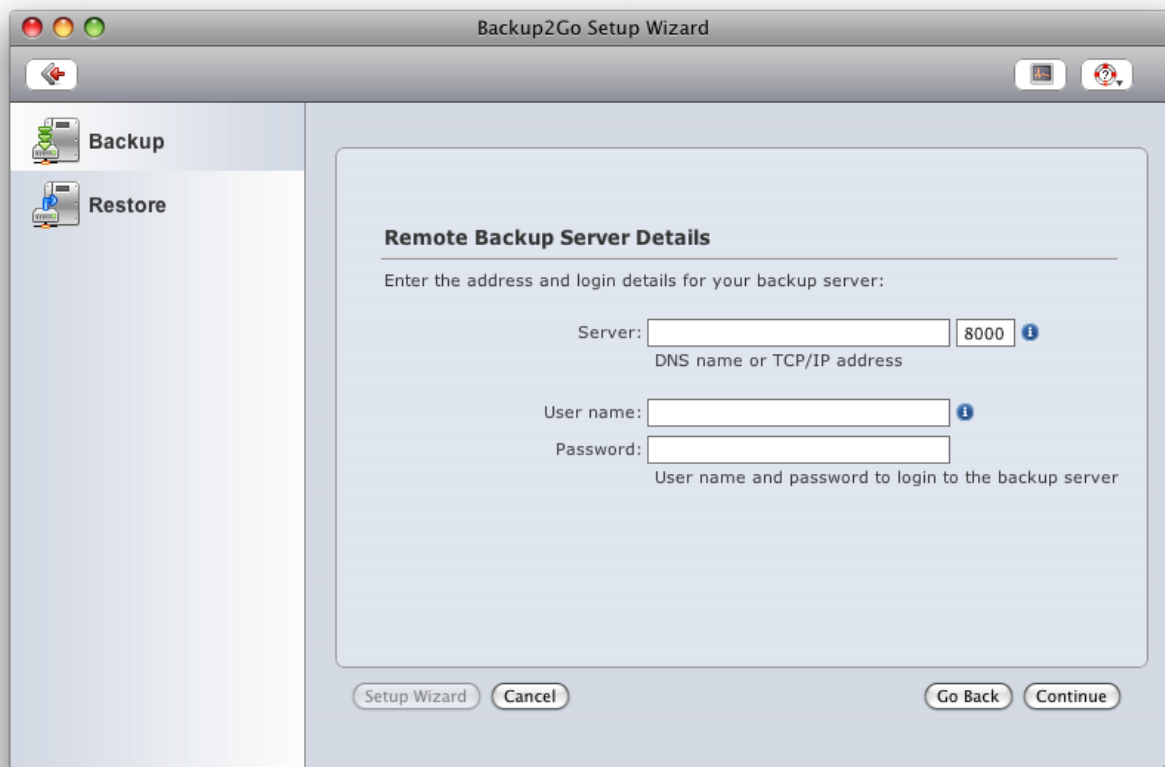
After installing, please use a web browser (for example Safari or Firefox) and the following URL to log into the P5 user interface:

<http://localhost:8000/workstation>





Click **Setup** to configure the workstation backup, which will take you to the workstation setup wizard:



To setup the required connection to the server, enter the **server name** or IP address of an existing Backup2Go server and enter the **User name** and **Password** to be used to connect to the server. This account must have log-in access to the Backup2Go server. The same server account is required later during restore to access the data. Click **Continue** to confirm the input.

You can define the directories that you would like to back up, provided that this field is not overridden by the administrator, who can explicitly define the directories to backup in the workstation template. By default, this field is filled with the home directory of the user that signed in to configure the workstation. To configure it, follow the wizard to the **Data to backup** section to define or modify the directories you want to include in the backup.

Check **Activate backup** and confirm the changes by clicking **Apply**.

The template on the server defines how often a workstation backup will be done. The backup for the workstation is run for the first time after the first such interval has expired. In case the workstation is offline or in standby mode at the scheduled time, the backup will start as soon as the Backup2Go server becomes accessible again.

Please find more information about the Backup2Go server in section 6.13 and the Backup2Go workstation in section 7.



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 2.5, The 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 intermediate 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**.

License Configuration: 16090198

License Type: PresSTORE 4 and 3

License	Status
Serial number: 16090198	Enabled

Scope of License

Selection: AWx511 – Workgroup Edition 24

Expires: 19-Nov-2011 (Interim license)

Maintenance expires: 28-Feb-2012 Format: dd-mmm-yyyy

System Identification

Host ID: c8a16796-0401

License key: tpfe-khrt-edbs-wthd

Buttons: Close, Apply

Serial number:

Please enter the serial number from your license paper. Each serial number is unique and is validated.

Status:

The serial number is automatically filled in when selecting a trial license. The status flag shows if this license and its resources, is 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.
Host ID:	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. 6. In case the host ID changes due to hardware modifications, please request a host key transfer from the Archiware license portal.
Selection:	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.
Version:	The version of the license – It does not have to match the installed P5 version. Enter the license version, not the PresSTORE version into this field.
Expires:	Applies only to the trial license. Shows when the license expires.
Maintenance Expires:	Shows when the maintenance expires. This date has no influence on the installed P5, but updates released after that date will be unavailable.
License Key:	Archiware confirms your registration by sending your License Key and Registration paper.

6.1.2 Importing a License

Starting with PresSTORE 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 Devices

Backing up means copying data from on-line storage, for example a RAID system, to one or more off-line storage units. Off-line storage are for example LTO tapes, which the user usually cannot access 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 backup data to. Each device carries several attributes, which describe the properties of the drive. Drives in changers also have 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.

6.2.1 Creating a Stand-Alone Tape Device

You can configure a device either manually or 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 (see note). Then you will have to configure the device manually.

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
Standalone Tape Drive

On Solaris, the 'sgen' driver scans the SCSI bus addresses 0 to 15, LUNs 0 and 1 for sequential devices. Should a device be configured on a higher LUN, you must extend the driver configuration file '/kernel/drv/sgen.conf' with the corresponding LUN entry. The following setup line in '/kernel/drv/sgen.conf' tells the driver to look for a SCSI device at address 120, LUN 2.

```
name='sgen' class='scsi' target=120 lun=2;
```



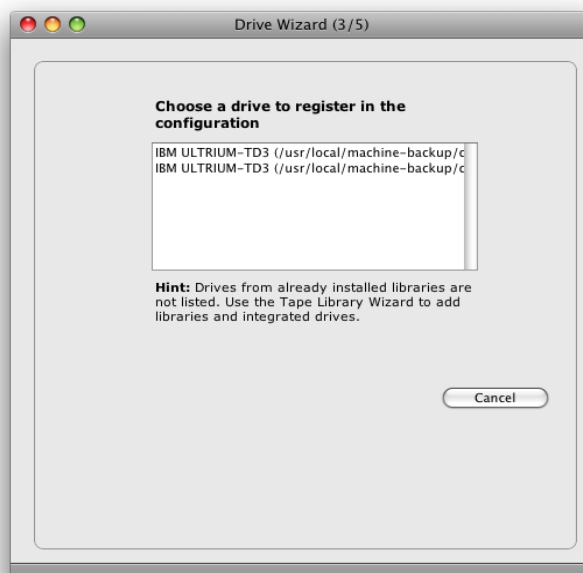
3. Creating a Device Using the Device Wizard.

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

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



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





- c. Following the scan, all the devices found will be listed. Devices that are already configured as part of media libraries, will be left out. See section 6.3.1, Creating a Jukebox. Select a device and click **Next** to confirm.



- d. 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.



4. Creating a Device Manually.

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

- a. Enter the required information.

General Setup

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 store 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'.

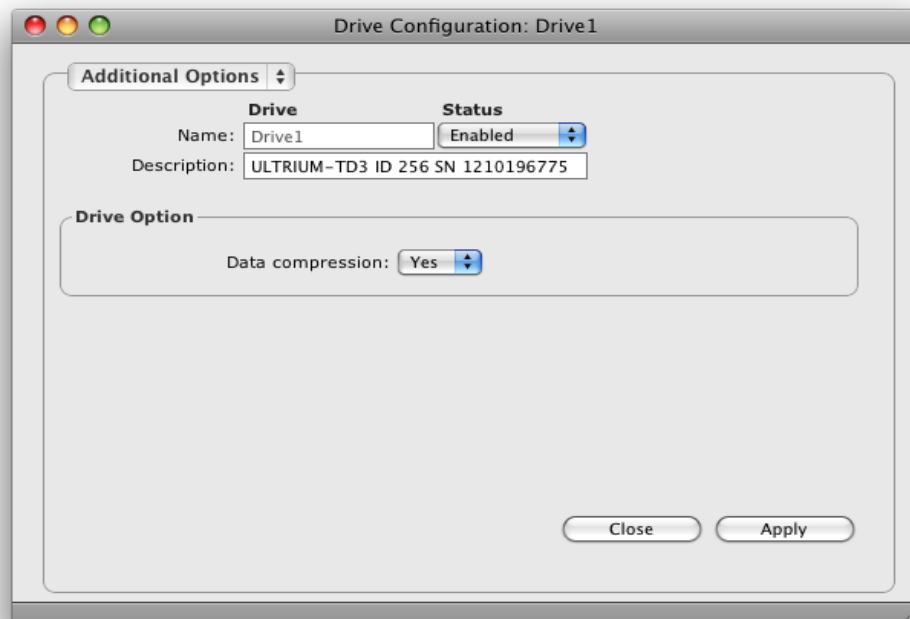
Cleaning required:

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. Drives in jukeboxes can be cleaned automatically as described in section 6.3.1, Creating a Jukebox .

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.



6.2.2 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 are written like single tape drives as one data storage. This method is intended for removable disks, not for bigger fixed disk arrays.

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 **Standalone Tape Drive**
3. Click **New Removable Disk**



4. 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.

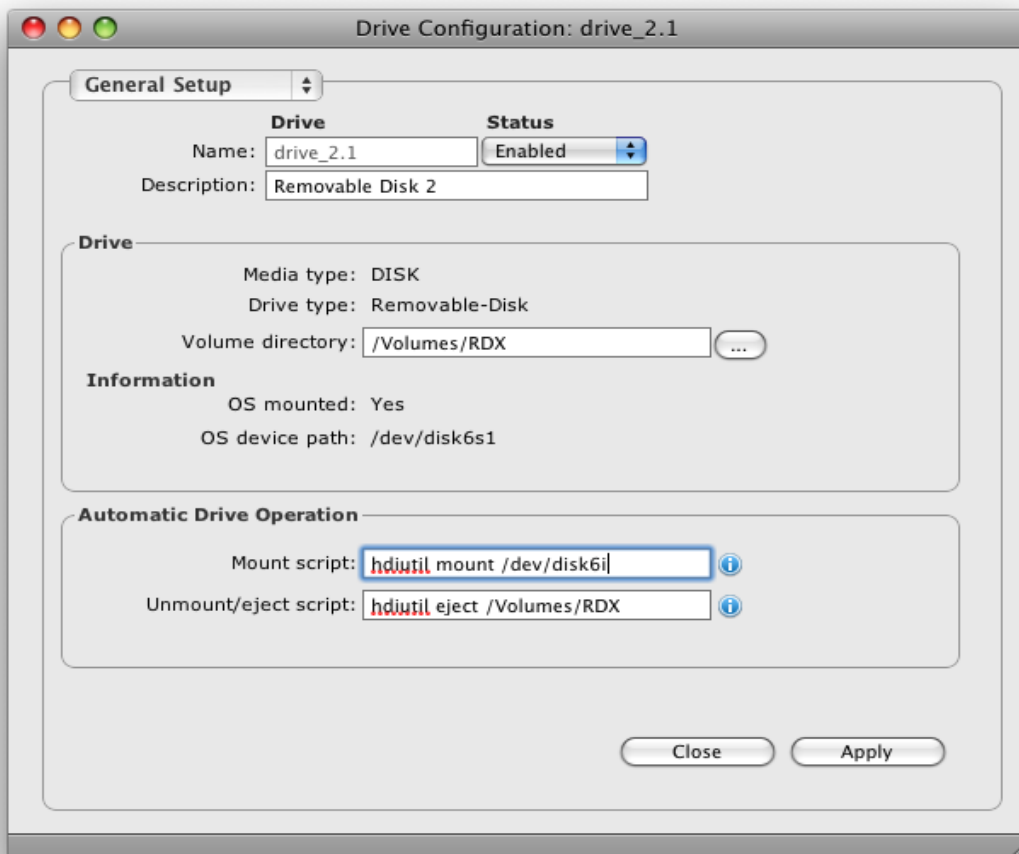


You may select "do not label" to skip that step.



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. P5 is capable in this setup to handle disks with different sizes 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.

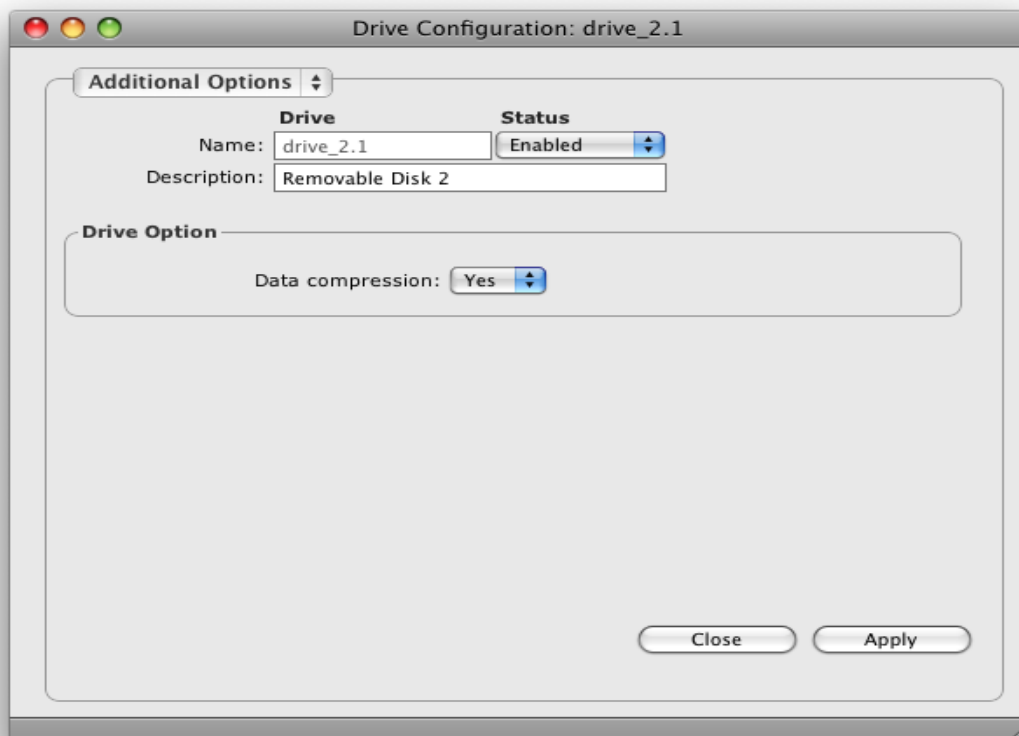


Unmount script: The unmount script can be used to unmount the disk after use. Devices supporting a mechanical eject command (for instance RDX drives) can be ejected with this script.



Unfortunately, the device node used in mount commands mounted at the given path (on Window: 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 practise to keep the assignment of device file and drive fix is to connect external devices in the same order.



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



6.2.3 Devices: Base Buttons

For each configured device, there are several options you can access with the buttons in the base bar. All the options are available in the *Gear Menu*:

- | | |
|----------------------------|---|
| <i>Delete</i> | Deletes the selected drive setup. |
| <i>Edit</i> | Edits the selected drive setup. |
| <i>Label Volume . . .</i> | Labels a volume for use with P5, see section 10.1, Labeling Media. |
| <i>Mount/Unmount</i> | Mounts or unmounts a volume, see section 10.2, Mounting Volumes. |
| <i>Drive Reset . . .</i> | Forces the drive to reset itself.
Before initiating a reset, you should check that no job is using the drive, since the job will be interrupted and data may be lost. |
| <i>Import Volume . . .</i> | Imports a P5 volume into your database that has been written using P5 at some other location, for example from a subsidiary. A separate index is created for imported volumes from other hosts. |



6.3 The Media Library (Changer or Jukebox)

Media libraries enable fully automated data save and restore operations without manual intervention. For large data volumes, a single media 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.

6.3.1 Creating a Jukebox

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.


1. Log into P5 using an Internet browser as an admin user.

See section 4.1, Starting the Browser.

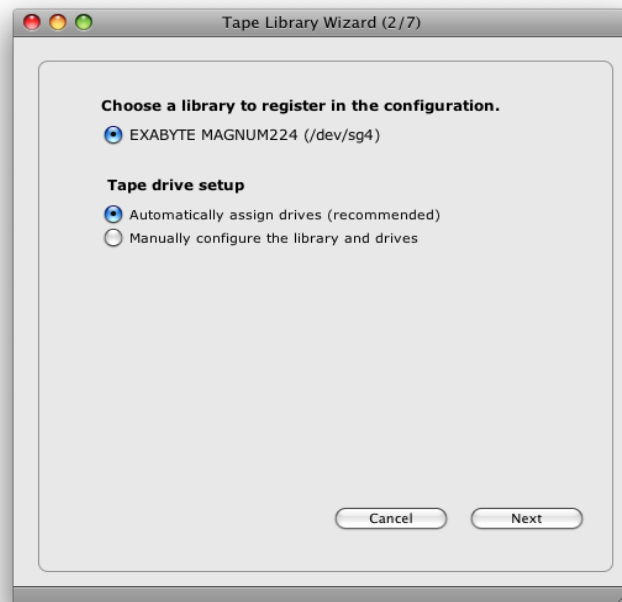
2. Select **Archive** or **Backup** from the top menu and then **Storage Manager**

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

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



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

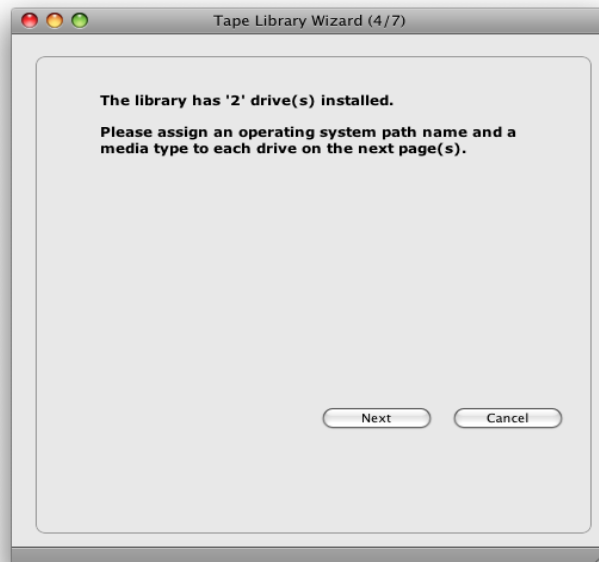


- c. The wizard assigns a name for the library and suggests a description. You may modify the description and/or confirm by clicking **Next**.

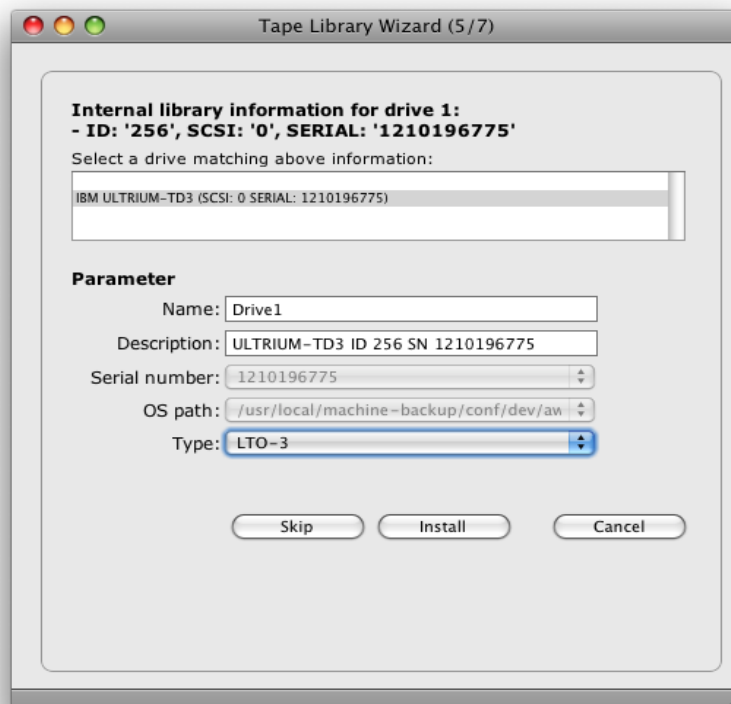
The Jukebox Name it is automatically assigned by the wizard and cannot be changed. The description can be changed later by editing the Jukebox definition.



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



- e. Now you will be then prompted to configure and create each of the drives. Complete or modify the information and click **Install** to confirm.





Should P5 not use or be aware of one of the drives, you can omit installing it by clicking **Skip** and continue with installing the next one.

Name: The drives resource name. May be changed

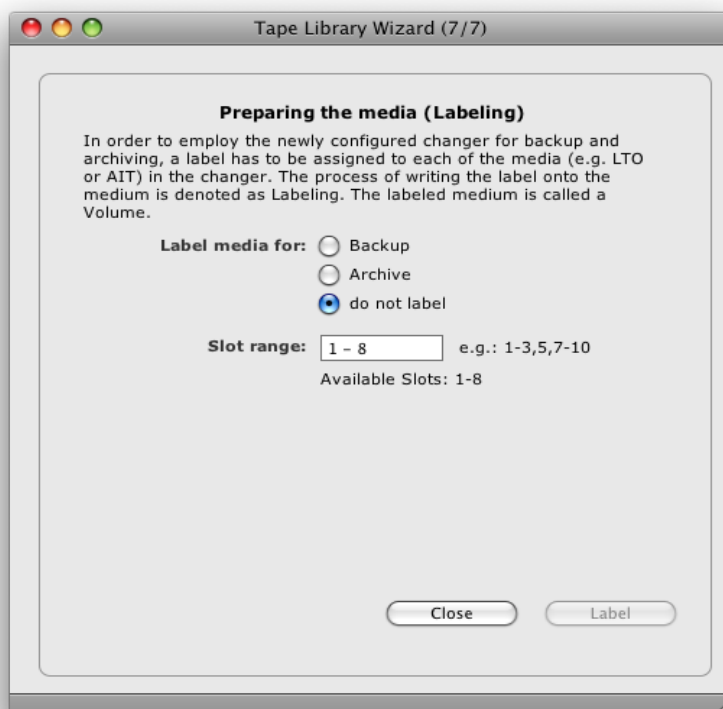
Description: Describes the drive. Changeable

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 jukebox installation will be saved by clicking **Cancel**, following the installation of the last drive.

2. The data browser will display the new jukebox if the configuration of the library completed successfully. You may then close the wizard window.
3. Following the jukebox setup, the wizard offers the possibility to label tapes:

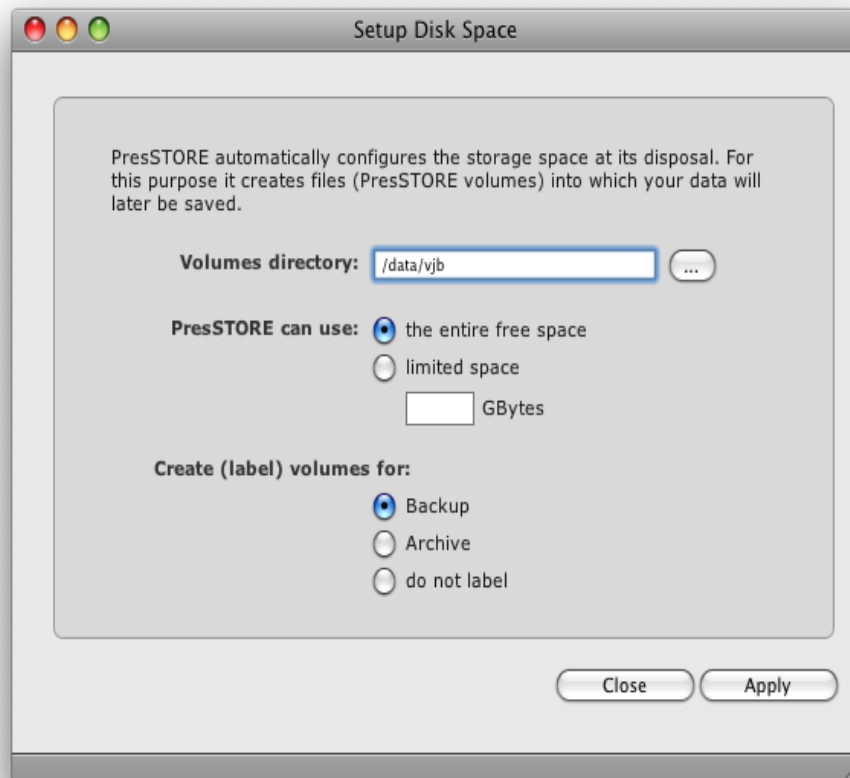




6.3.2 Create a Virtual Jukebox

The virtual jukebox handles disk files like tapes of a physical library, thus making possible 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 drives associated with the changer are also automatically configured by the wizard.

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 *Storage Manager*
3. Choose the option *New Disk Storage...* from the bottom bar
4. The wizard will display the virtual jukebox configuration.



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.

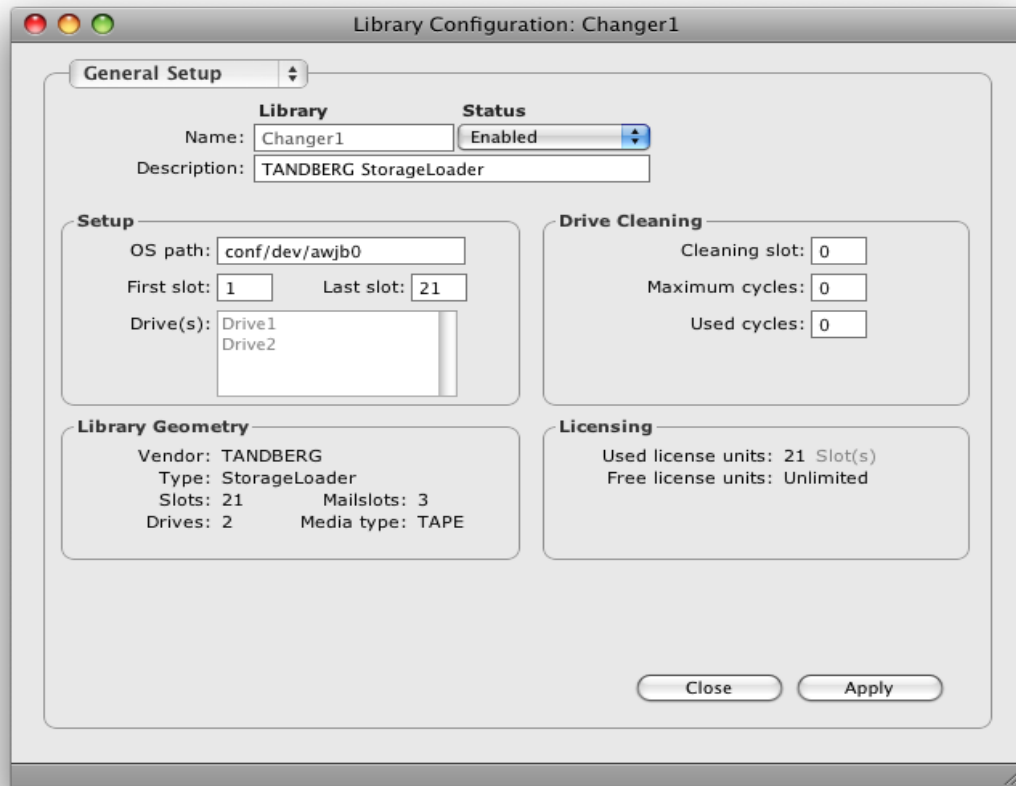
In order to resize a virtual jukebox without losing data:

- 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.

6.3.3 Modifying a Jukebox Setup

The jukebox describes for the most part its physical properties. Should the configuration of the jukebox have changed, for example a new drive or more slots got added, it is best to delete the existing entry and install a new one. This way, all the corresponding device entries will be correctly actualized and referred to.

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 **Storage Manager**
3. Double click the jukebox you wish to modify or mark its line and choose the **Edit** option from the **Gear menu**.
4. Change the desired fields using the jukebox manager dialog.



Jukebox name: The name of the jukebox. It is assigned during the initial configuration and cannot be modified.

Jukebox Status: The status flag defines the availability of this changer. The states are:
 » Enabled: The changer is available for use by the system.
 » Disabled: The changer will not be used by the system.

OS path: The operating system path to the device's inode. For example '/dev/rmt/awjb1 '

Slot Range: If P5 should use only a certain range of slots, it can be constricted here by entering the first and the last slot of this range. **Caution:** you cannot reconfigure the jukebox to extend the number of slots here.

Media Drives: This is a list of one or more devices (drives) that are built in the changer and may be used by the system. If a device should be added or removed, you must delete the library and configure it anew using the wizard.



- Cleaning slot:** P5 supports automatic drive cleaning. This field carries the slot number where the cleaning media can be found.
- Max cln cycles:** This defines the maximal number of cleaning cycles per cleaning media.
- Used cln cycles:** This field is incremented by one each time the cleaning media is used, until the "Max cln cycles" are reached. When the limit is reached, you must replace the cleaning media and reset this field.

Additional Options:



Change only if Archiware recommends it.

Library Configuration: Changer1

Additional Options ▾

	Drive	Status
Name:	Changer1	Enabled ▾
Description:	TANDBERG StorageLoader	

Hardware Options

Note: Change only if recommended

Eject medium before transport: Yes ▾

Force resetting all drives: Yes ▾

Take inventory on failed volume requests: Use Barcode ▾

Close Apply



<i>Eject media before transport:</i>	Some compact libraries eject the media from the drive automatically before transporting it, without requiring additional SCSI commands to do so. These devices may require this switch to be set to 'No' to function correctly.
<i>Force resetting all drives:</i>	When this option is set, P5 attempts to un-mount all the drives of a library during reset, even though the jukebox inventory does not report the drive to be mounted. This option is useful for libraries, that falsely report the drives to be empty in their inventories, although the drives are actually mounted.
<i>Take inventory on failed volume requests:</i>	When P5 detects that its view of the library's inventory does not correspond to the actual inventory of the media within it, P5 can attempt to automatically correct its inventory. This can be done either by comparing the bar codes to identify the volumes or by reading the label off the volume, which requires mounting the volume into a drive. The latter is necessary for libraries not equipped with a bar code reader.

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

6.3.4 Storage Manager: Gear Menu

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

6.3.4.1 Label Volumes...

For a description read section 10.1, Labeling Media.

6.3.4.2 Take Library Inventory...

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 un-mount the volume.

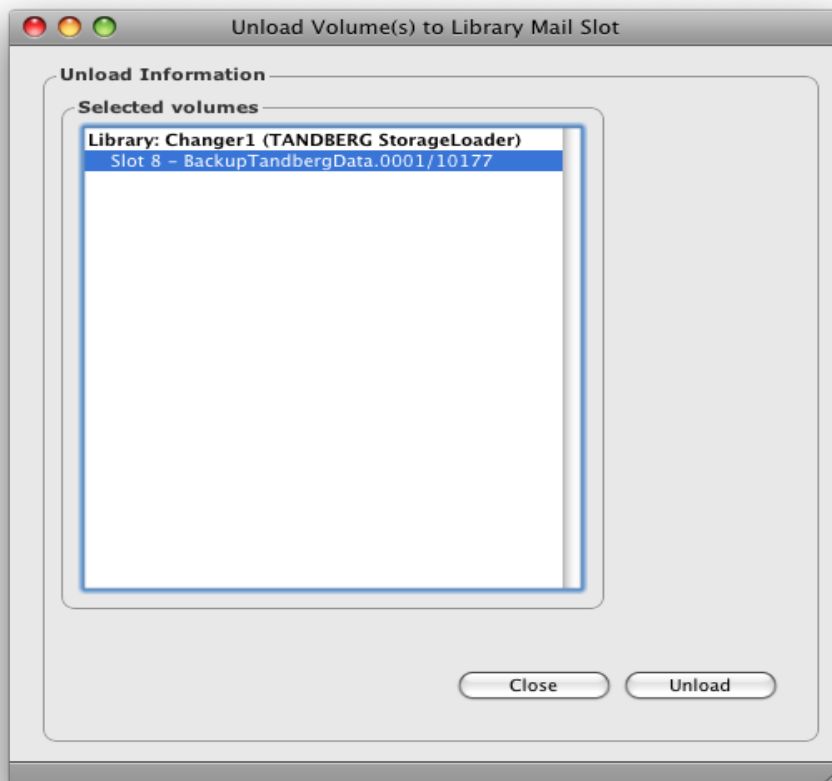
If you remove media from the library or replace it for off-site storage, or the exchange of cycles, you have to run an inventory for the library in order to make the changes known to P5. By default, all the slots of the jukebox are listed in the slot selection field of the inventory window. You can take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

6.3.4.3 Load/Unload Slot...

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.



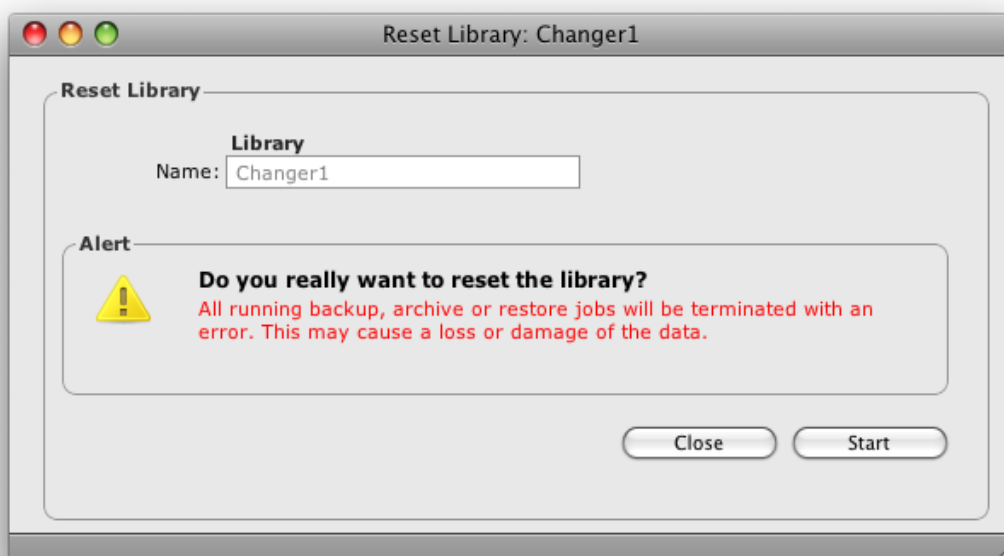
6.3.4.4 **Import Volume...**

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

6.3.4.5 **Reset Jukebox...**

This option 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.



6.4 The Volume Pool

The volume 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.

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.



6.4.1 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:

 Copy data to a second set of media.' At the bottom right are 'Close' and 'Apply' buttons."/>

Pool Configuration: Media2011

General Setup

Pool **Status**

Name: Media2011 Enabled

Description: Media 2011 Tape Archive

Usage

Use for: Archive

Media type: TAPE

Drives

Use listed drives only:

Drive name:

Clone data: Copy data to a second set of media.

Close Apply

Pool Name:

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: Used both for backups and archives.
- » DISK: Used both for backups and archives.
- » LTFS: Special Media pool for LTFS formatted tapes, see section 12, Handling LTFS volumes with P5

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.

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.

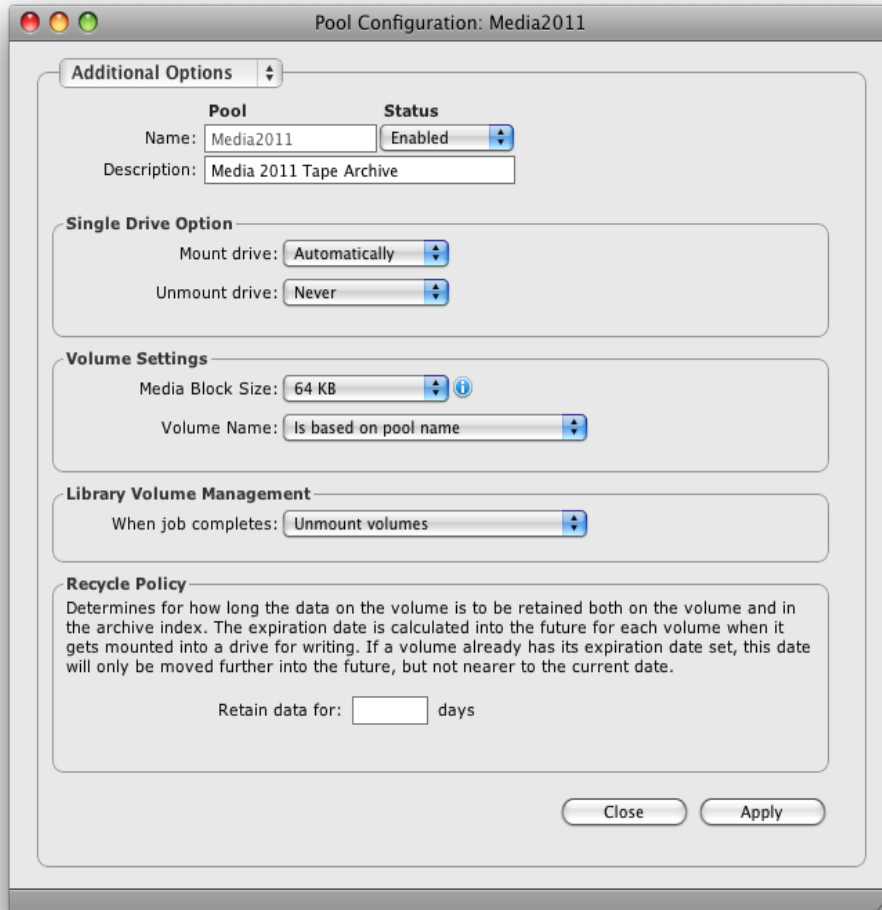
Clone Data:

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.



6.4.2 Pool Configuration: 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. Un-mounting releases the drive.

- » Manually: The mount of a volume in a single drive (device) must be done manually.
- » Automatically: Each of the drives given 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 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.

Moment 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.

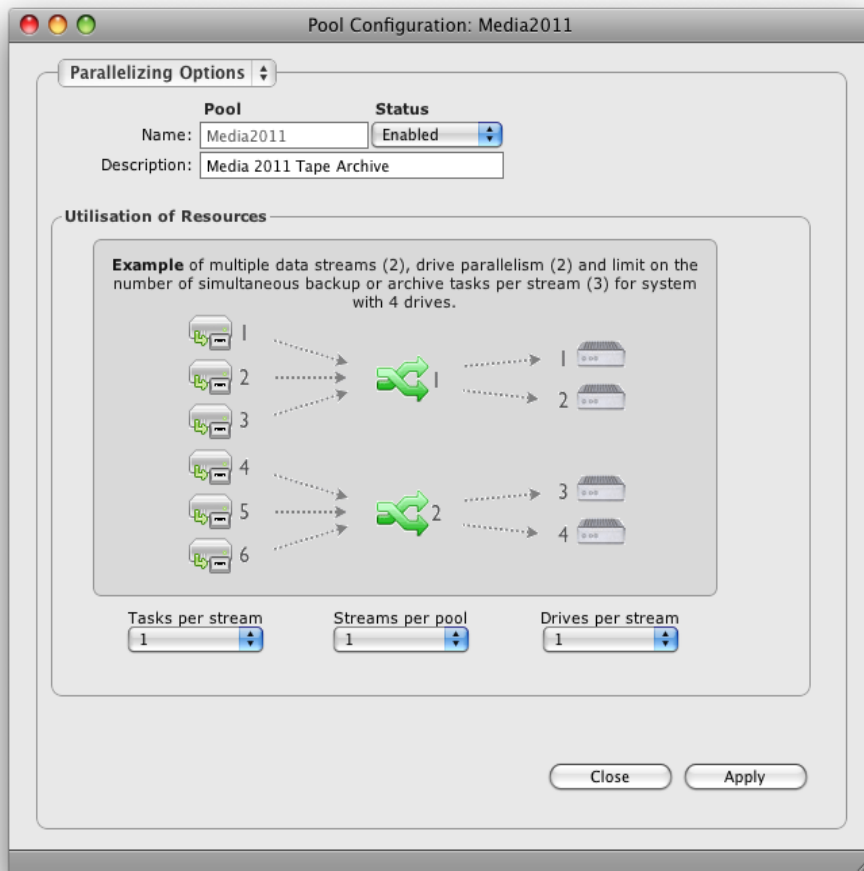
» When volumes are needed:

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.

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

6.4.3 Pool Configuration: 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 set configured separately for each pool. The following window allows to define:



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.



6.5 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.

6.5.1 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.

6.5.2 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 Information: 10182

Volume Information

Label
Name: ArchiwareProgressive.0006

Status
Enabled

Operation
Condition: Ok
Mode: Recyclable

History
Date labeled: 12.Oct.2011 02:28
Last used: 17.Oct.2011 02:32
Date Filled: 17.Oct.2011 02:28
Expiration: 21.Oct.2011 01:24

Location
Offline Location:
Location:
Library:
Position:
Drive:

Media and Pool Infos
Pool: ArchiwareProgressive
Type: LTO
Kind:
Bar code: A00056L3
Org ID: 0
Block size: 64 KB

Error Counter
Nonrecovered write: 0
Recovered write: 2
Nonrecovered read: 0
Recovered read: 0

Performance
Approx. Used: 605.51 GB
Time used: 24:21 hours
Used count: 13

Close Apply

**Volume Label:**

The volume name is automatically generated during labeling. It consists of the name of the pool 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, 'Usage' 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.
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 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 could be corrected.



Non-recovered read The number of reading errors that could not be corrected.
Warning: See 'Non-recovered write'

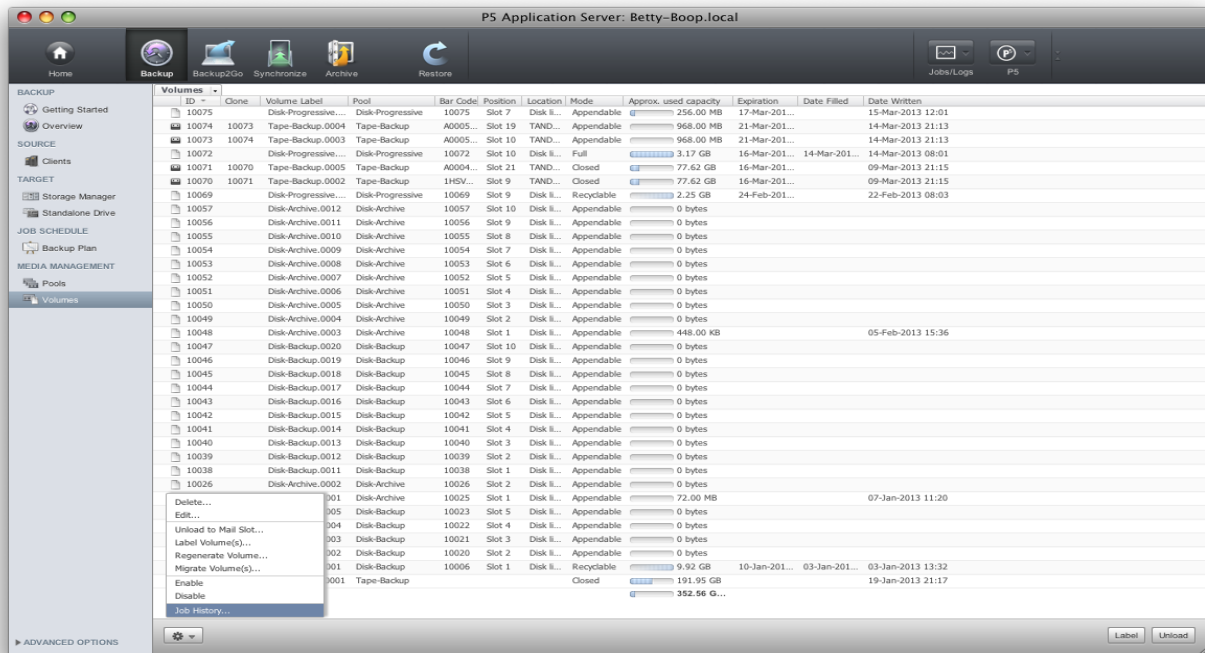
Recovered read The number of read errors that could be corrected.

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



6.5.3 Volume Manager: Gear Menu

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



6.5.3.1 Label Volume

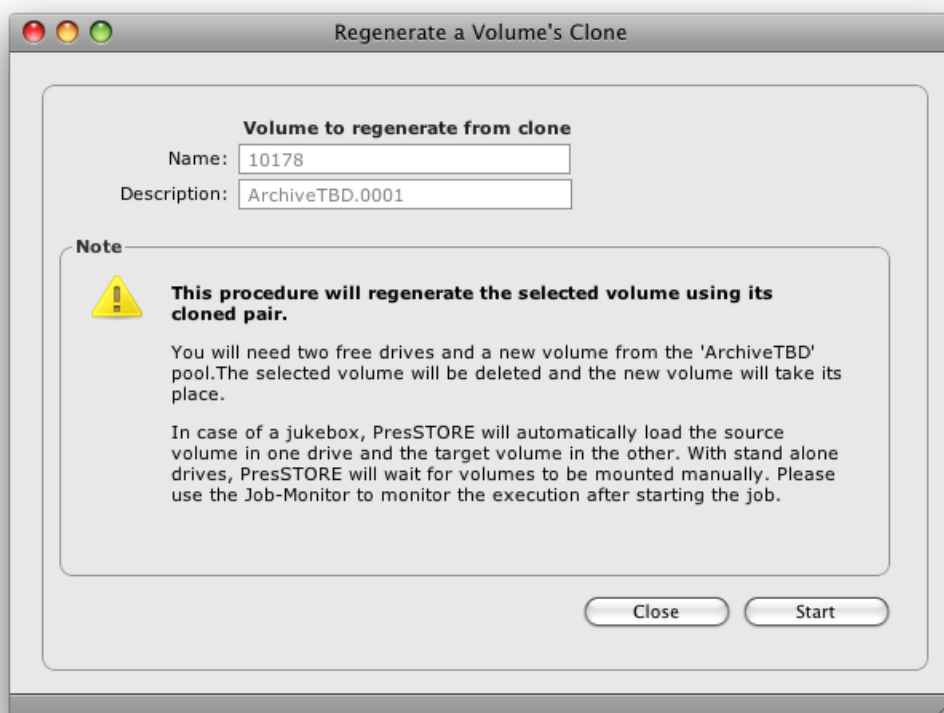
For a description refer to section 10.1, Labeling Media



6.5.3.2 **Regenerate Volume ...**

This option allows you to recreate a volume of a backup or an archive pool. To start this process use the context or the gear menu.

To regenerate a volume, the volume must have been written with cloning and the clone volume must be made available. You should use this when one of the clones gets broken.



Before starting a volume regenerate job, 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.

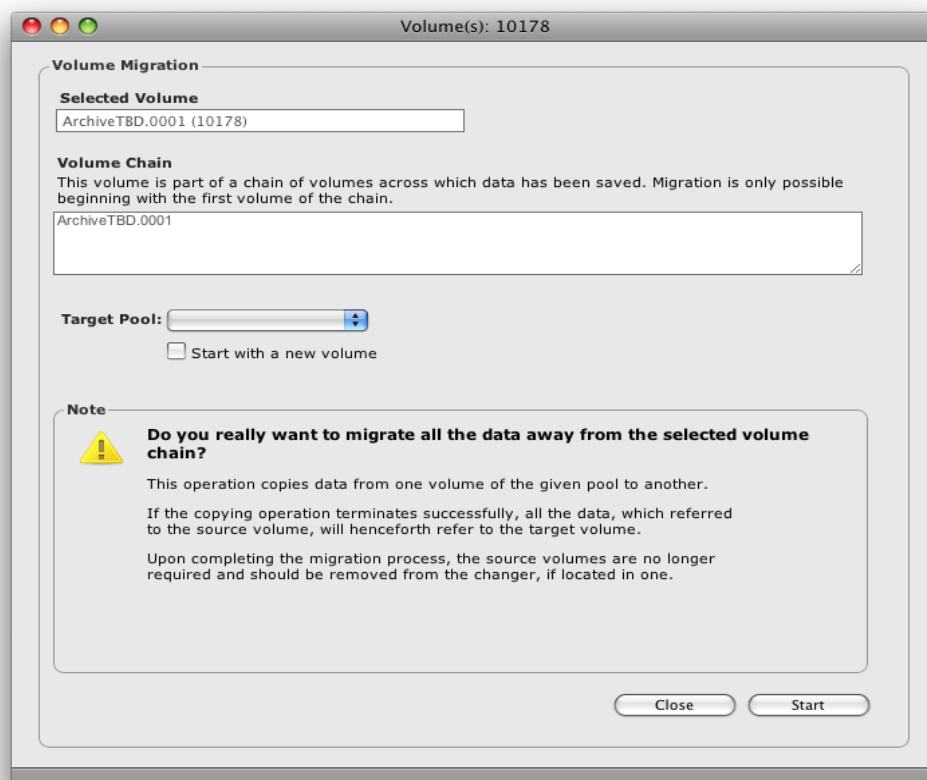
See section 6.4.1 Creating a Pool for where to set the cloning option.



6.5.3.3 *Migrate Volume(s)...*

You can use this function to migrate your data from one type of media to another, for example from DAT to LTO tapes. This therefore enables a technological upgrade of your archive. Migration is supported for archive media only.

The data of the source pool will be migrated to the target pool and all the references to the data will be adjusted. After completing the migration, the source volumes are not required any longer and should be removed.



When volumes are written, data at the end of the tape that does not fit completely is split and the rest of the files is written onto the next tape. This creates chains of volumes that should be migrated together in order to also migrate the files split across two volumes. P5 will follow the tape order through these chains during migration.



Note that files at the end of the migrated tape may have been continued on the next subsequent written tape. If so, that follower tape is requested during migration, too.



6.5.3.4 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 is cannot be assumed that the complete data written by the named job resides on this tape.

Plan name/operation	Client	Start time	End time	Job
Save backup index for client 'Mickey'	localhost	04.May.2011 04:40	04.May.2011 04:44	26012
Progressive	Mickey	04.May.2011 00:55	04.May.2011 04:40	25967
Save backup index for client 'fileservr'	localhost	04.May.2011 01:14	04.May.2011 01:16	26000
Progressive	fileservr	04.May.2011 00:55	04.May.2011 01:16	25965
Save backup index for client 'kompilierwin'	localhost	04.May.2011 01:14	04.May.2011 01:14	26001
Save backup index for client 'SunSparc98'	localhost	04.May.2011 01:07	04.May.2011 01:09	25986
Progressive	SunSparc98	04.May.2011 00:55	04.May.2011 01:07	25964
Save backup index for client 'LinuxX86-96'	localhost	04.May.2011 01:05	04.May.2011 01:07	25984
Save backup index for client 'MacPPC-53'	localhost	04.May.2011 01:04	04.May.2011 01:04	25982
Progressive	LinuxX86-96	04.May.2011 00:55	04.May.2011 01:04	25963
Progressive	MacPPC-53	04.May.2011 00:55	04.May.2011 01:03	25969
Save backup index for client 'SunX86-99'	localhost	04.May.2011 01:02	04.May.2011 01:03	25978
Save backup index for client 'iMacX86-104'	localhost	04.May.2011 01:02	04.May.2011 01:03	25979
Progressive	iMacX86-104	04.May.2011 00:55	04.May.2011 01:01	25968
Save backup index for client 'fmserver'	localhost	04.May.2011 01:01	04.May.2011 01:01	25976
Save backup index for client 'Studenten'	localhost	04.May.2011 01:00	04.May.2011 01:01	25972
Progressive	fmserver	04.May.2011 00:55	04.May.2011 01:00	25966
Progressive	portal	04.May.2011 00:55	04.May.2011 00:59	25970
Progressive	Studenten	04.May.2011 00:55	04.May.2011 00:59	25971
Progressive	Mickey	03.May.2011 00:55	03.May.2011 05:10	25909
Labeling for pool ArchiwareProgressive	localhost	26.Apr.2011 13:48	26.Apr.2011 13:54	25324



6.6 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*.

6.6.1 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.

Archive Plan: 10003

General Setup

Description	Status	Auto start
Archive Plan: Media2011	Enabled	Disabled

Data

Target pool: Media2011 Set of volumes for archiving data

Target index: Default-Archive The target database for indexing archived data

File filter: Use a filter to exclude data from archiving.

Access

List of group(s): A list of valid groups to allow access to this archive plan. If empty, all groups are authorized.

A valid group: - +

Close Apply

**Description:**

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

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.

Target pool:

The pool defines onto which group of volumes the data will be saved. See section 6.4, The Volume Pool

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 6.14, Index Databases.

Filter:

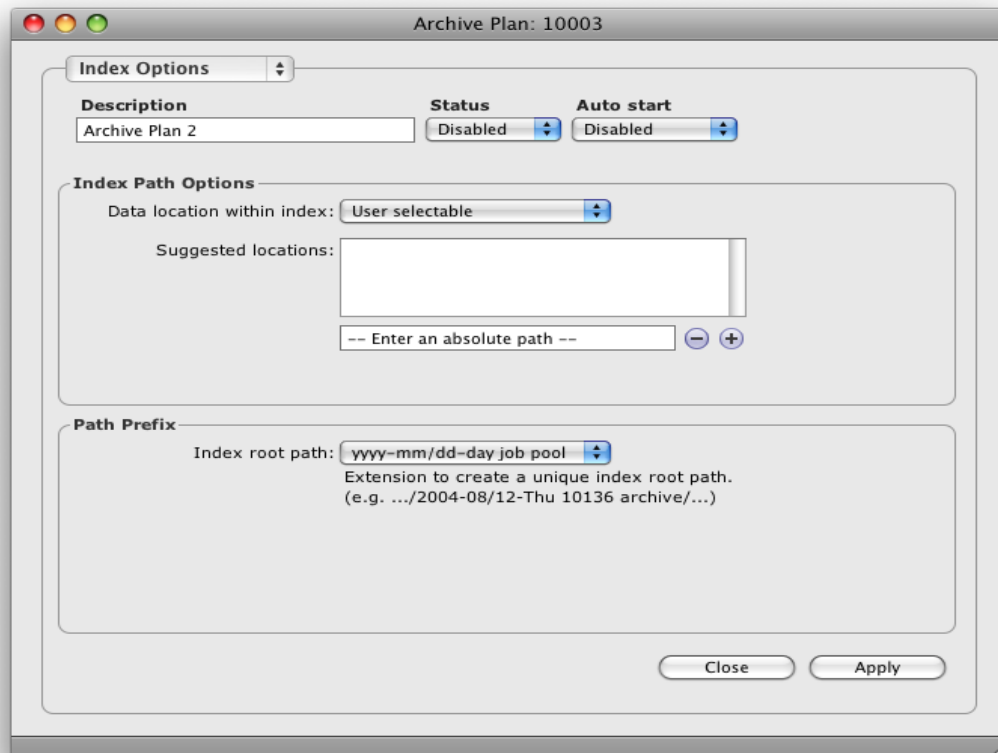
If you wish to filter out certain data from the storage process, you can choose the appropriate filter here. See section 6.9, Filters.

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.



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 two possibilities for this setting:

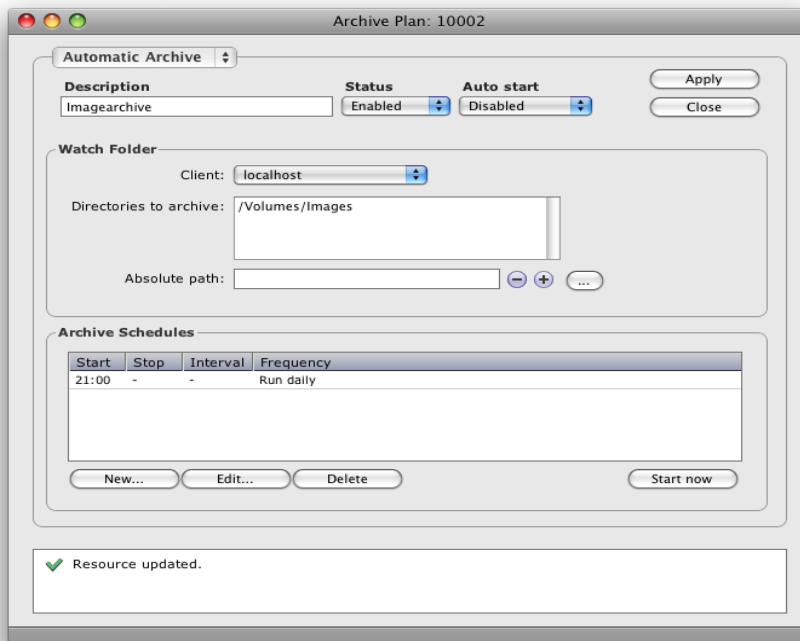
- » Predetermined for user: Absolute Path:
A single path is defined under which all the data archived with this plan will be kept in the index. The user cannot change this path or choose an alternative.
- » User selectable: Suggested Locations:
A list of paths is suggested to the user when starting an 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 anyway 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 that is when manually archiving), this option is not required.

1. 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 included 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 6.12, P5 Clients.

Directories to archive:

All the folders that you wish to back up have to appear here.



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

2. 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, at which time it should be run. It is possible to rerun the archive in intervals by specifying an interval and an end time. According to the frequency 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.

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.

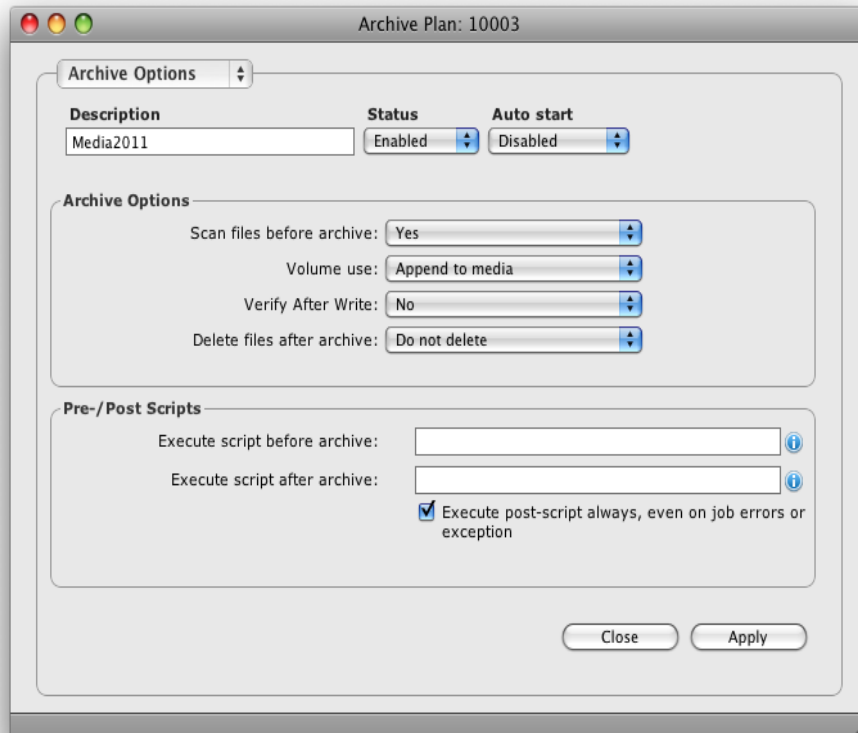
Exception:

Exceptions can be set 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



Scan files before archive:

This field controls whether a list of files to be archived will be obtained before the actual operation takes place. If set, you can observe the progress of the plan in terms of size and file count while the plan is executed. If disabled, P5 immediately starts the operation and cannot display this information. The choices are:

- » **Enabled**: A list of files and sizes will be collected.
- » **Disabled**: A list of files and sizes will not be collected.

Volume use:

Normally the data of consequent 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.

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:

Choose 'yes' to remove the files from the on-line storage after having been successfully archived.

Directories will not be removed. The choices are:

» Yes: Remove the files.

» No: Do not remove the files.

Archive Level:

Choose '*full*' to archive all files again. Choose '*incremental*' to archive files only in case these are not yet contained in the archive. Note that if a file is explicitly selected, it is added even if it already exists in the archive. 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
```

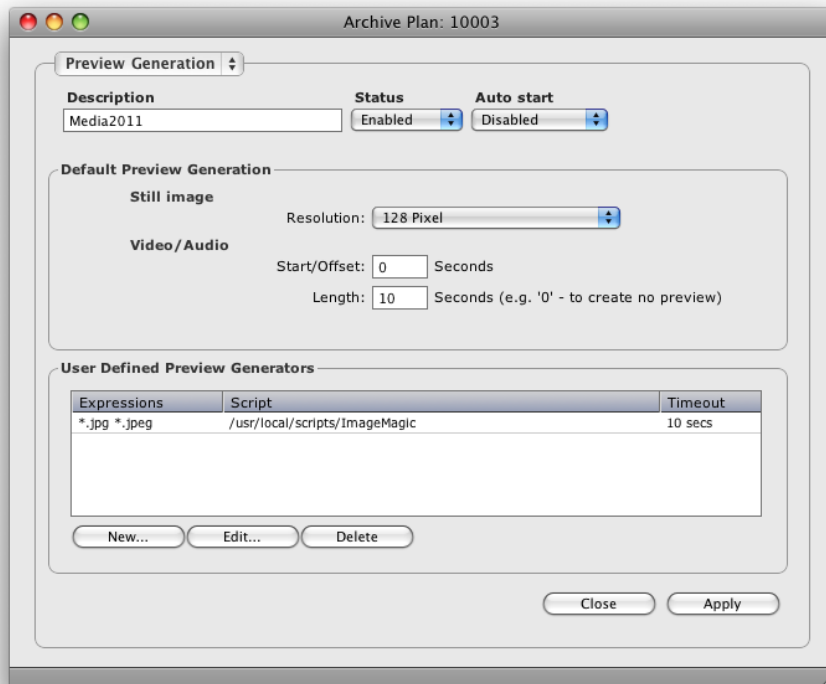
on a windows machine:

```
/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.



Still image resolution:

Select the image resolution here

Video/Audio Offset and Length:

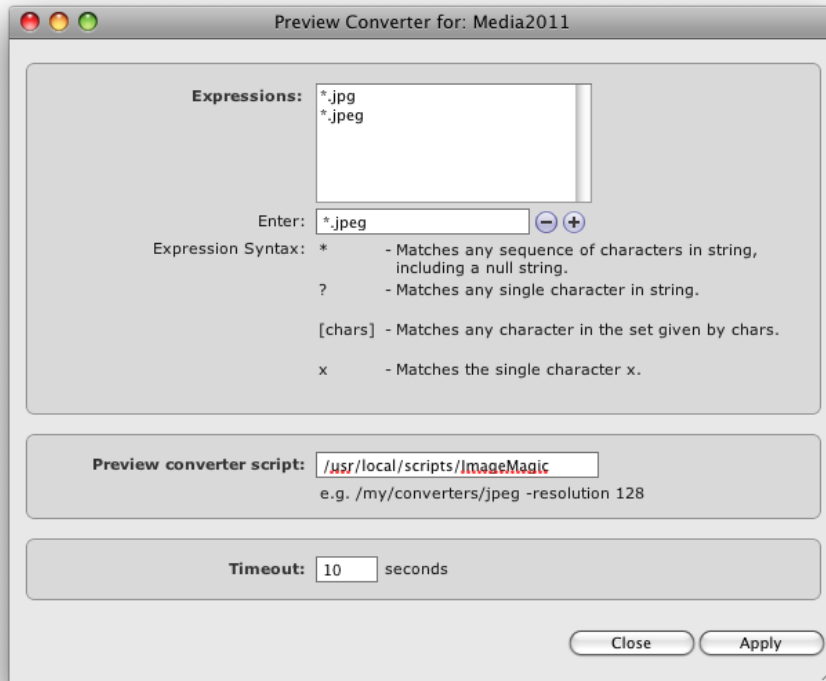
Define the offset and length for audio and video data previews.



Note that preview files in P5 are accepted up to 128 MB in size per preview file.



User Defined Preview Generators



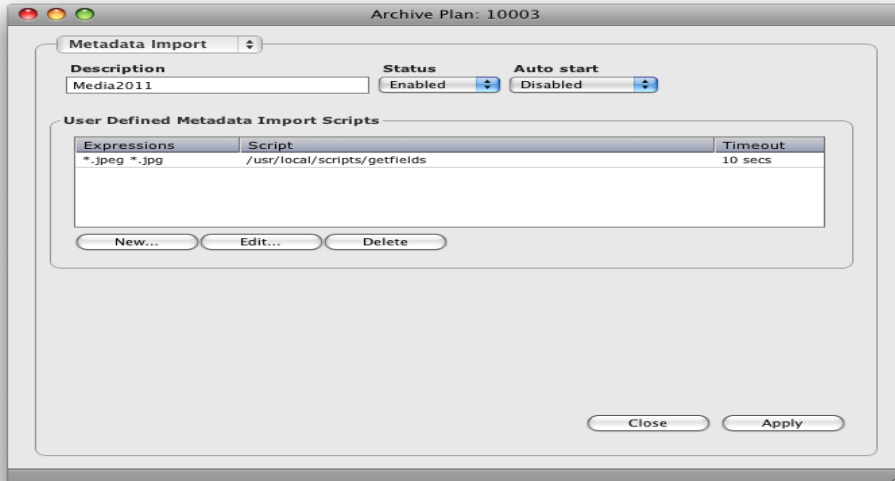
As an alternative to the default preview converter, it is possible to configure your own converter commands. For each converter definition, a case sensitive list of extensions and a command can be given. For files with one of the listed file name extensions, P5 will call the converter command. The converter 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.

The length, size and offset parameters of the built-in converter do not apply.

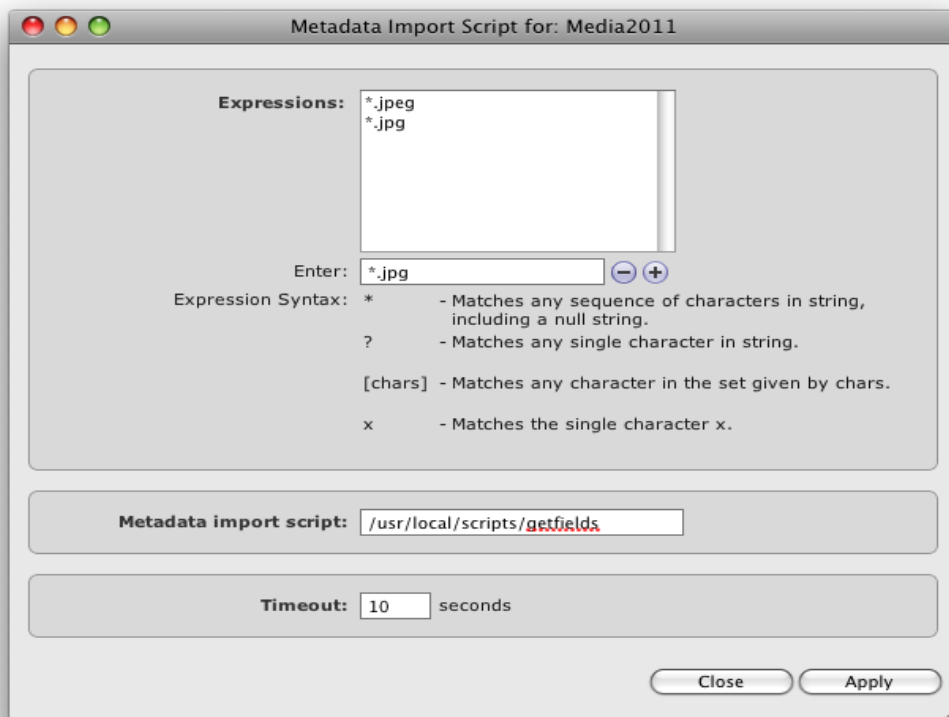


Automatic Metadata Import

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.



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 filename 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.

By example an output line like

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

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

See section 6.14.2.1 Fields... for a description how to add your own metadata fields to an archive index database.

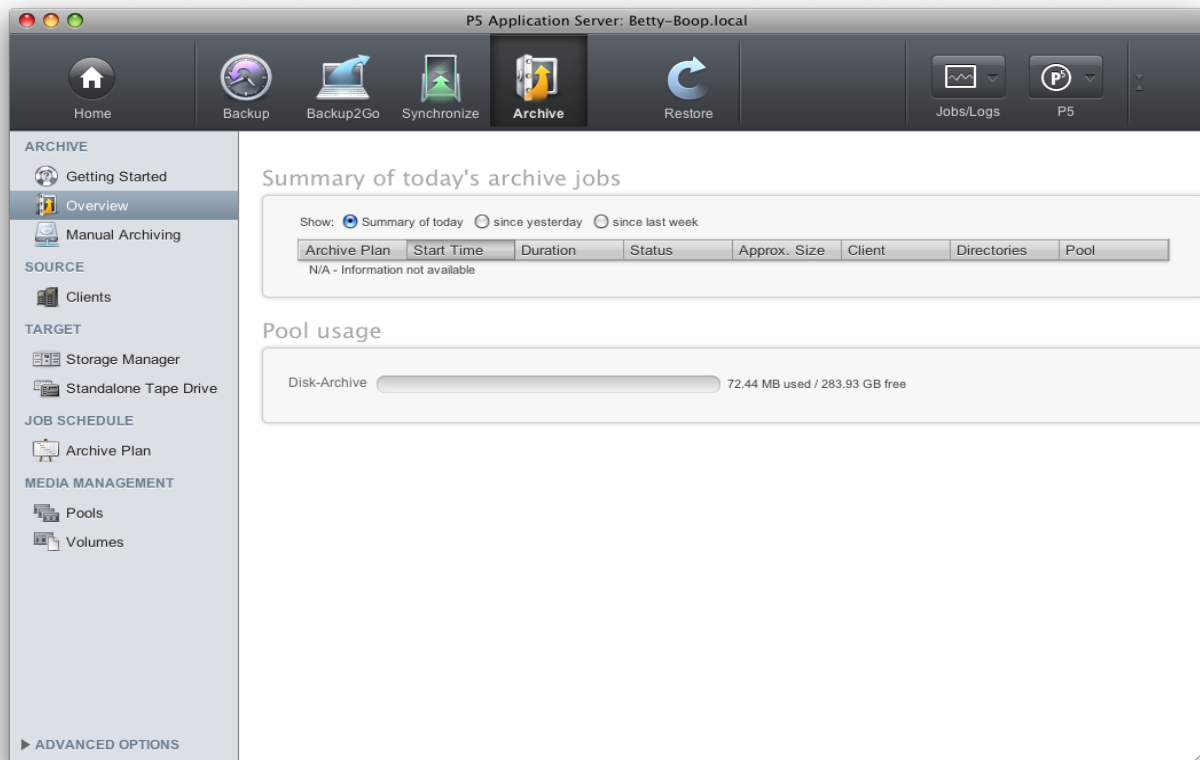
To confirm your input, click **Apply**. In case all the entered data is correct, the plan will be updated.



6.6.2 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.





6.7 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 to save the data to is a volume pool. A backup run can save either entire directory structures, or only data (files) that have 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 back up are defined in the task.

6.7.1 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.

The screenshot shows a web-based configuration window for a backup plan. The title bar reads "Backup Plan: 10009".

General Setup

Description: Progressive

Status: Enabled

Backup Tasks

Client	Directories to backup
Pluto	'ALL'
kompllierwin	'/X/Ibrahim'
MarcMiniMac	'/Volumes/MMBSync250'
SunX86-99	'/usr/local/homes', '/usr/local/bin', '/usr/local/etc', '/usr/local/hellos', '/usr/local/include',
LinuxX86-96	'/usr/local/homes/zv', '/usr/local/bin', '/usr/local/etc', '/usr/local/hellos', '/usr/local/includ
SunSparc98	'/', '/space/bin', '/space/etc', '/space/homes', '/space/include', '/space/lib', '/space/ns', '/

Buttons: New..., Edit..., Delete

Backup Schedules

Level	Start	Time	Frequency	Target Pool
Incr	27 Oct 11	00:55	Run daily	ArchIwareProgressive

Buttons: New..., Edit..., Delete, Start now

Buttons: Close, Apply

**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 this backup:

Enter when the backup plan should run.

- » autostart disabled: the plan will not be run automatically, it has to be started manually
- » at start time: the backup plan will be run according to the schedule

When using chained backup plans with the start option after previous plan, the very first plan in the chain defines the start date of the whole chain. P5 searches for a backup schedule event in the subsequent plan that starts on the same date as the plan at the beginning of the chain which runs at the same or later time. If such a backup schedule exists, the backup level and media pool are taken from that event. Otherwise the execution for the corresponding plan is skipped.

Backup Plan: Additional Options

Backup Plan: 10001

Additional Options

Description: Backup Plan: Classic localhost Status: Disabled

Backup

Scan files before backup: No

Volume use: Use new media

Job timeout for waiting on volumes: Unlimited

Recycle Policy

Defines how long data written by this backup plan should be kept available for restoring. Each time a volume is written, the expiration date on that volume is set according to this retention period. Once the expiration date is reached, the tape will be set as 'recyclable' and can be recycled (reused) by the system.

Retain data for: 7 days

Reschedule backup of data:

- by a full backup event, or when volumes get recycled (traditional backup)
- when the retention time is reached (progressive backup)

Logging

Job activity log: only errors

Close Apply

**Scan files before backup:**

This field controls if a list of files to be backed up will be obtained before the actual operation takes place. If set, you can observe the progress of the plan in terms of size and file count while the plan is executed. If disabled, P5 immediately starts executing the plan and cannot display this information. The choices are:

- » Yes: A list of files and sizes will be collected.
- » No: A list of files and sizes will not be collected.

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 to not 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.

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.

Recycle Policy, backup renewal:

This option defines how the renewal of the saved data shall take place:

- » Traditional Backup: Renewal is done by running full backups (scheduled or manually), expired media may be recycled. Data missing due to recycling is added automatically with the next incremental backup
- » Progressive Backup: Data renewal is done automatically, data from expired media is saved again before recycling. Media with data that would be lost are not set to recyclable.

Job activity Log:

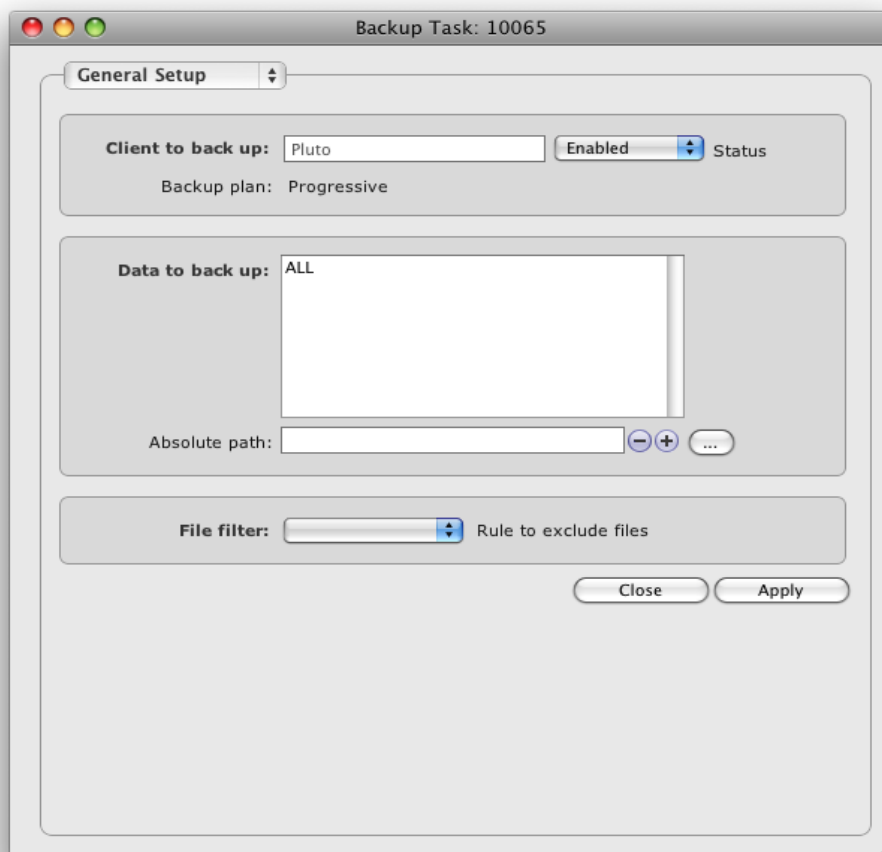
By default, only errors are kept in the activity log. Select here *Errors and file operations* in order to also log each file that is added to the backup. This setting will be applied for a single job execution and reset afterwards in order to avoid huge logs.

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



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 6.12, P5 Clients.

Data to back up:

Add all the client's folders to back 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

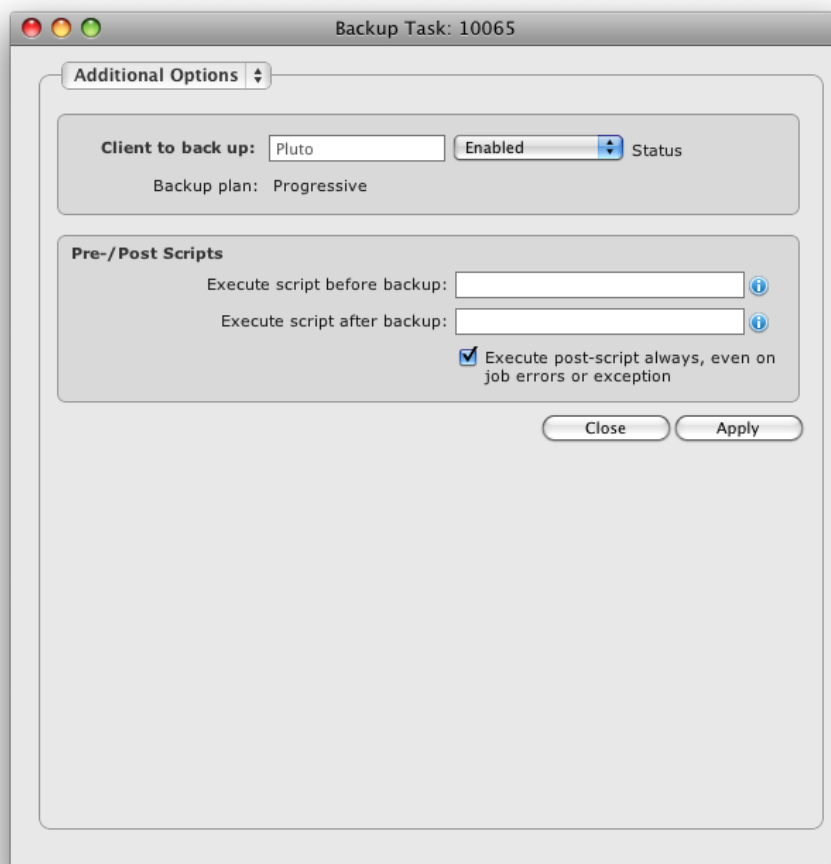
- The selected path will be removed from the list

**File System (Volume, Partition):**

The contents of mounted file systems within the given folders will not be saved, only the mount points themselves will.

Note: the path of every file system to be saved has to be entered into the list as described above.

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, or 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:

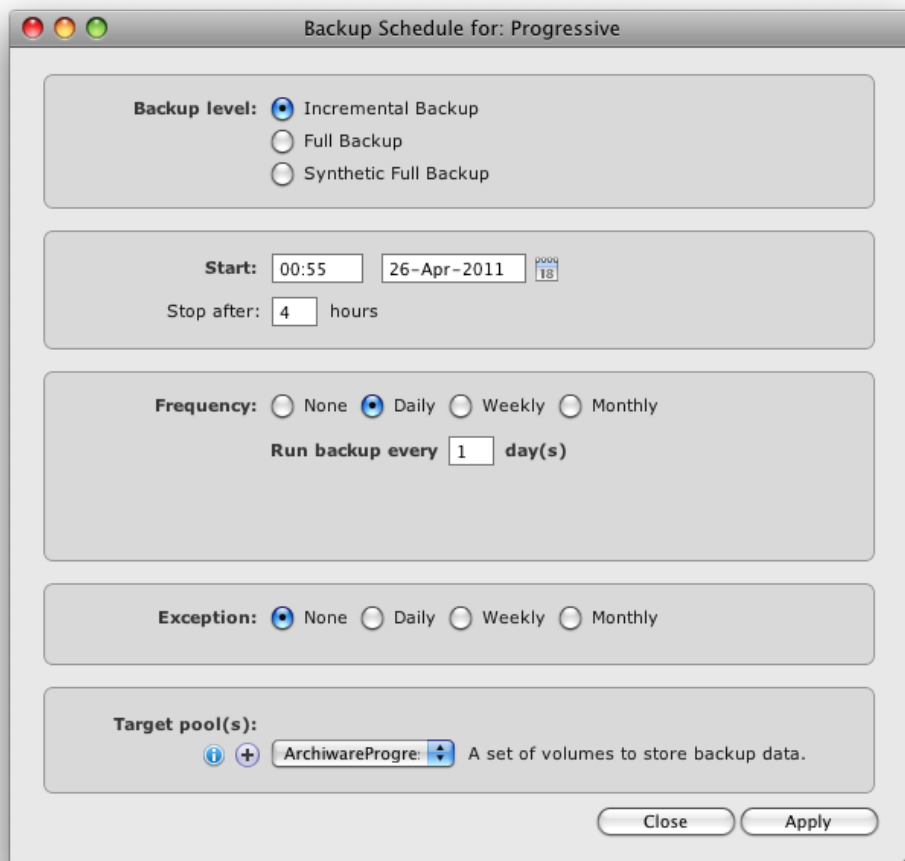
```
dbserver:/usr/local/scripts/db-stop
```

on a windows machine:

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

1. 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 changed since the last full or incremental backup are included in the backup.
- » Synthetic A full backup is done, but the files are taken from existing backup media, not from the file system.

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 pull-down-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 according media pool are saved again.

In general, it does not make sense to specify backup events with one and such with multiple target pools in 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.

6.7.2 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

Backup overview

Show: All Successful Other

Client	Backup Plan	Last Run	Last Successful	Directories	Last Pool	Next Run	Next Pool
redstone	Progressive Redstone	Error	9 days ago (0 bytes)	/usr/bin	Disk-Progressive	16-Mar 08:00	Disk-Progressive
	Classic localhost	Error	9 days ago (12.68 GB)	/	Tape-Backup	15-Mar 21:00	Tape-Backup
localhost	Classic localhost	Error	9 days ago (1.29 GB)	/	Tape-Backup	15-Mar 21:00	Tape-Backup
	stoerfried	Finished	29 days ago (0 bytes)	/Users/tannir	Tape-Backup		

N/A - Information not available






Pool usage

- Disk-Backup: 0 bytes used / 297.92 GB free
- Disk-Progressive: 3.42 GB used / 4.00 GB free
- Tape-Backup: 349.07 GB used / 1.61 TB free






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, that the plan is currently running.

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

-  Green: The job finished ok
-  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 ok the last time:

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

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



6.8 Synchronize 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 or which later get created 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.

6.8.1 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.

The screenshot shows the 'Sync Plan: 10001' configuration window. It is divided into three main sections: 'General Setup', 'Source of Data', and 'Destination for Data'. Below these is a 'Synchronize Schedules' section with a table and buttons.

General Setup

Sync Plan: Status: Auto start:

Source of Data

Client:

Directories:

Absolute path:

Destination for Data

Client:

Absolute path:

Synchronize Schedules

Mode	Start	Stop	Interval	Frequency	Filter
Not currently scheduled. Click New to create a new synchronize event.					

**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.

**Note:**

Either 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.

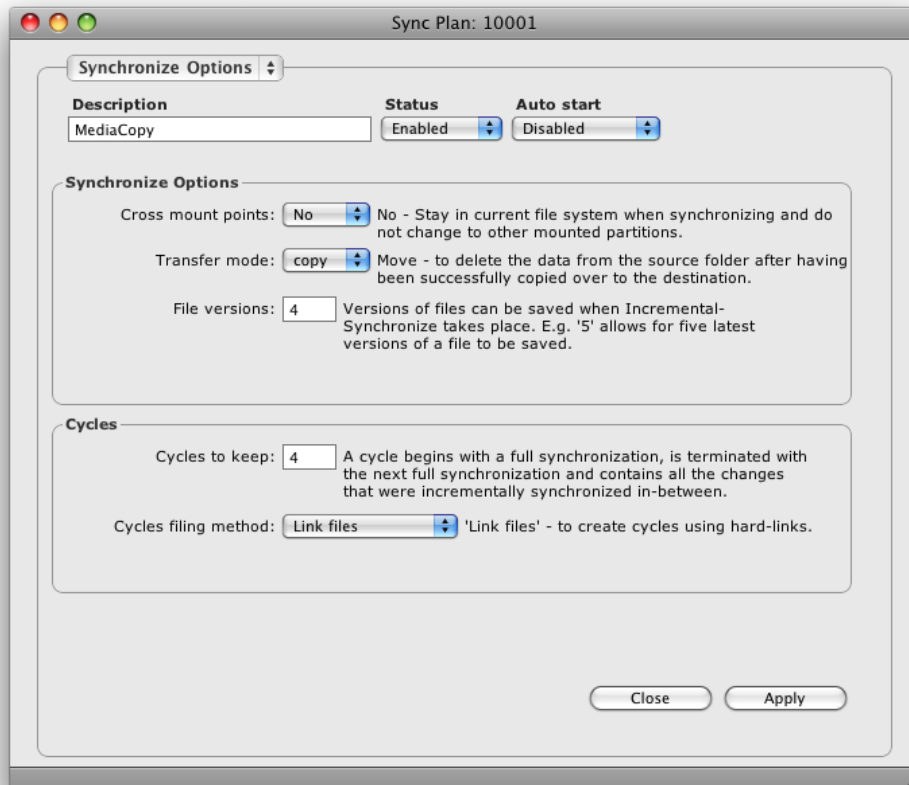
**Note:**

P5 Synchronize allows the user to specify cloud storage as data source or target, thus allowing to transfer data to and from a web based storage.

In order to do so, a Cloud storage must be defined first, please see section 6.10, Cloud Storage.



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 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, an entirely new copy of the data will be created on the target. Later, in the Synchronize Schedule events, you can specify if you would like to keep the cycles while executing the event.

Cycles filing method:

When retaining a cycle, P5 can either make a physical copy of each file in the cycle or employ hard links to reduce the size used on the file system.

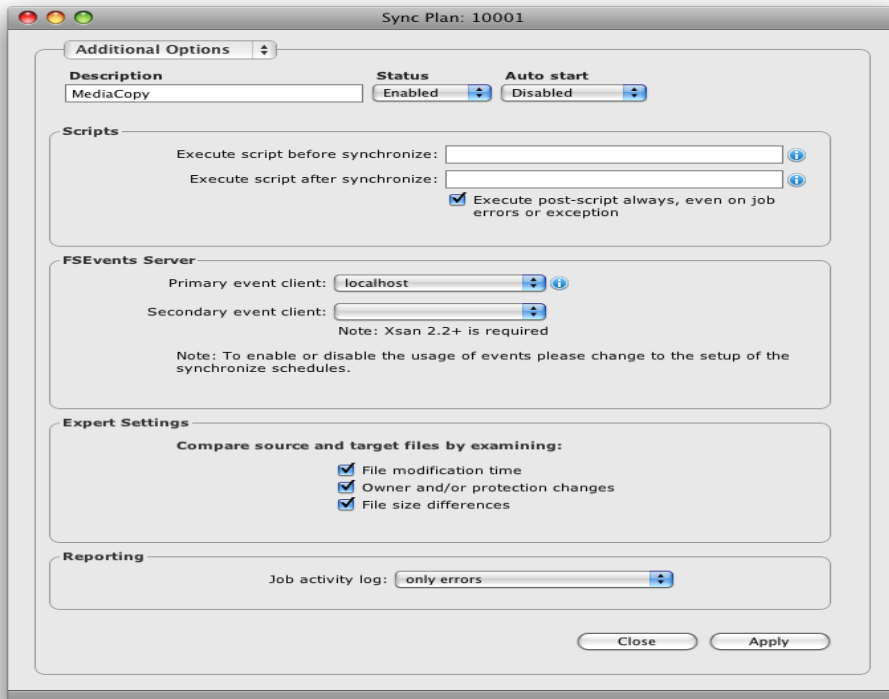
FSEvents Server:

This setting is specific to the XSan 2.2 file systems. 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 a FSEvent server for filesystems local to the selected client.

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
```

Expert Settings

These settings are used to tell Synchronize which attributes of the file to use to make its decision 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).

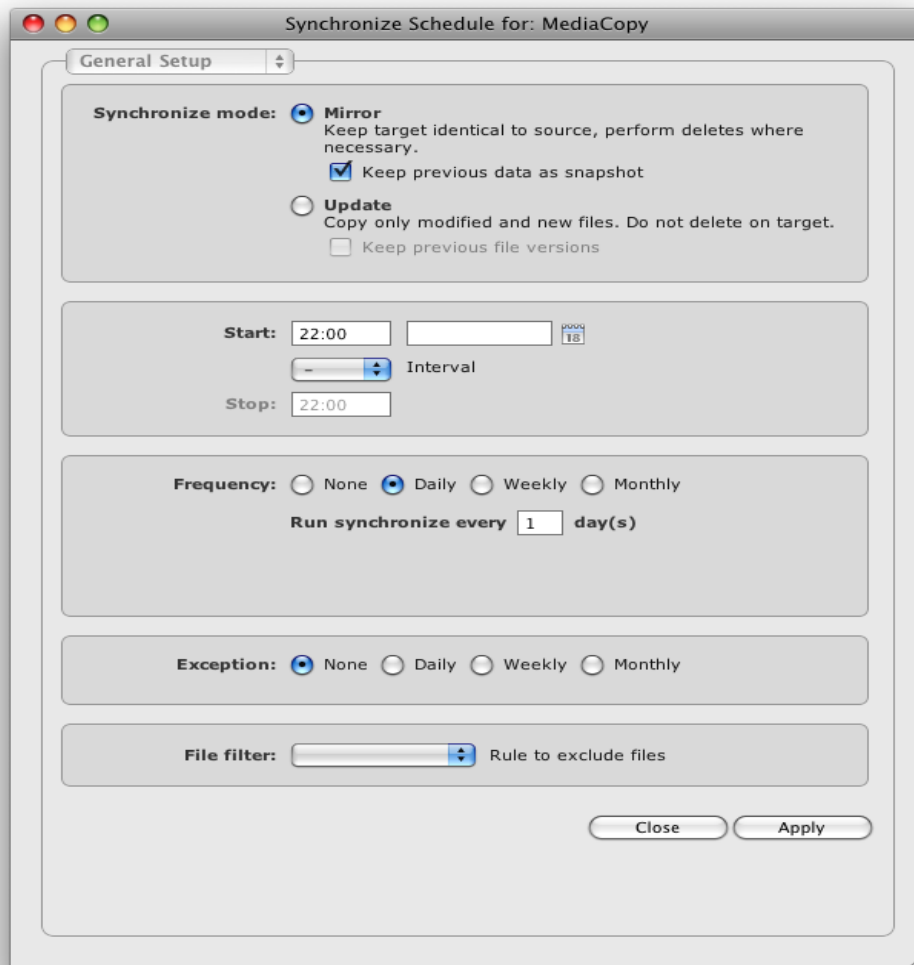
When using the "Owner and/or protection changes" for file comparison, the target filesystem MUST support extended attributes (xattr). Current standard filesystems (like HFS, EXT3/4, NTFS) do so.

Job activity Log:



By default, only errors 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 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 time, 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 over 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.

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 Mac OSX 10.4 or later with 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 with kernel version 2.6.13+ 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.

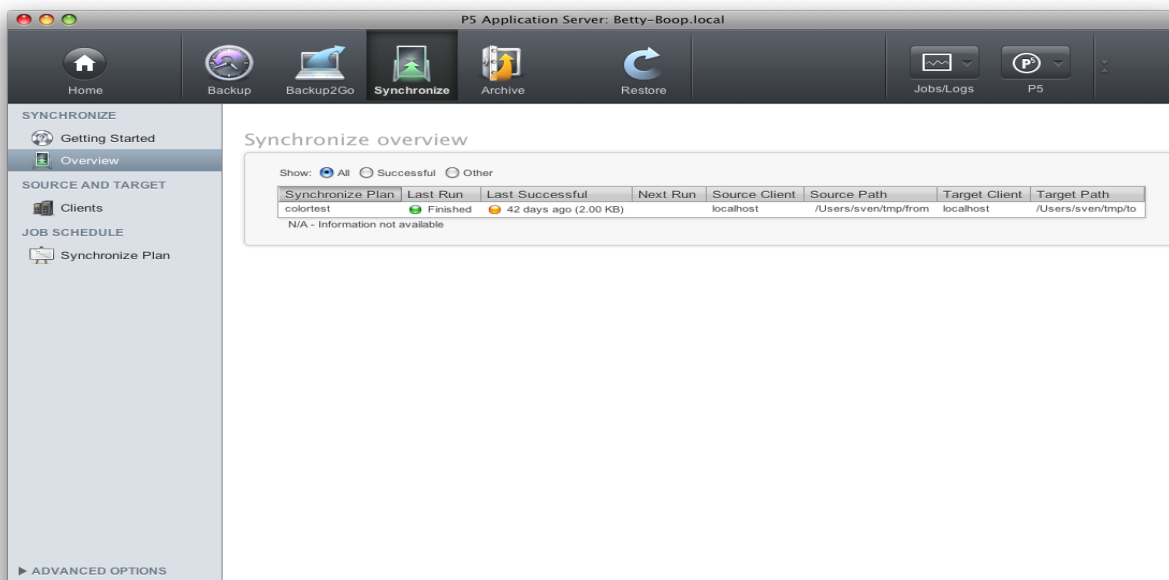


6.8.2 Synchronize Overview

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, 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, that the plan is currently running.

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

-  Green: The job finished ok
-  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 ok the last time:

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

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



6.9 Filters

While executing a backup, a synchronize or an archive plan, the entire tree consisting of all the files and folders within the selected folder will be saved in an in-place recursive walk of the tree. 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 *File Filter* from the *bottom bar* to create a new filter.
4. Fill in the required fields.

Filter Configuration: Archive-Helios

General Setup

Name	Description	Status
Archive-Helios	EtherShare Archive Filter	Enabled

Filter condition to reduce the file list

Select files matching this expression

Exclude files matching this expression

```
-name '*.lay' -type file  
-name 'layouts' -type directory
```

Add Entry

File type	File name	Modification date	Size
--- Select ---	--- Select ---	--- Select ---	--- Select ---

Close Apply

**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 to-be-processed files 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.

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.

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.

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.

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.

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 Cloud Storage

P5 supports synchronizing data to or from a cloud storage accessible via the WingFS interface.

WingFS is an adapter that allows addressing cloud storage like a file system, developed by ARCHIWARE.

6.10.1 Defining a Cloud Storage

1. Log into P5 using a web browser as an admin user.
2. See section 4.1, Starting the Browser.
3. Select the *Synchronize* module from the top menu
4. Select *Cloud Storage* from the sidebar.
5. Click **New** in the Cloud Manager's bottom bar to define a Cloud Storage.
6. Fill in the required fields.

Cloud Storage:

Cloud Service Provider: SFTP (SSH)

Please enter your login information:

Host: FTP_NAS

User name: root

Password: ****

Close Apply



Defining and using a cloud storage in P5 Synchronize requires WingFS to be installed. For the download of WingFS and for further info, see <http://www.wingfs.com>



6.10.2 Supported Cloud Storage

P5 supports addressing the following cloud storage:

- Amazon S3
- Google Drive
- Google Cloud Storage
- SFTP(SSH)

WingFS is currently (Spring 2013) available for Mac OSX 10.6 +.



6.11 Logs

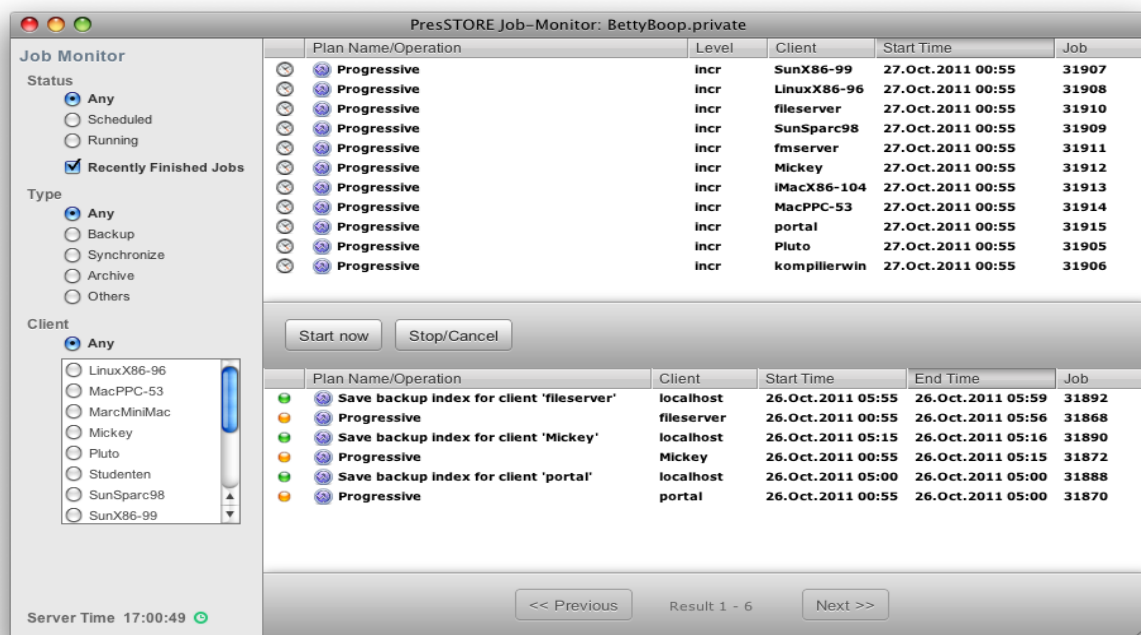
P5 automatically creates and manages logs and protocols. Every job results in a log entry upon completion. You can view this entry with the data browser.

Job results can also be e-mailed to selected user(s), see 11.2, Customizing Reporting. At the same time, notifications can be sent to Macs where growl must be configured to receive the notifications and to Windows hosts using bubble notifications in the Windows task bar.

6.11.1 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:

Status: Show only planned or only running jobs.
Clicking the check box *Recently finished jobs* will split the screen and show the recently finished jobs in the lower half.

Type: 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)

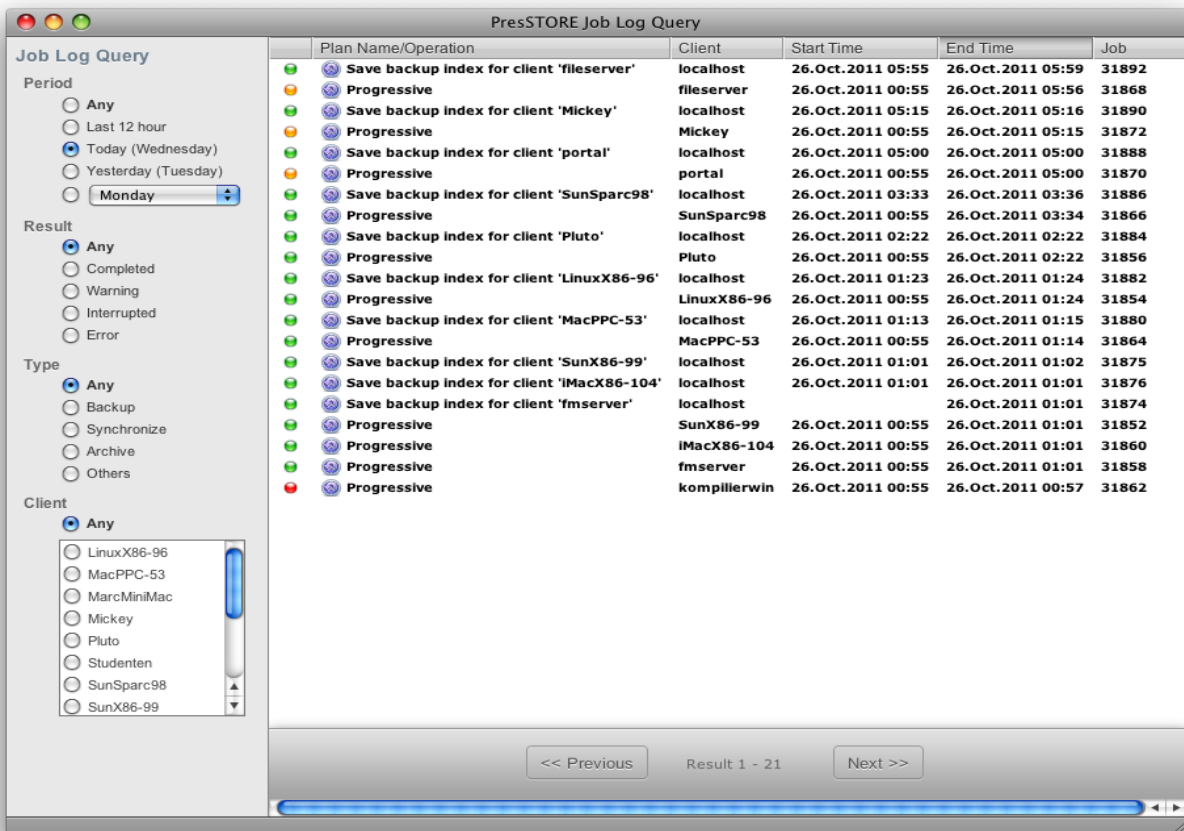
Keeping the job monitor window open will also prevent the current log-in session from expiring.



6.11.2 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



By default, the job monitor shows the jobs that ended in 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.



6.12 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.

6.12.1 Creating a Client

7. Log into P5 using a web browser as an admin user.
8. See section 4.1, Starting the Browser.
9. Select the **Archive**, **Backup** or **Synchronize** module from the top menu
10. Select **Client** from the sidebar.
11. Click **New** in the client managers bottom bar to create a Client .
12. Fill in the required fields.

The screenshot shows a web-based configuration window titled "Client: Pluto". The window is divided into several sections:

- General Setup:** Contains a "Client" section with "Name: Pluto" and "Status: Enabled" (dropdown). Below it is "Description: Neuer Linux Web Server (Coeno)".
- Host:** Contains "Computer is licensed as: Virtual Server" (dropdown), "TCP/IP Address: 192.168.111.4" (with examples like 192.168.10.50), "Port: 8042" (with example 8000), "User name: root" (with example root), and "Password: *****".
- Communication:** Contains "Compress data transfer: No" (dropdown, with example Yes), "Encrypt data transfer: No" (dropdown, with example Yes), and "Limit bandwidth: 100%" (dropdown, with example '100%').
- PresSTORE Upgrade:** Contains "Automatic upgrade: Yes" (dropdown, with example Yes).

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

**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 type:

Use `Server` for a server operating system and `Workstation` otherwise. P5 will not allow workstation license to be employed for a server. Using a server license for workstation is allowed.

- » 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.

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 log-in name of the user who has sufficient permissions to access the data to be saved from, usually `root`.

Password:

The password corresponding to the log-in. 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 here if the data should be encrypted for transferred. 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.

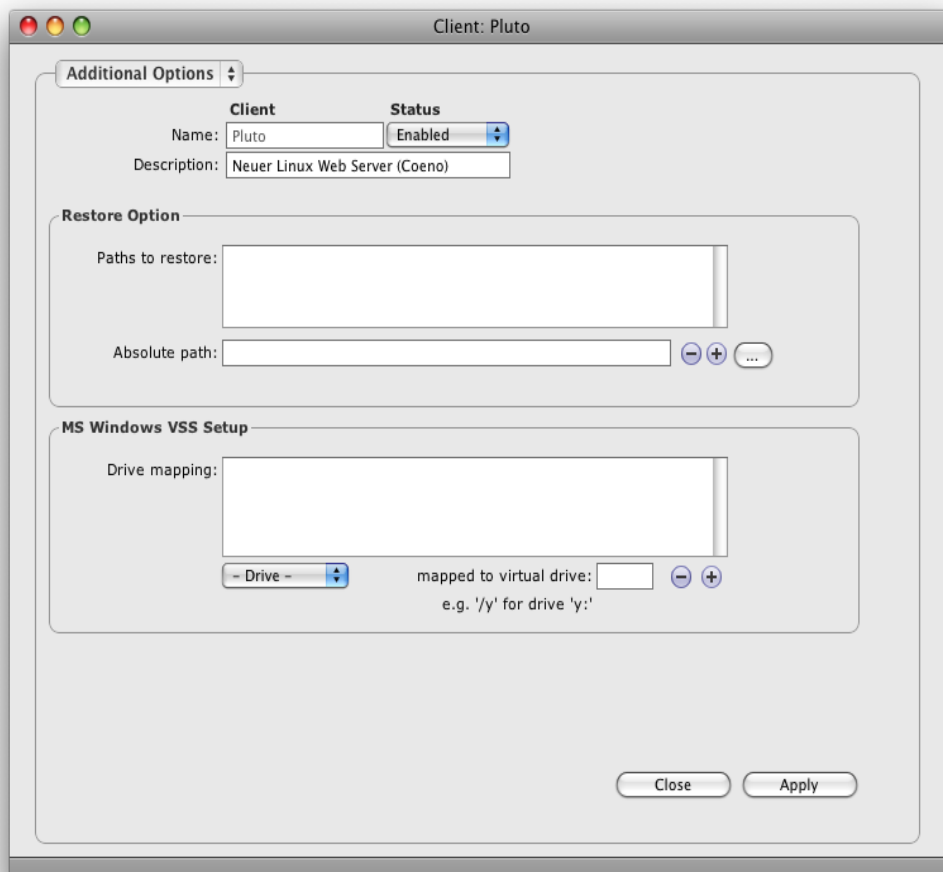
Limit bandwidth:

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.

Automatic upgrade:

Allow automatic upgrades for this client.

Client Additional Options:



Paths to restore:

To simplify restore, the user can be offered a list of paths to choose where to restore the required data. You can define this list here.

MS Windows VSS Setup:

Define for each Windows client's local NTFS file system whether a virtual shadow copy (VSS) should be

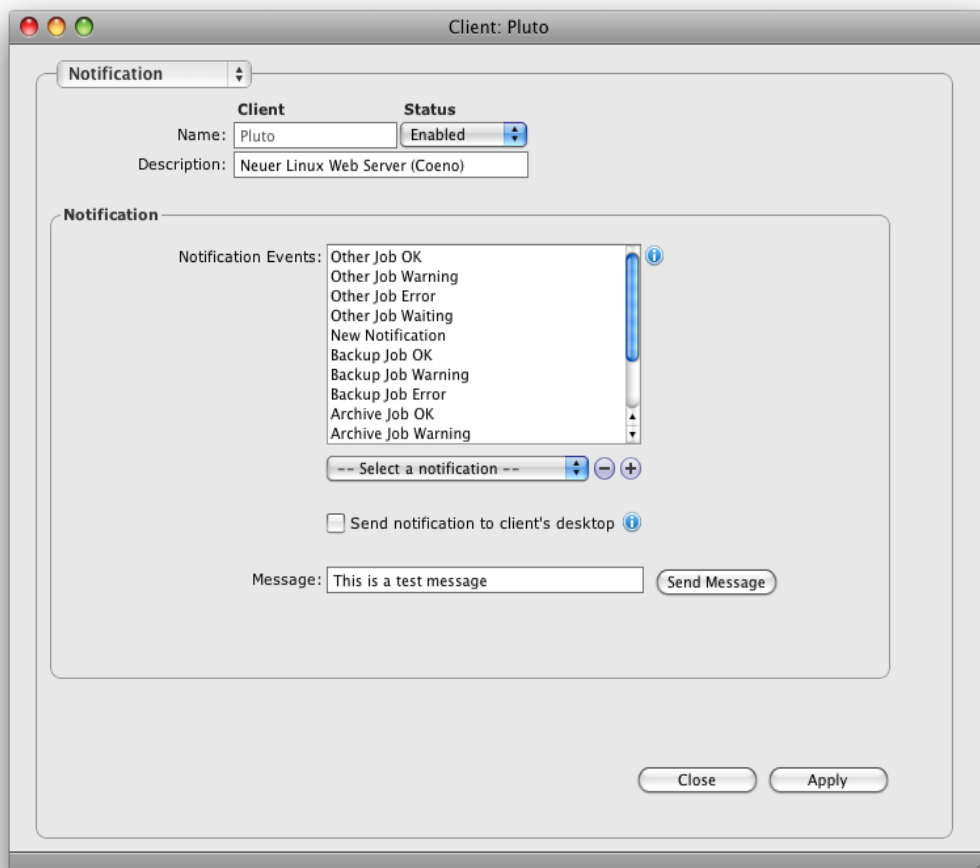


created during processing. For each corresponding drive letter define an unused drive letter to map it to. P5 will then create a virtual shadow copy for that drive and access it with the defined mapped drive letter.

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.

Windows workstation (XP, Vista, 7) allow only one shadow copy at a time.

Client Notifications:



In addition to the e-mail notifications available in previous versions, P5 notifications can now be sent to Mac and Windows hosts. On Macs, Growl must be configured to receive notifications. On Windows, the P5 Server manager must be active to show notifications as balloon tips.

To enable notifications for the selected client, check **Send Notifications to client's desktop** and select the notification events that you desire to receive.

For each P5 client, a different set of events can be defined. 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 (than 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 is set to suspect (due to an error)
The notification	Is sent on a B2G Workstation when:
Workstation Backup Job Ok	A Workstation Backup job ends OK
Workstation Backup Job Warning	A Workstation Backup job ends with warning
Workstation Backup Job Error	A Workstation Backup job ends with error
Workstation Restore Job Ok	A Workstation Restore job ends OK
Workstation Restore Job Warning	A Workstation Restore job ends with warning
Workstation Restore Job Error	A Workstation Restore job ends with error



6.12.2 Client: Interactive Functions

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

6.12.2.1 *Ping*

Ping tests whether a P5 client is available.

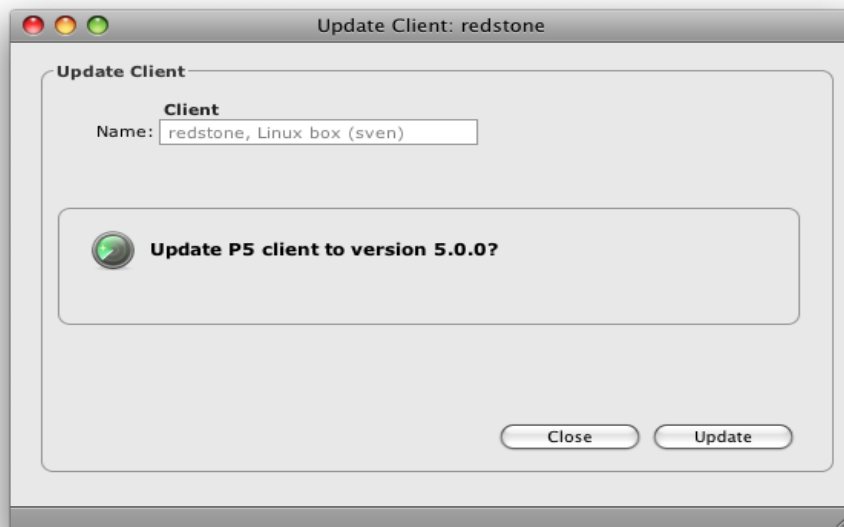
A click on ***Ping*** starts the test. The result is displayed at the bottom of the ping window. A check is made whether the client's version is compatible with the server.





6.12.2.2 Update Client...

Running an update will force the P5 version on the client to correspond with that of the server. Any jobs involving this client will be terminated, which can lead to a loss of data. It is therefore recommended to check for running jobs employing the client before initiating an update. Clicking **Test** will check if the client can be reached



Note that during the update, the client is restarted. During the restart, the P5 server will lose the connection to the client.

In case automatic update is selected in the configuration for this client, please note that only required updates are applied, patch updates may be skipped by the automatic update.

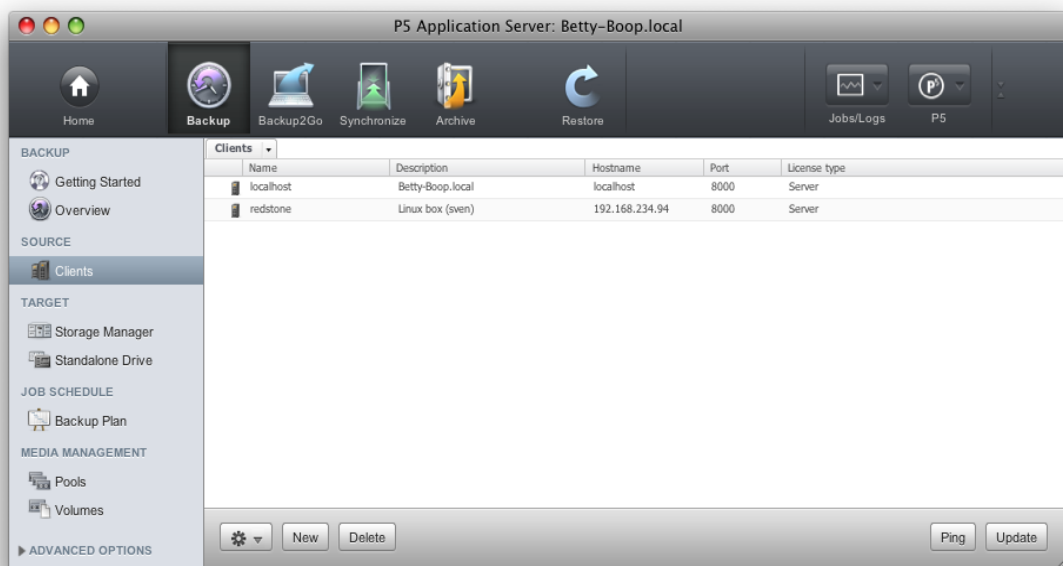


6.12.3 Setting up Encryption for a Client

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, setting it up 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 log in 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 the P5 server, use the default *localhost* entry.



This is the only set up that is done on the client. Entering a client license and configuring a client is done on the P5 server.



6.13 The Backup2Go Server

The Backup2Go Server and Client modules offer a new backup system that is designed to backup workstations that are not permanently available for backup. The generic template defines a scheme on the server that allows to entirely skip configuration if desired and makes it trivial otherwise.

This section describes the setup on the Backup2Go server.

6.13.1 Editing the Backup2Go Template

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the *Backup2Go* module from the top menu
Select *Workstation Groups* from the sidebar .
3. Double click the generic template to edit the template.

The screenshot shows a configuration window titled "Template: 1000". It contains several sections:

- Backup Configuration:**
 - Description: Generic template
 - Status: Enabled
- Server storage options:**
 - Data storage directory: /Volumes/B2Go/Generic_template
 - Retain data for: 2 weeks
- Pre-defined workstation parameter:**
 - Backup options:**
 - Directories: /Users, /C/Users
 - Absolute path: (empty)
 - File filter: workstation
 - Backup interval: 4 hours
 - Use file-system events: - workstation decides -
 - CPU throttle: - workstation decides -
 - On battery power: Backup is disabled
 - Communication options:**
 - Compress data transfer: - workstation decides -
 - Encrypt data transfer: - workstation decides -
 - Limit bandwidth: - workstation decides -
 - Network mask: 192.168.234.255

Buttons for "Close" and "Apply" are located at the bottom right.



Each time a new workstation connects to the Backup2Go server, a new record is created on the server and linked to the *Generic template*. When the generic template is *Disabled*, the server will not accept new workstations. Further templates can be created to group workstations with different setups.

Description and Status:

This section holds the preset template name and the status. For user defined templates, a meaningful name can be set.

Server storage options:

This section holds parameters that can be predefined on the server. These parameters are used from the template to which the workstation's record is linked to when created. If you wish to modify them individually for a particular workstation, you can break the link, in which case at that point, the parameters are copied into the workstation's record.

Data storage directory:

This entry defines the path that is used as the place for storing data on the Backup2Go server. This field is mandatory. The data backed up from the workstations is saved into this directory.



To take advantage of the full power of P5 Backup2Go, use a file system that supports snapshots for storing the workstation's data on the server.

P5 supports

On Mac OSX Journaled HFS+ file systems.

On Solaris ZFS (Zetabyte file system) with snapshot support
Native Unix file systems (*)

On Linux Btrfs (Btree FS) with snapshot support
ZFS (Zetabyte file system) with snapshot support
Native Unix file systems (*)

On Windows NTFS (VSS snapshots are not supported)

(*) Native Unix fs includes all file system on Solaris and linux that support:
Soft- and Hardlinks, user level extended attributes and UTF-8 character sets

As a rule of thumb, when there are more than 5 million files backed up over all workstations and snapshots, the repository should be based on ZFS or btrfs for acceptable access speed

Retain data for:

Define for how long the workstation data must be kept on the server. This value is disregarded on file systems that cannot support snapshots, where the data is never removed.



Predefined workstation parameters:

This section holds parameters that can be predefined on the server. These parameters are used from the template to which the workstation's record is linked to when created. If you wish to modify them individually for a particular workstation, you can break the link, in which case at that point, the parameters are copied into the workstation's record. Leaving these fields empty either in the template or in the individual record on the server, allows to define each of the parameters individually on the workstation. Parameters set on the server always override those of the workstation.

Directories:

This list holds all the paths of the workstation that are to be saved by Backup2Go. If left blank, the workstation user can define what is to be saved.

Absolute path:

Use this field to add or delete entries in the list of directories.

File filter:

Specify here an optional filter to include or exclude specific files and/or directories to or from the backup. See section 6.9, Filters, for more information about filters. There is no filter defined by default.

Backup Interval:

This parameter defines the interval between two backups. The initial backup is started after the interval elapses for the first time after registering the workstation with the server. In case a backup could not be started on time, the next backup is started as soon as the workstation is able to connect to the server again.

Use file-system events:

With *Yes* or *No* choose to use file system events on the workstation.

CPU throttle:

Use this control to limit CPU load used by the workstation for Backup2Go. Altering this setting may also reduce the I/O load on the workstation. 100% is set by default for no limit.

On battery power:

Use this switch to define whether or not the backup is enabled while running on battery power..

Compress data transfer:

With *Yes* or *No* choose to use transfer data compression. When turned on, data is compressed during transfer over the network. This reduces the required line bandwidth, but produces more CPU load both on the workstation as well as on the server.

Encrypt data transfer:

With *Yes* or *No* choose to use transfer data encryption. When turned on, data is encrypted during transfer over the network. This produces more CPU load both on the workstation as well as on the server, but enhances the security of transfers over public lines.



Limit bandwidth:

Use this control to limit the bandwidth used by the workstation to transfer data. Altering this setting may also reduce the I/O load on the workstation. 20% is set by default.

Network Mask:

A Network mask can be specified to limit workstation access only from the given network. Specify the network address, set the value 255 for "all addresses". For instance 192.168.1.255 allows access for all workstations which IP address starts with 192.168.1.

Note that for simplicity, the network mask works byte-wise only. Sub-byte values cannot be masked.



Expert settings:

Server communication parameters:

This section defines the parameters used for remote workstation configuration:

DNS Name:

Define the IP address or the DNS resolvable name of the Backup2Go server the workstations should connect to.

Port:

Define the port of the Backup2Go server the workstations should connect to. The default value is 8000.

Time window to accept Backup:

This section allows to limit the backup time. Use this in order to stop the workstation backup while for instance a backup to tape of the Backup2Go repository shall be done.

Start/end hour:

Set the start and end time as hour of the day.



Adjusting Server load:

This section defines the parameter used to limit the server load caused by Backup2Go.

Maximum active workstations:

This parameter allows to limit the number of the currently active workstations. The default value is -1 for no limit.

Limiting the number of workstations may be required in case too many workstations access the server simultaneously and drives the server I/O or CPU load to a value that brings down the overall performance.

Verify data integrity:

The workstation index is maintained on each workstation and is required to calculate which files are to be saved. This index is rebuilt automatically from time to time in order to ensure the index's integrity. Do not change these values if unsure.

Maximum index creators:

This parameter allows to limit the number of workstations that can rebuild their indexes simultaneously. The default value is 1, thus allowing one index to be rebuilt at a time.

Wait for free index slot:

This parameter defines how long a workstation will wait until it may rebuild its index. The default value is 120 seconds.

Network speed limit for regenerating index:

This parameter limits the network speed used to transfer the index to the workstation. The default value is 1024 KB/sec.

Automatic rebuild Workstation index:

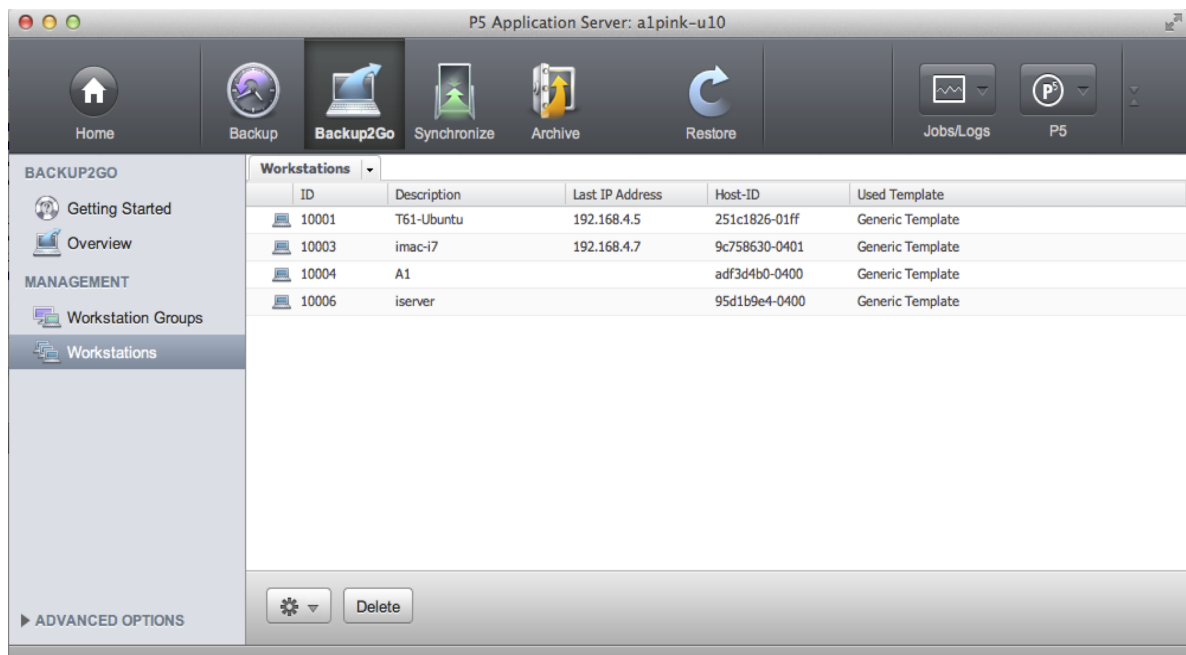
The workstation index should be rebuilt in order not to risk integrity loss between the saved data and the index reflecting which data has been saved. The default value is weekly.



6.13.2 The Workstation Manager

When a workstation connects to the Backup2Go server for the first time, a new record is created for this workstation using the current backup template.

The workstation manager shows all the workstations and displays the time and the result of the most recent backup.



The workstation manager allows to edit, delete or duplicate existing workstation records as well as to create new records.



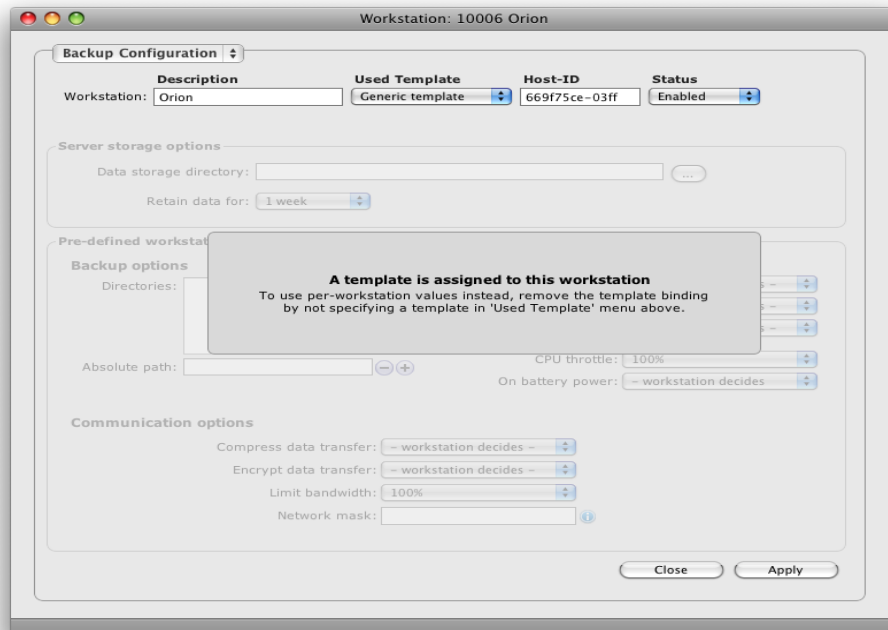
Deleting a workstation record will permanently disable the corresponding workstation from connecting to the server.

Do not delete the record of a workstation that you plan to use with Backup2Go later, since it will otherwise not be able to reconnect to the Backup2Go server when the time comes.

6.13.3 The Workstation Record

A workstation record is created and linked to the generic Backup2Go template when a workstation connects to the server for the first time. By default, the record remains linked to the template. This means that any changes in the template automatically affect the workstation.

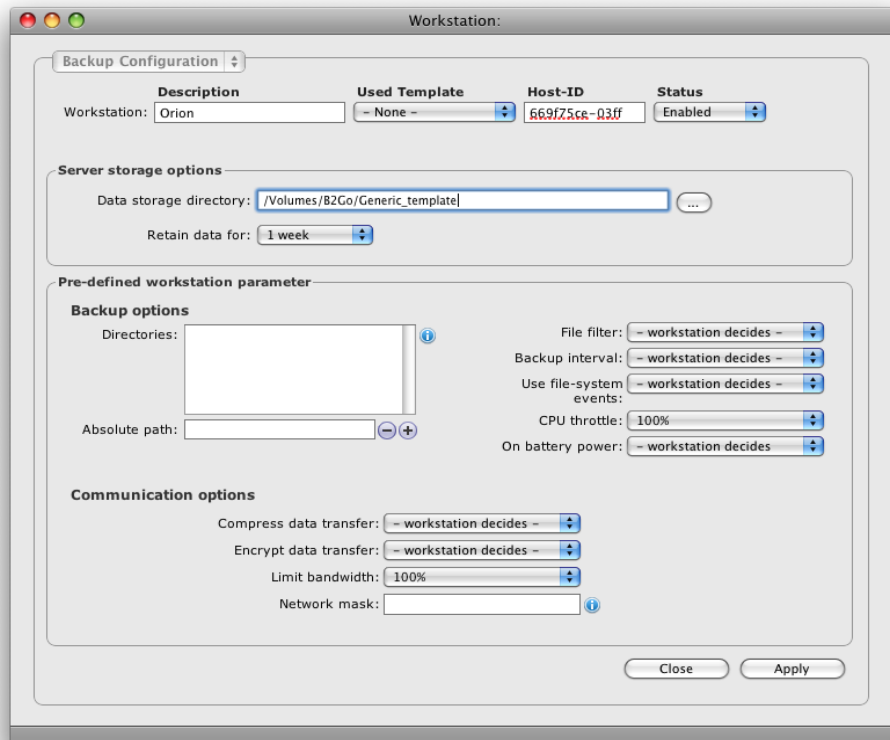
As long as the template remains assigned to the workstation record, the record is locked against changes as indicated by the following screen:



In this state, only the fields of the workstation configuration can be edited.



In case the workstation configuration is decoupled from a template, values can be set for the workstation separately:



Workstation Configuration:

These parameters are unique for each particular workstation:

Description:

This field describes the workstation. In case the workstation record was generated from the template, the workstation's host name appears here by default. You can replace this entry with a more descriptive name .

Used Template:

This field shows the currently assigned template. As long as there is a template assigned, the individual parameters are locked against editing and the values from the template are used instead. To break the link to the template in order to set individual values for the particular workstation, clear this field. You may want to define multiple templates in order to group workstations with common and different settings. You can change the template assignment here.

Host-ID:

The Host-ID is a unique identifier of each workstation. This value is required only if the workstation has been replaced or if it has lost its local configuration. In case a new workstation has to access data



stored by another workstation, the host-ID must be edited to fit to the value visible on that workstation's main screen.

In normal operation, a workstation is identified only by the ID of the workstation's record on the server. This ID is stored in the workstation. This way, P5 can identify the workstation even with a different network address or Host-ID, for example when the workstation connects from a different network interface.

Status:

Workstation records are enabled by default. Disabling the workstation record causes the workstation's backup to fail since the workstation cannot connect any more.

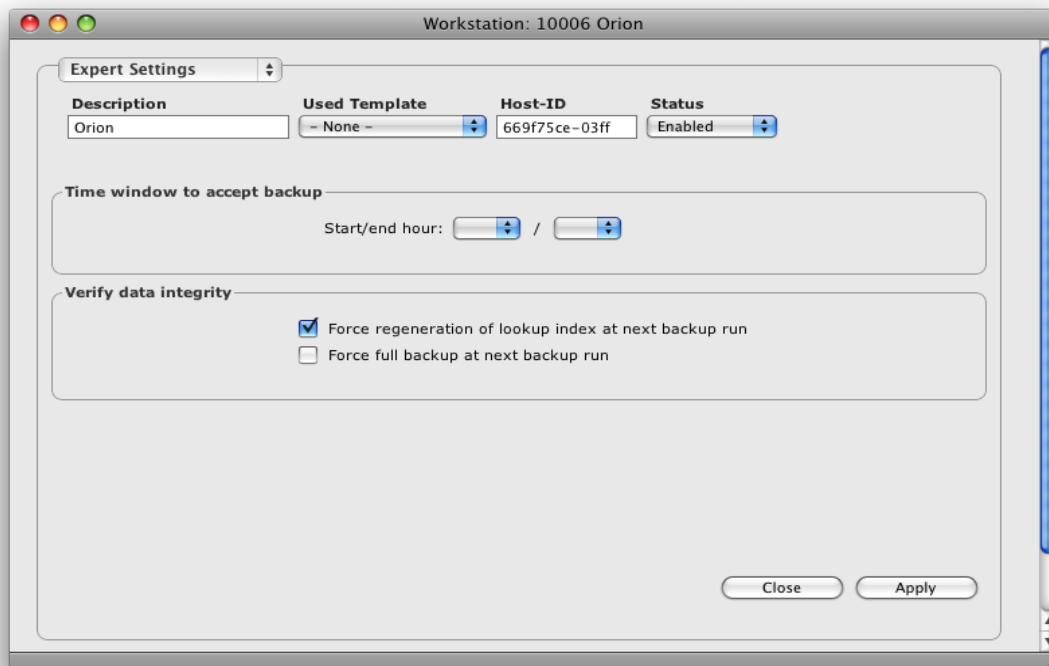
Server storage options:

This section holds the storage parameters on the server. See section 6.13.1, Editing the Backup2Go Template for a detailed description.

Pre-defined workstation parameter:

This section holds the individual parameters for the particular workstation. See section 6.13.1, Editing the Backup2Go Template for a detailed description. Values assigned in the template override values assigned in the individual workstation records on the server, which in turn override the values defined on the workstation.

Expert Settings:





These parameters are unique for each particular workstation:

Time Window to accept Backup:

Specify start and end time in case the backup shall not be constantly available. This entry can be used to disable Backup2Go at specific times, for instance to perform a Backup of the server itself.

Force Regeneration of lookup index on next backup run:

Check this flag to make P5 recompare the files on server and workstation instead of relying on the index. This is required after changing or deleting files in the Backup2Go repository on the server.

Force Full Backup on next backup run:

Check this flag to make P5 save all files again on the next job execution.



6.13.4 The Workstation Overview

To obtain a status report of all workstations at a glance, visit the Workstation Overview

To reach the workstation overview:

1. Log into P5 using a web browser as an admin user.
See section 4.1, Starting the Browser.
2. Select the *Backup2Go* module from the top menu and select *Backup2Go Overview* in the sidebar

Backup2Go overview

Show: All Successful Other

ID	Description	Last Backup	Last Successful	Used Template
10006	iserver	Never	Never	Generic Template
10004	A1	Never	Never	Generic Template
10003	imac-i7	Okay (0 bytes)	3 hours ago	Generic Template
10001	T61-Ubuntu	Okay (0 bytes)	6 days ago	Generic Template




N/A - Information not available

3. The display shows the workstation description and the workstation group (template) it is assigned to. The Last Backup column shows the return state of the last backup execution:

- Green: The job finished ok
- Orange: The job terminated with warnings.
- 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 ok the last time:

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



6.14 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, in which all the archive indexes will be maintained 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.

6.14.1 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 **Advances Options** sub-menu.
3. Click **New** from the bottom bar.
4. Fill in the fields of the Database dialog window.

Database Configuration: Default-Archive

Database Configuration

Database	Status
Name: Default-Archive	Enabled
Description: Default archive database	

Type and Usage

Driver type: AixTree
Usage: Archive

Directory

conf/index/archive/Default-Archive

Close Apply

**Database Name:**

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

1. Uniqueness: No other database 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 database here.

Status:

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

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

Driver type:

Holds the type of the tree driver, which will be applied to manage this database.

Usage:

Type of the data that this database references.

Directory Path:

Enter the path in the file system where this database should be created. You may also simply accept the suggested value.



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 6.15.1 Creating a Log-In Area .

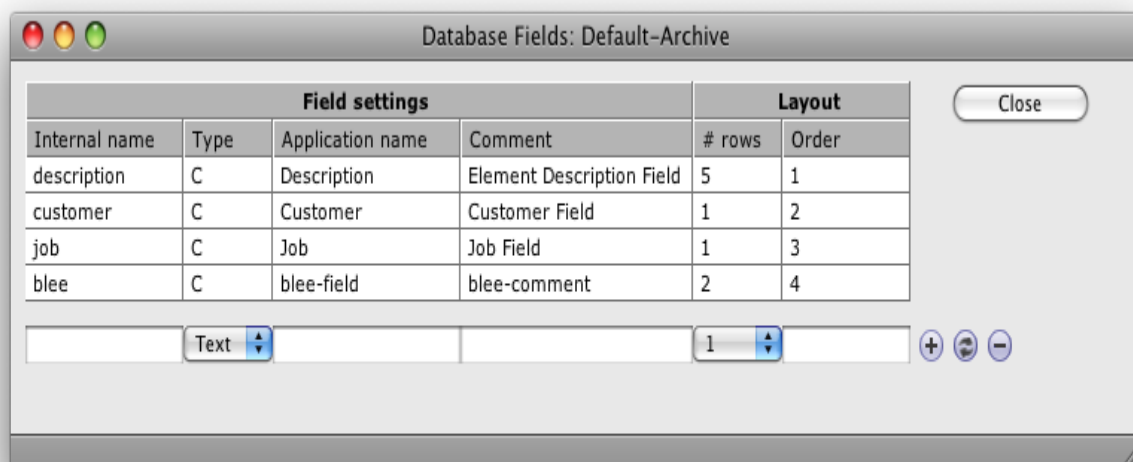


6.14.2 Database: Functions

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

6.14.2.1 *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).

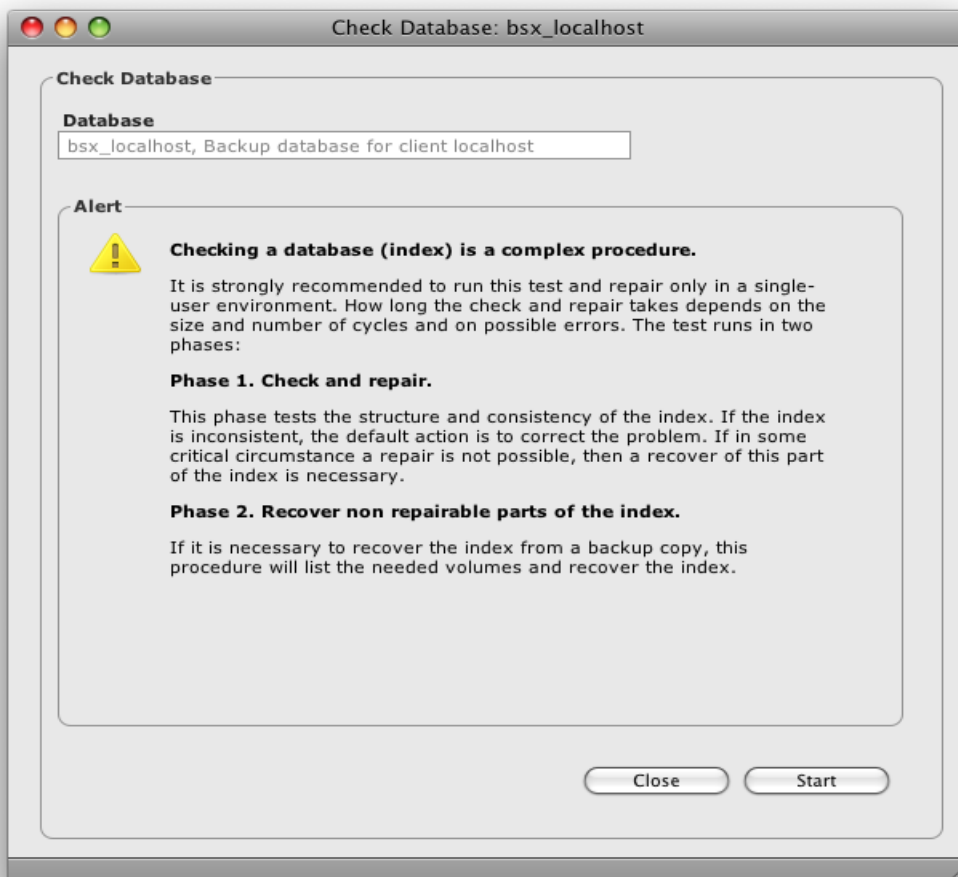


6.14.2.2 Check

This function is used to check the selected backup or archive databases, which means the backup indexes of a given client or of a given archive. 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 recover 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 that could not be corrected by the check index function, cannot be recovered automatically, since it does not automatically get saved. It is therefore recommended to back up the archive indexes by means of a backup plan, so that in the event of a corruption, you can recover it from the backup.



6.15 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 disc. 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. Authorizing is done via standard user/group mechanisms.

6.15.1 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 **Advances Options** sub menu
3. Click New from the bottom bar.
4. Fill in the required fields.

Login Area Configuration:

General Setup

Description **Status**

Media2011 Enabled

Index or File System

Selection: Archive index: Default-Archive

Path: /Media/2011/ (e.g. /data/raid1/vol1)

Access

Allow access for group(s): Archivers

Group name: a valid OS group

Close Apply

**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:

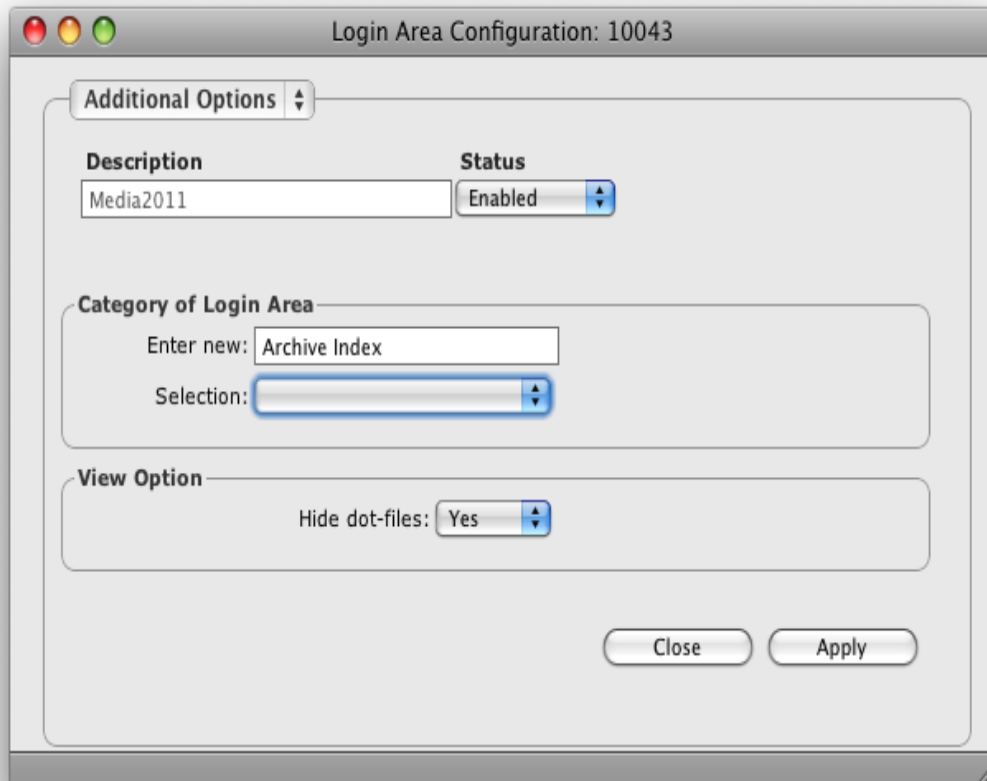
The database you choose contains a directory structure. Enter the path in this structure, i.e. the point from which you wish to enable access into it. Should this be the whole database, then enter '/'.

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.



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.



6.16 User Preferences

P5 doesn't manage the users itself. Authorizing is done using standard operating system mechanisms of the server. From the user preferences area, you can allow or deny the user certain functions, 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.

6.16.1 Managing User Preferences

1. Log into P5 using a web browser as an 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 *Advances Options* sub-menu.
3. Click **New** in the bottom bar
4. Fill in the required fields.

User Setting: sven

General Setup

User Name: sven Status: Enabled

Description: Support

Login

Allow log in

Allow local user to log in automatically

Login Behaviour

Open Login area:

Operation

Allow user to restore

Allow user to archive

Close Apply

A user preference is created automatically when a new user logs in.

***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.

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.

Allow log in:

Check this box to allow that user to log in into the P5 interface.

By disabling this check box in the **generic** user's preferences, new users will not be allowed to log in

Allow local user to log in automatically:

Click this check box to skip the authentication procedure and automatically log in the user, provided he is working with a browser that is running directly on the machine. For safety reasons, automatic log in has no effect from other hosts.

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.

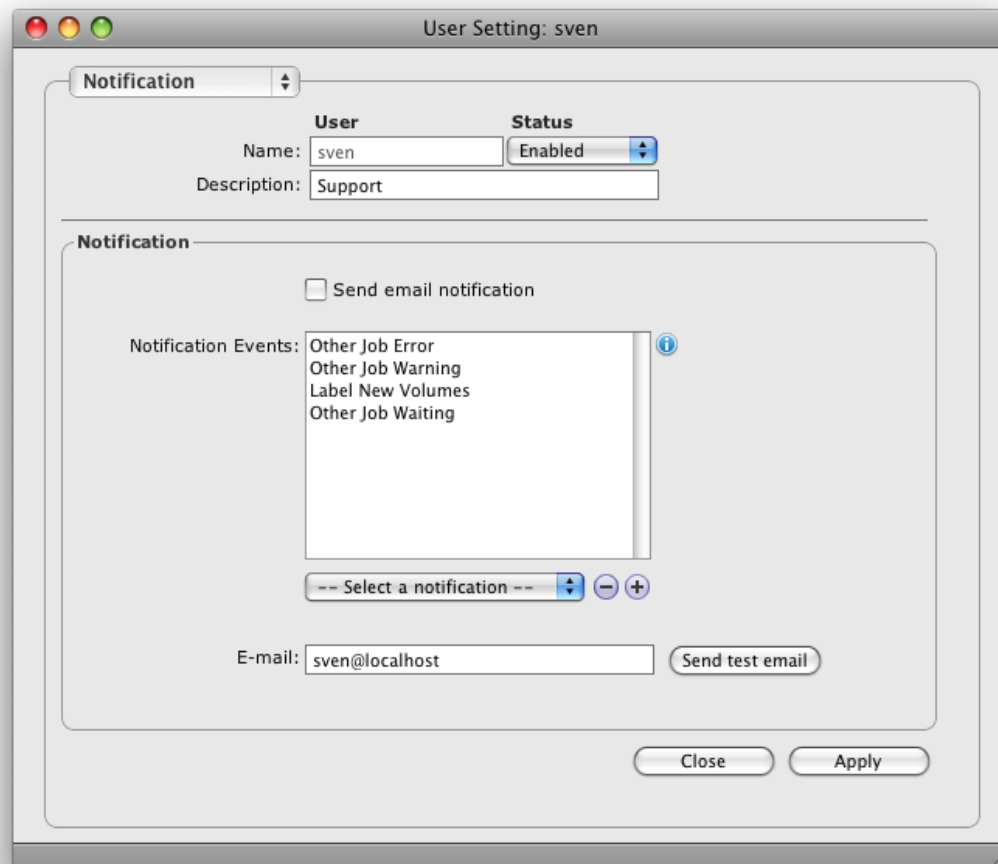
Allowed operations:

The function can be restricted individually for each user. The restriction is always associated with a log-in area .

- » **Restore:** Allows the user to restore data from the backup or the archive.
- » **Archive:** Allows the user to archive on-line data.



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 Email notifications

Check this box in order to enable E-mail notification for the given 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:

The users E-mail address

Send test E-mail

Use this button to send a test e-mail to the configured user.



In order for e-mails to be sent correctly, a valid mail server must be defined in the setup of the *localhost* client. P5 does not support mail server authentication mechanisms. In order to send mails via a mail server that requires authentication, you may want to setup a local mail server like PostFix or SendMail without authentication to forward the mail or to use Growl's mechanisms on Mac.

As an alternative or in addition to e-mails it is possible to setup Growl notifications on Mac hosts or bubble notifications on Windows hosts. See section 6.12.1, Creating a Client for more.


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 (than 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 is set to suspect (due to an error)

The notification	Is sent on a B2G Workstation when:
Workstation Backup Job Ok	A Workstation Backup job ends OK
Workstation Backup Job Warning	A Workstation Backup job ends with warning
Workstation Backup Job Error	A Workstation Backup job ends with error
Workstation Restore Job Ok	A Workstation Restore job ends OK
Workstation Restore Job Warning	A Workstation Restore job ends with warning
Workstation Restore Job Error	A Workstation Restore job ends with error



Users preferences – Additional Options



The image shows a dialog box titled "User Setting: sven". It has a tab labeled "Additional Options". The dialog is divided into three sections:

- User**: Contains a "Name" field with the value "sven" and a "Status" dropdown menu set to "Enabled". Below these is a "Description" field with the value "SUPPORT".
- Display Option**: Contains a "Date/Time" dropdown menu with the value "31.Oct.2001 17:23" and the text "(Date/Time format)" to its right.
- Restore Option**: Contains a "Fixed restore path" field with a browse button ("...") to its right, and a "to client" dropdown menu with the value "localhost".

At the bottom right of the dialog are two buttons: "Close" and "Apply".

Date/Time Format:

Offers a list of possible date/time formats to choose from



Restore Option, path and client:

These fields define the default values for the location where to restore data.



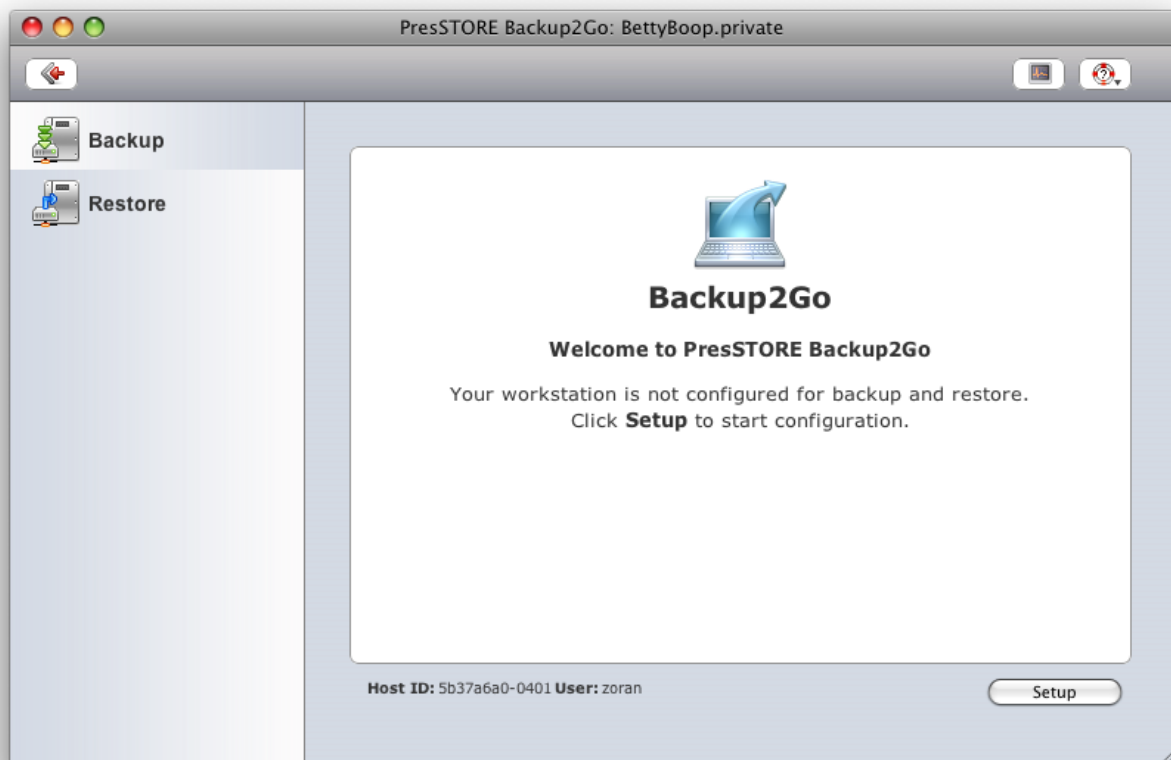
7 Configuring the P5 Backup2Go Workstation

This section describes the Backup2Go Workstation interface. In contrast to the P5 client described in section 6.12, the Backup2Go workstation requires to be set up and has it's own browser interface.

-  You have to install a copy of P5 on each machine that is to act as a P5 Backup2Go workstation. No license is required on the workstation.
-  To configure Backup2Go, please use Safari, Mozilla Firefox or Microsoft Internet Explorer. For more information see section 3.2, System Requirements.

Enter the workstation's user URL <http://localhost:8000/workstation> into the address field of the browser. The first dialog will prompt you to enter the *User name*: and the corresponding *Password*:

Please log in with administrator privileges to proceed with the Backup2Go workstation setup:





7.1 The Workstation Setup Wizard

Click *Setup* to configure the workstation. When no setup is present, the setup wizard starts automatically and takes you through the required steps:



To set up the workstation for backup, please choose **Backup**.



The following mandatory fields are shown:

**Server name:**

Please enter the name or the IP address of the Backup2Go server. In case a server name is entered, please make sure the name can be resolved into an IP address through a DNS server or the workstation's */etc/hosts* file. P5 supports the Bonjour service to locate the Backup2Go server.

Communication port:

This is the TCP/IP port used by the Backup2Go server. By default, this is 8000. Do not change this value unless the port has been changed on the Backup2Go server.

User name / Password:

Enter the user name and password of an account on the server. The given account must have permission to log in to the server.

Provided that the backup administrator predefined the setup through server settings, further options are not required.

7.2 Manual Workstation Setup

Once the initial setup has been done, the manual configuration screen can be reached by clicking **Setup**. The Setup Wizard can be reached through a button in the bottom bar.

The **Backup Server** tab shows the server settings as entered in the Setup Wizard:

The screenshot shows a window titled "Backup Server" with three tabs: "Backup Server", "Data to Back Up", and "Advanced". The "Backup Server" tab is selected. The main area contains the following text and form elements:

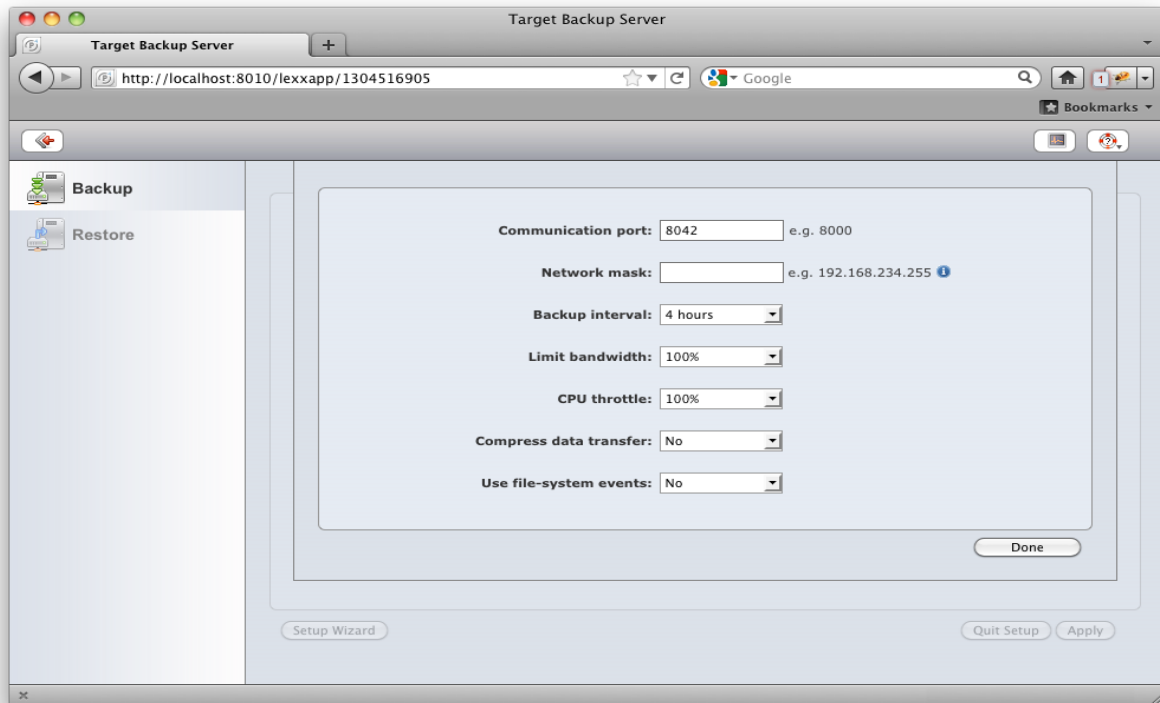
- Header: "Enter the address and login details for your backup server:"
- Field: "Server: 192.168.1.200" with a sub-label "DNS name or TCP/IP address" and a "Advanced Options" toggle.
- Field: "User name: [encrypted]"
- Field: "Password: [masked with dots]" with a sub-label "User name and password to login to the backup server"
- Bottom bar: "Setup Wizard" button on the left, and "Quit Setup" and "Apply" buttons on the right.

See section 7.1, The Workstation Setup Wizard for a description of the values.

The **Advanced options** allow you to change the following values. Note that values preset by the backup



administrator cannot be set on the workstation.



Communication port:

Set here the TCP/IP port of the Backup2Go server. The default value is 8000.

Backup Interval:

An attempt is made to back up each workstation in regular intervals. This parameter defines the interval between two backups. The first backup is started once the interval time elapses for the first time upon registering the workstation with the server. In case a backup could not be started on time, the subsequent backup is started as soon as the workstation is able to connect to the server again.

Limit bandwidth:

Use this control to limit the bandwidth used by the workstation to transfer data. Altering this setting may also be used to reduce the I/O load on the workstation.

Reducing this value means that

- the backup will be less noticed since it consumes less workstation resources.
- the backup may take longer.
- there is more network bandwidth left free to access resources from other servers, for examples files on shares.

**Network Mask:**

A Network mask can be specified to limit workstation access only from the given network. Specify the network address, set the value 255 for "all addresses". For instance 192.168.1.255 allows access for all workstations which IP address starts with 192.168.1.

CPU throttle:

Use this control to limit the CPU load used by the workstation for Backup2Go. Altering this setting may also reduce the I/O load on the workstation. 100% is set by default for no limit.

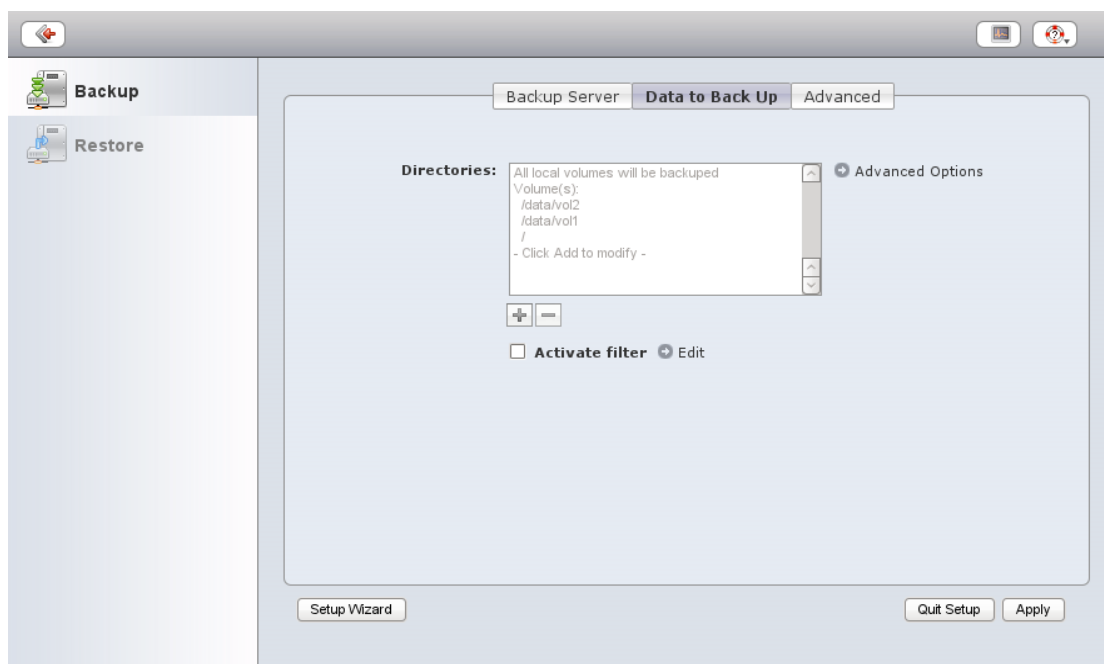
Compress data transfer:

With *Yes* or *No* choose to use transfer data compression. When turned on, data is compressed during transfer over the network. This reduces the required line bandwidth, but produces more CPU load both on the workstation as well as on the server.

Use file-system events:

With *Yes* or *No* choose to use file system events on the workstation.

Select the *Data to Back Up* tab to specify which directories are to be included in the backup:

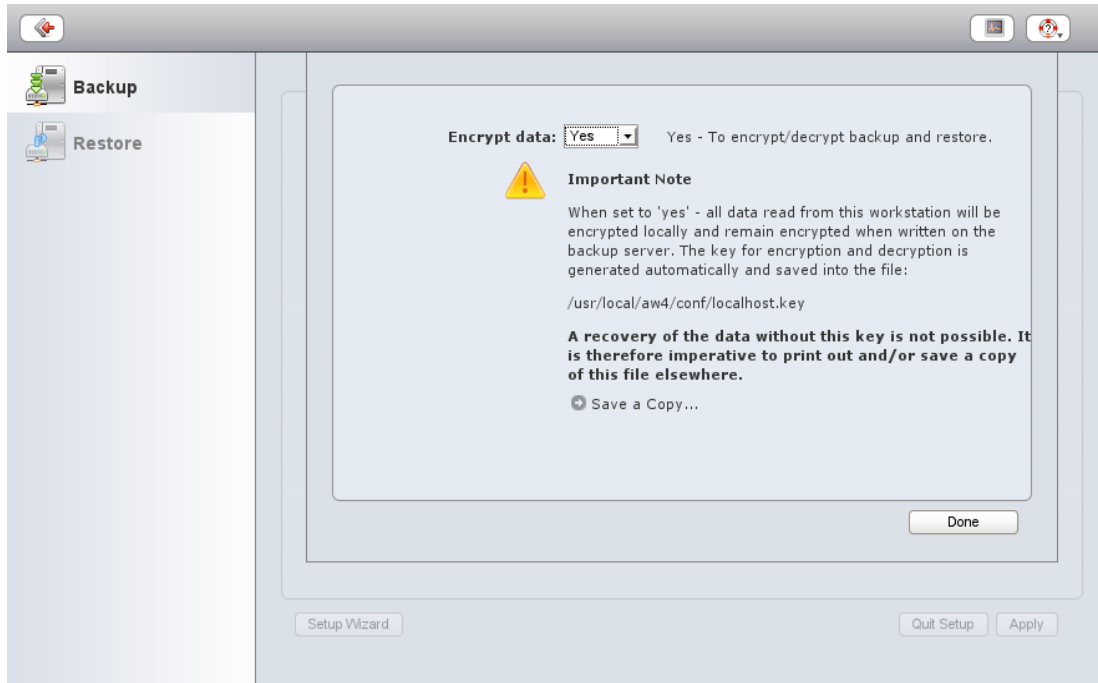


If not overridden by the server's administrator, additional directories can be added with **+** and deleted with **-**. Please note that P5 Backup2Go does not follow mount points. In order to backup mounted volumes, they must be specified separately.

In order to activate a filter to include or exclude specific files during backup, select *Activate Filter* and click *Edit* to setup the filter details. See section 6.9, Filters for details about the filter functionality. Click **Apply** to save your changes. Click **Quit Setup** to get back to the main screen.



The **Advanced** options allow to change the encryption:



Encrypt data:

With *Yes* or *No* choose to use data encryption. This option encrypts all the data leaving this workstation and is stored on the Backup2Go server in encrypted form.

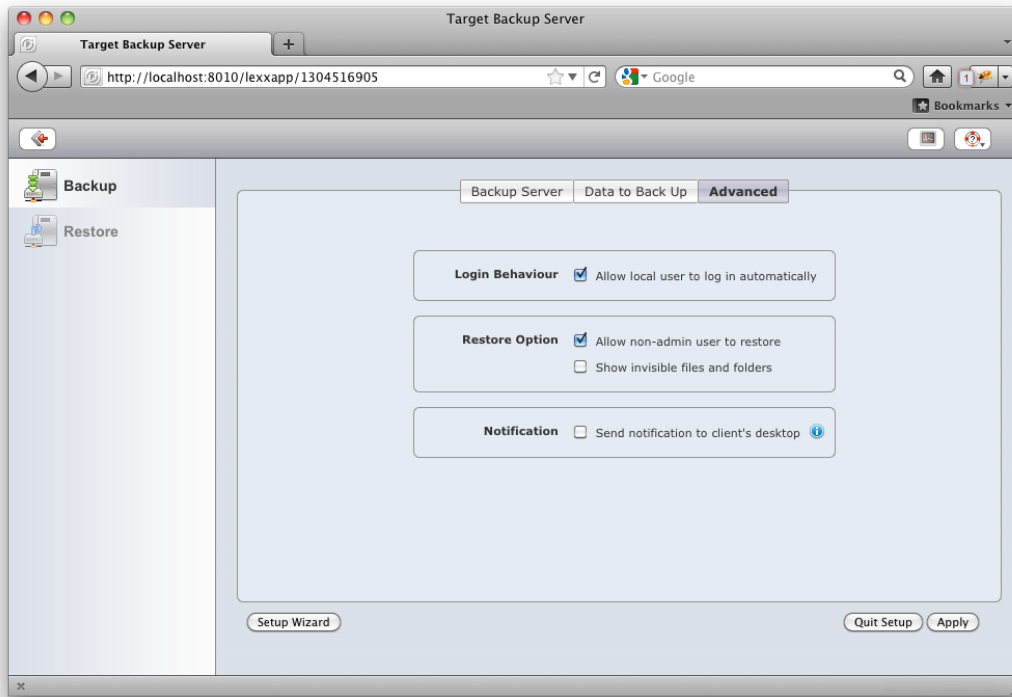


The key for encryption and decryption is generated automatically and saved in the file `/usr/local/aw/conf/localhost.key`

A recovery without this file is not possible. Please print out the key or keep a copy of the file in a safe place.



Select the *Advanced* tab to specify user permissions:



Login Behavior:

Check ***Allow local users to log in automatically*** to allow the user logged in to the local host to access P5 without log-in authentication.

Restore Options:

Check ***Allow non-admin users to restore*** to allow access to the ***Restore*** section. By default, only administrators are permitted to restore files.

Check ***Show invisible files*** and folders to show also files normally hidden by the operating system on the restore screen.

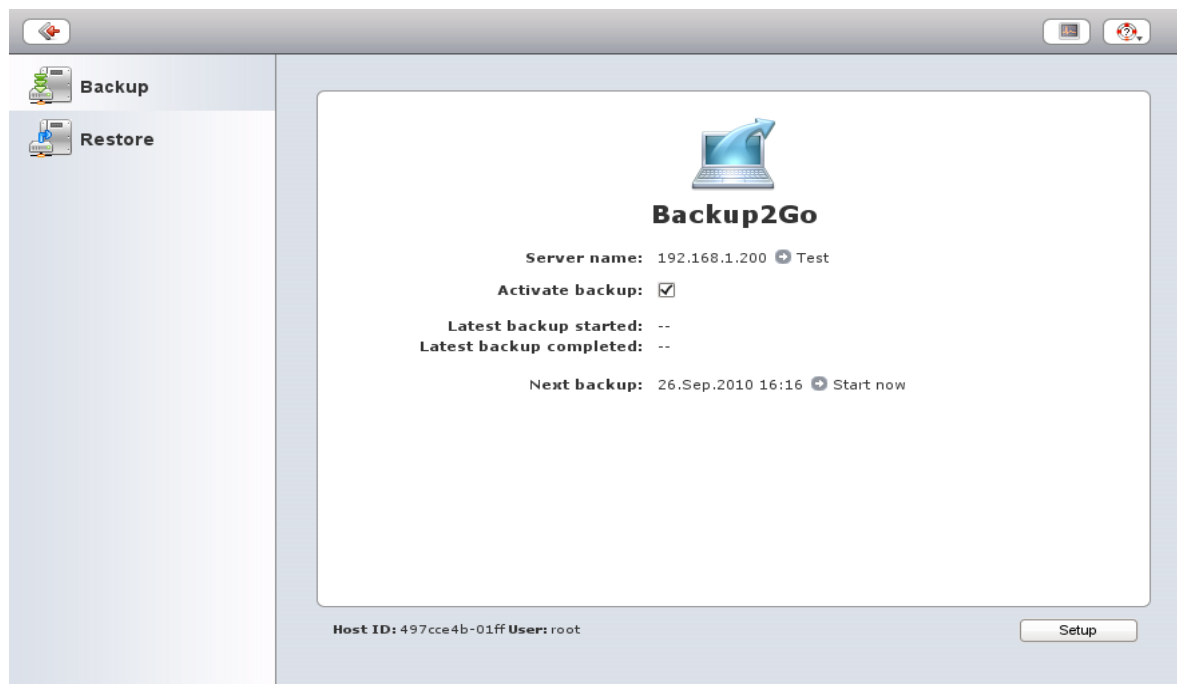
Notification:

Check ***Send notification to client's workstation*** to receive Growl notifications about the Backup2Go jobs on this workstation.



7.3 The P5 Workstation Main Screen

Once set up and activated, the workstation's main screen displays the current state of the backup.



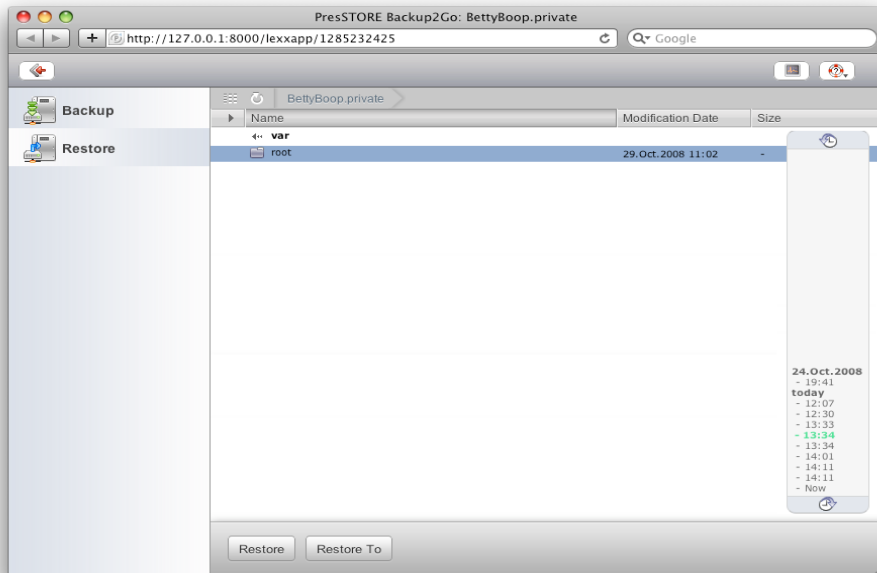
The P5 status and the dates of the last successful backup are listed here. If you are unsure if the server connection is set up properly, click *Test connection* to check. Use *Start Now* to start the backup immediately.

Please remember that administrator access is required on the workstation in order to run the setup or to manually start the backup.



7.4 Restoring Files on the Workstation

To restore data, select the **Restore** tab.



When logged in as a workstation administrator, all the workstation's directories that have been saved on the server will be displayed. A non-privileged user is restricted to his home directory, provided that it had been saved.

In case the account that is used to connect from the workstation to the Backup2Go server, the advanced options allow to select another workstation on the server to restore. The other workstation's data will then be restored to the current workstation.





8 Restoring saved files

8.1 Restoring Data From Backup

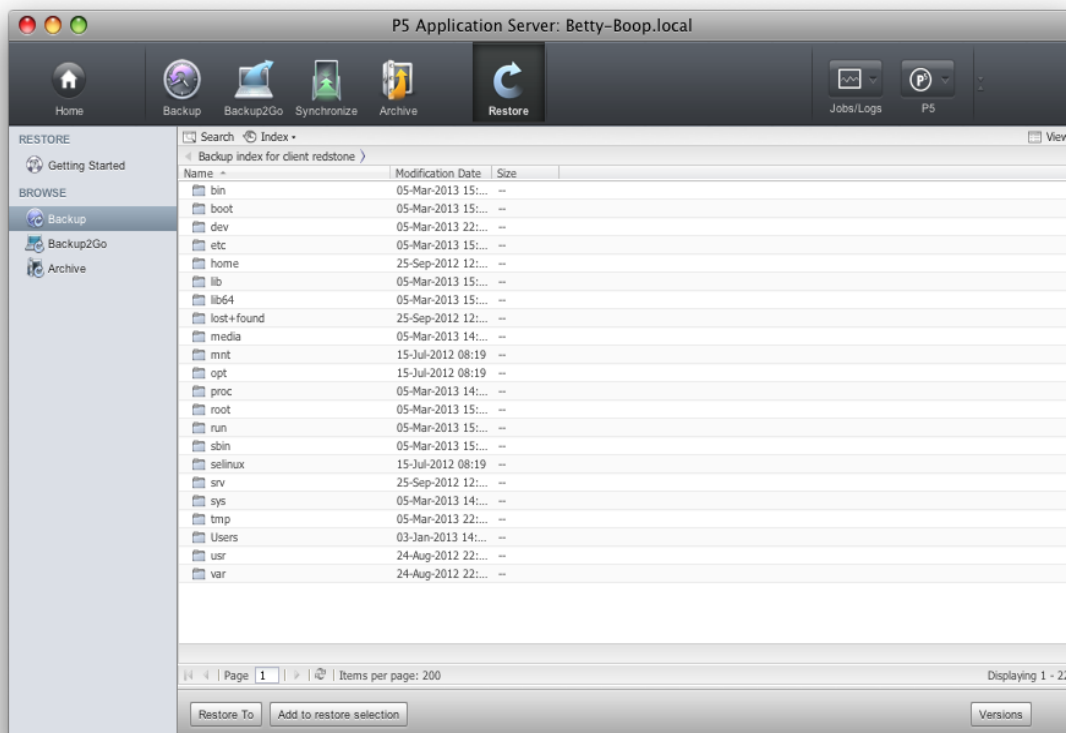
Restoring data is done through the P5 data browser. Each backed file has a corresponding entry in the index. The index tells P5 on which volume the file is located and where, and carries other meta-data describing the file. A machine (client) that was never backed up will have an empty backup index.

You can search and browse through backed up directory structures with the data browser. Data can be restored in any of the states they 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 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 is replaced with an SQLite database implementation. Indexes are no longer represented by backup cycle and run, but rather snapshots are shown 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.





8.1.1 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 6.15, 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 it has been save from during backup.

When restoring files from multiple position in the index, it may be more optimal 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.

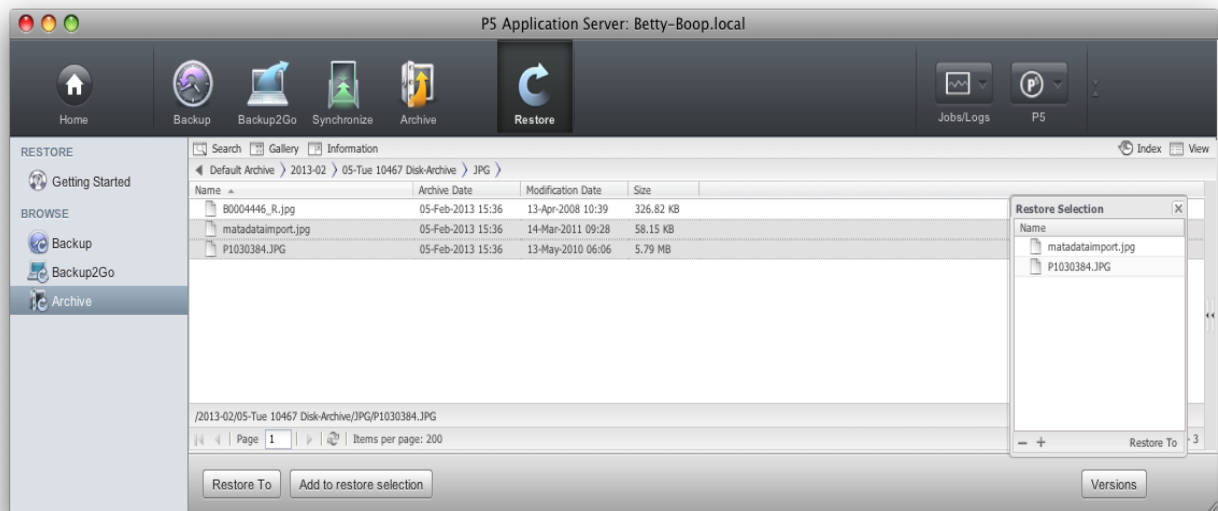


8.1.2 The Restore Selection

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 folder 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.

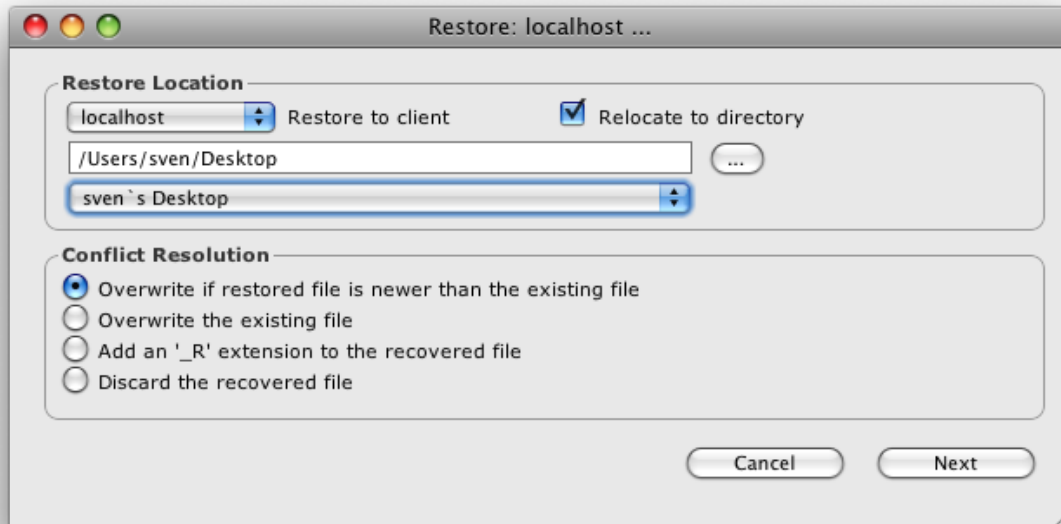
The restore selection remains active in the background until the user logs out from the P5 GUI. There is a separate Restore Selection for each index. In order to restore files from multiple clients, P5 uses a separate Restore Selection for each client, each selection must be restored with a separate restore job.





8.1.3 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 in order 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 6.16) and the defaults defined for the respective client (see section 6.12).

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.

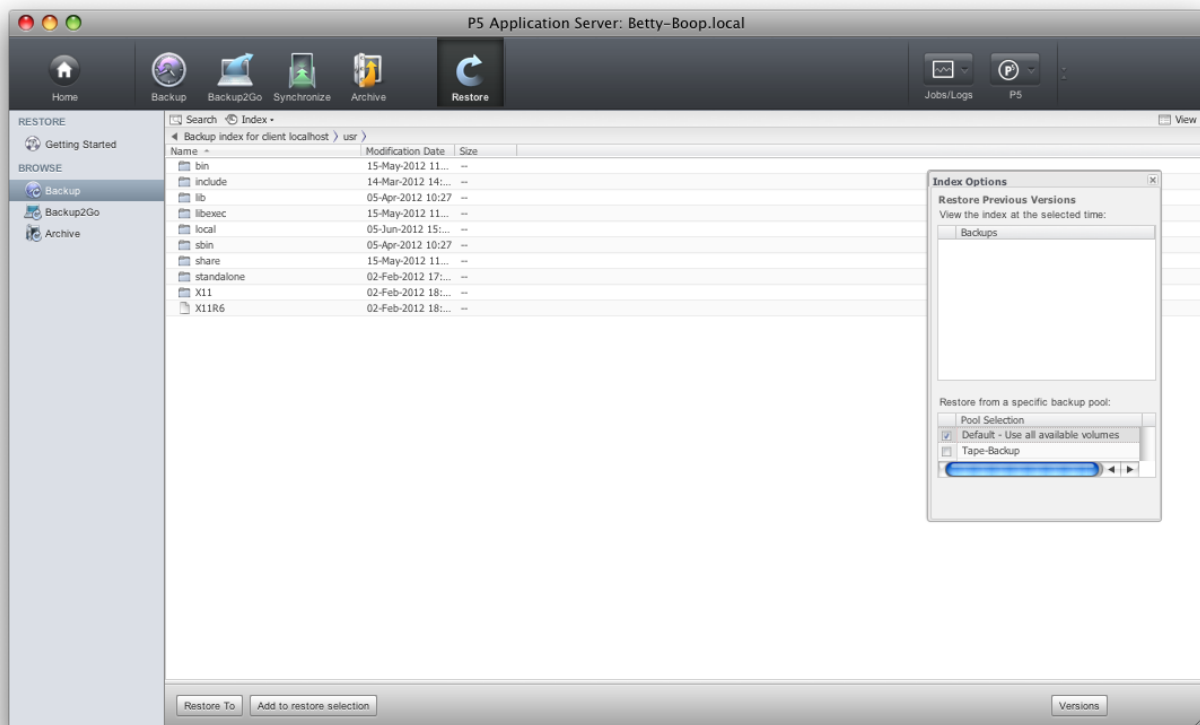
Before the actual restore takes place, you are informed about the volumes required to complete the restore operation and their availability. If all the required volumes can be found in a library that is online, P5 will manage volume changes autonomously, otherwise missing volumes will be requested and 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 6.11.1, The Job Monitor.

8.1.4 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 backup job creates a new snapshot when running 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 *Index* section in the window heading. The item *More Options* opens a sub-window to select the snapshot.





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.



Index snapshots may refer to files that were not saved when the snapshot was created but earlier, for instance when the file is unchanged and already on tape. So during a restore, tapes also can be requested that were written before the snapshot was created.

8.1.5 The Versions Button

Use the *versions* button to list all the versions of a particular 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.

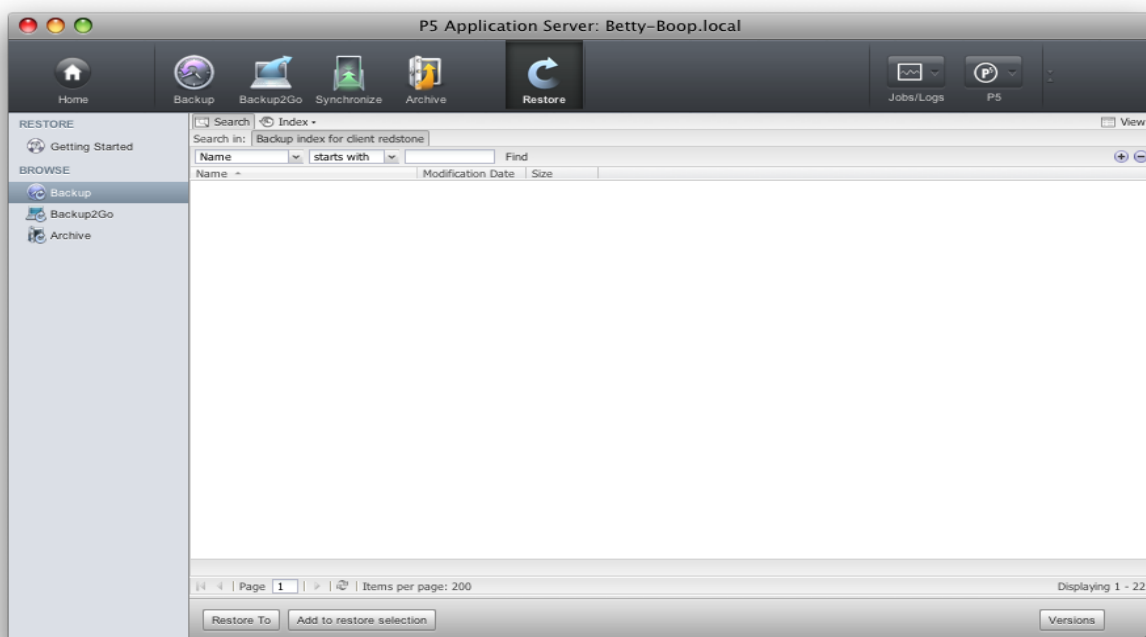
The version button is intended to select a specific version of a single files. In order to select the previous state of a folder tree, please select the appropriate snapshot.



8.1.6 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.

Next ,select what to search:

File name:

When searching for a file, enter the required name or expression here. The pull-down selection enables you to choose the matching mode.

- » starts with: match the beginning of the file name with the entered name.
- » contains: match any part of the file name with the given name.
- » matches: match the entire file name with the given name.

Modification date:

The result will contain all the data that had been modified no later than the given date.



The date format is: dd.mm.yy

It is possible to select multiple search expressions, for instance

- Name contains "Yesterday"
- Modification date is before Jan. 01 20013

The expressions must then all be met to find the file(s).

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.




8.2 Restoring data on Backup2Go Server

Data saved by a Backup2Go workstation is usually restored on the workstation itself via the workstation interface, see section 7.4, Restoring Files on the Workstation.

As an alternative, it is possible to restore the data on the Backup2Go-Server directly:

Follow these steps:

1. Log into P5 using a web browser.
See section 4.1, Starting the Browser.
2. Select **Restore** in the top menu, then choose **Backup2Go** from the left sidebar.
The main screen will show the backup workstations available to restore from.



3. As on the Backup2Go Workstation, the timeline allows to select the snapshot to restore from.



Restrictions:

In case the workstation backup has been stored encrypted, it is not possible to restore that workstations data via the server interface due to security reasons.

This restore function is implemented as a download via the browser. It is intended for quick accessing single files or folders, it is thus limited in size to 4 GB (compressed download size).



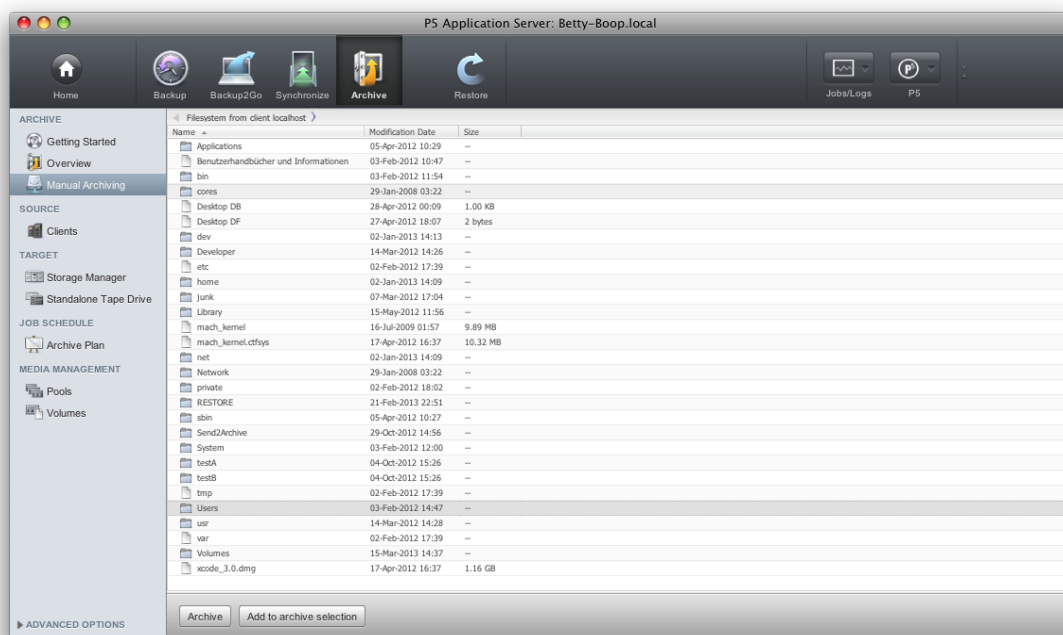
9 Archiving and Restore from Archive

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 6.6, 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 6.15, 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.

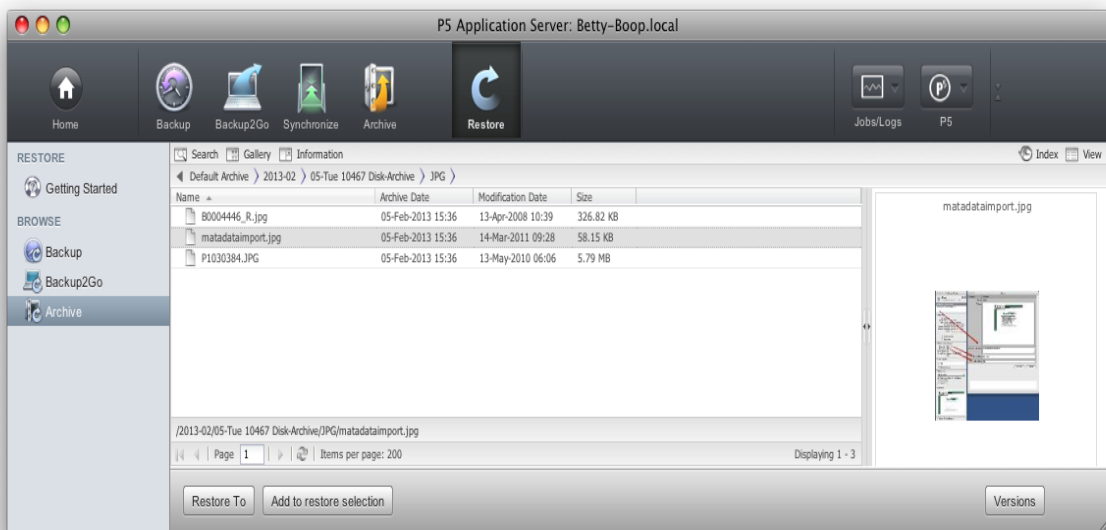
To monitor the restore process, start the *Job Monitor* from the *Jobs/Logs* section of the top menu, see section 6.11.1, The Job Monitor.



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.



You can limit the access to archive or backup indexes as needed. A log-in area for each index is required to grant or deny the access to that index. See section 6.15, 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 a second file by holding down the Shift key and clicking its name. You can select multiple elements this way.
4. 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 6.11.1, The Job Monitor.



9.2.1 The Restore To Button

Use *Restore To* 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:

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 6.16) and the defaults defined for the respective client (see section 6.12).

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 6.11.1, The Job Monitor.

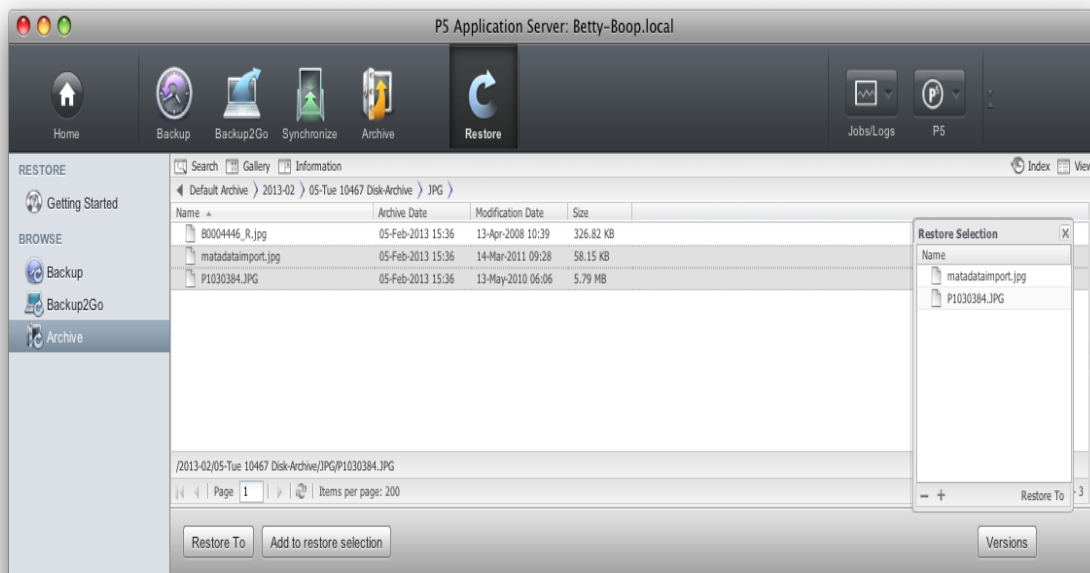


9.2.2 The Restore Selection

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 folder to restore with 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.

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.



9.2.3 The Versions Button

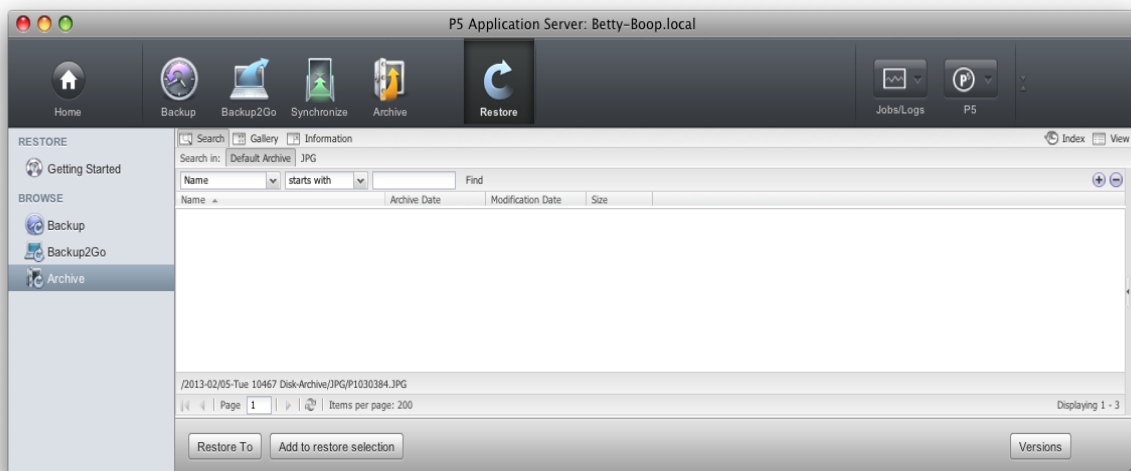
Use **Versions** to list all versions of a particular file that have been archived. To do so, select the desired file or folder from an archive index and click **Versions** in the bottom bar. All the available versions of a file will be displayed in an additional window.



9.2.4 The Search Button

To simplify the search 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.

The following items can be searched:

File name:

When searching for a file, enter the required name or expression here. The pull-down selection enables you to choose the matching mode.

- » starts with: match the beginning of the file name with the entered name.
- » contains: match any part of the file name with the given name.
- » matches: mach the entire file name with the given name .

Modification date:

The result will contain all the data that have been modified no later than the given date.

The date format is: **dd.mm.yy**

Archive date:

The result will contain all the data that was archived up to the given date. The date format is: **dd.mm.yy**.



It is possible to select multiple search expressions, for instance

- Name contains "Yesterday"
- Modification date is before Jan. 01 20013

The expressions must then all be met to find the file(s).

Click **Find** to start the search.

Files and folders can be selected for restore directly, same as in the index browser.

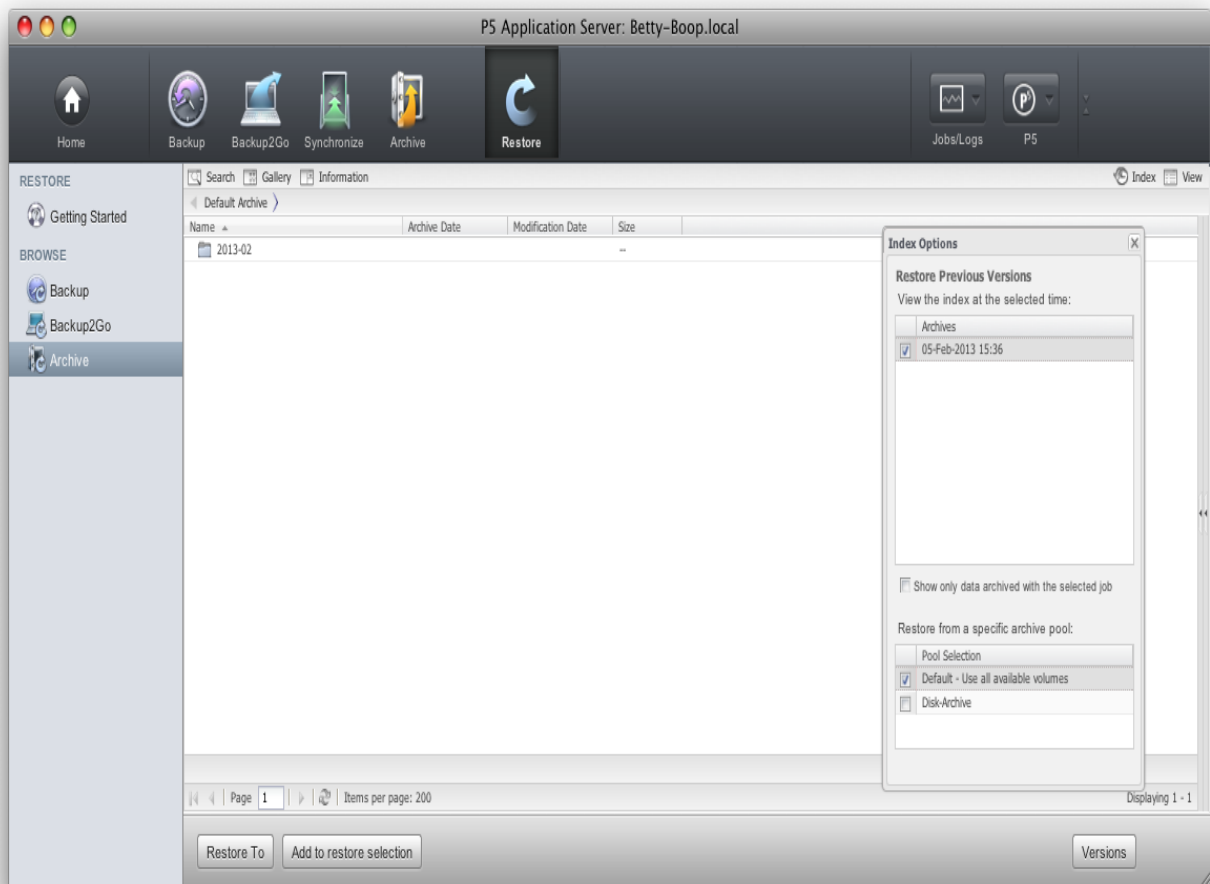
Double clicking a found folder takes you directly to that folder, allowing to restore files or folders from there or to start a new search in that sub folder.



9.2.5 Index Options

Each archive job creates a new snapshot when running. By default, P5 always shows the last 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.

Check the box **Show only data archived with the selected job** to focus on files written with the specific job that created this snapshot.

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.



10 Operating Procedures

This section describes the tasks necessary for an optimal and smooth day-to-day operation of the system. This includes:

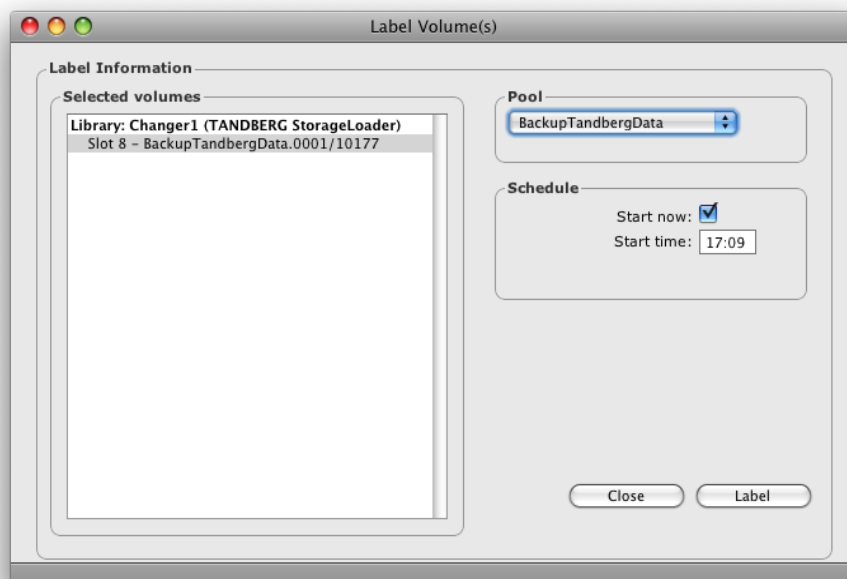
- checking the backup.
- managing the backup devices and volumes.

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 Volumes(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.

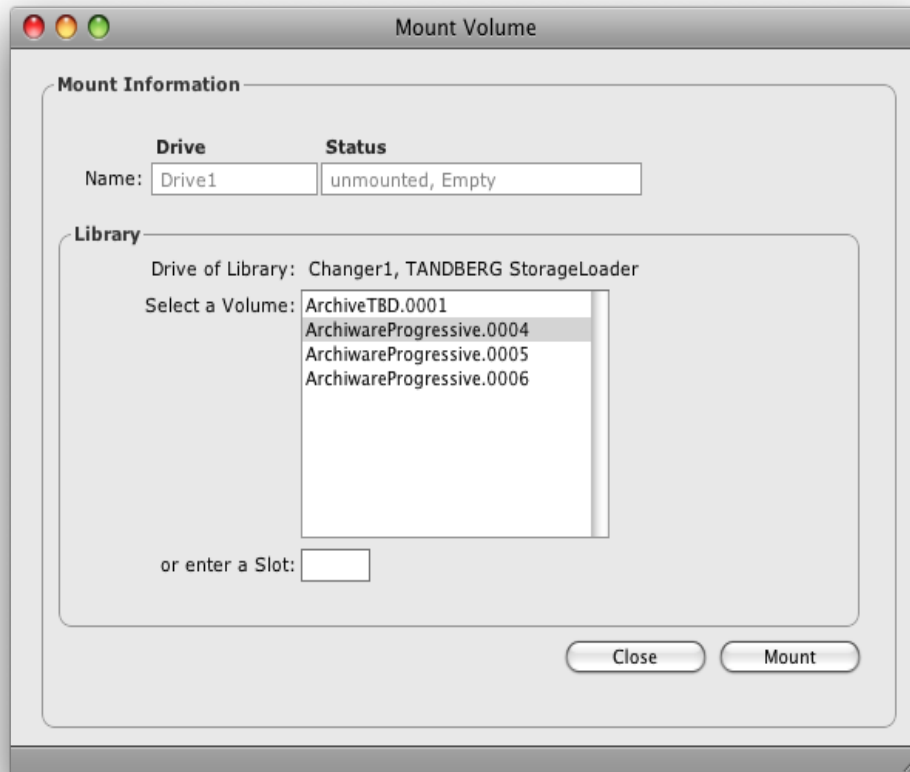
If a volume has been stored outside of the system (the library), it has to be put into the system (the library) again. Then an inventory of the library has to be run to make the volume position and availability known to P5 again. See section 6.3.4.2, Take Library Inventory....

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 Un-mounting 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 **Storage Manager** 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 the 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 to 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. In any backup or archive job, media pools are addressed. P5 Does not address single tapes when writing. To assign tapes to media pools, the tape 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 for several purposes 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 written 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 or to disk onto "virtual tapes". 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 tape 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. So it is possible to separate for instance an LTO-6 pool and an LTO-5 pool in the same changer so 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 to write 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 trivia to say that only one pool's tape can be in one drive at a time. This allows to restrict the total system load and regulate the data flow in 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.



11 Customizing

This section describes the customization possibilities in P5, so that each user can quickly and easily complete his P5 related tasks.

These include:

- customizing log-in areas.
- customizing the home log-in area.
- user groups.
- customizing reporting (mail delivery and notifications).

11.1 Customizing Index and File System Access

To create a log-in area, see section 6.15, Index Access Control (Login Areas).

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 need to perform one of a few tasks, this simplifies their task of finding the correct functionality and navigating through the application.

To configure user preferences, see section 6.16.1, Managing User Preferences.

To restrict each individual user to a particular log-in area, set as *Login Behavior* the corresponding log-in area 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.

User Groups

For security reasons, P5 recognizes only users who are authorized to access the system that P5 is installed on. This ensures consistency with security and access rights of the hosting system and enables the use of those contemporary authorization schemes that that host supports, such as LDAP, PAM, NIS and Active Directory Services. Within P5, the user groups can be associated with individual log-in areas, which allows the given group access to a defined log-in area and disallows 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 but not all log-in areas.

In order to restrict access to log-in areas to defined groups, add these groups in the log-in area record, see section 6.15, Index Access Control (Login Areas). Log-in areas with no group entered are not restricted.



11.2 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.

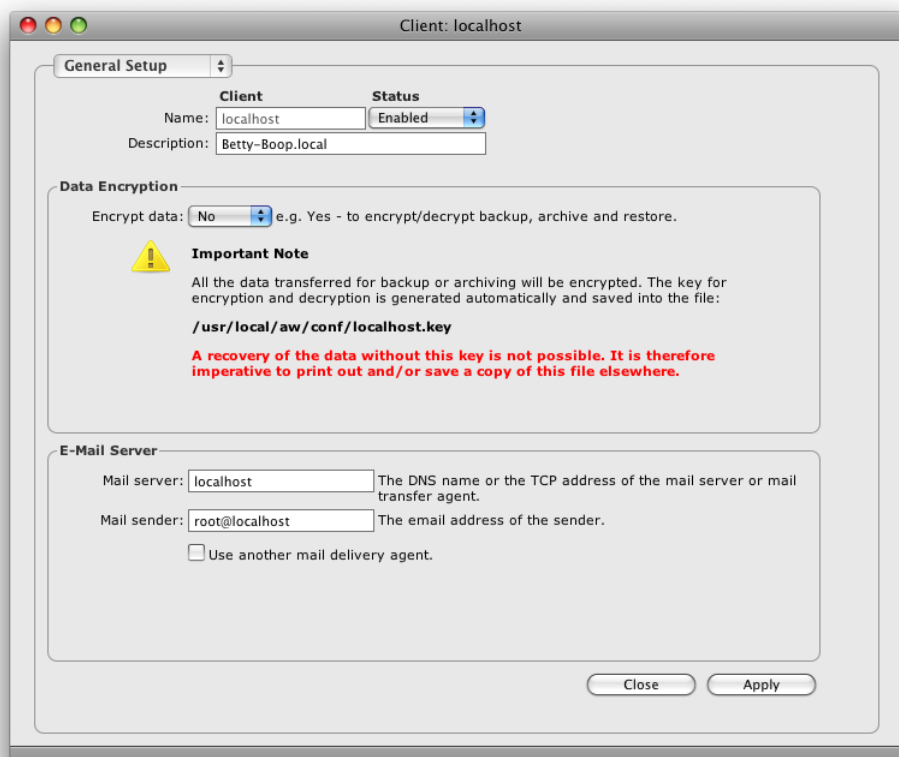
The protocols can be retrieved by selecting *Jobs/Logs* from the top menu.

See more about logs in section 6.11.2, 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 job 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 emails to be sent. The setup is described in detail in section 6.16.1, Managing User Preferences.

The prerequisite for a successful mail delivery is that a mail server (for example Postfix) is installed on the hosting system where the P5 server is located. If an alternative mail server should be used, it has to be entered in the *localhost* client configuration.





Enter as *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. See section 6.12.1, Creating a Client for how to find the localhost client configuration.

Please note that P5 expects a mail server that does not require authentication.

In addition to mail notifications sent to specific users, on screen notifications can be set for each defined P5 client. These notifications are sent as an alternative or in addition to emails. The setup of on screen notification is described in detail in section 6.12.1, Creating a Client.



- In order for on screen notification to be shown on the client:
On Mac clients, growl notifications are implemented, so the growl utility must be set up
- On Windows clients, bubble notifications are implemented. In order to enable these, the P5 Server Manager utility must be running.
- On screen notifications for other operating systems are not available.



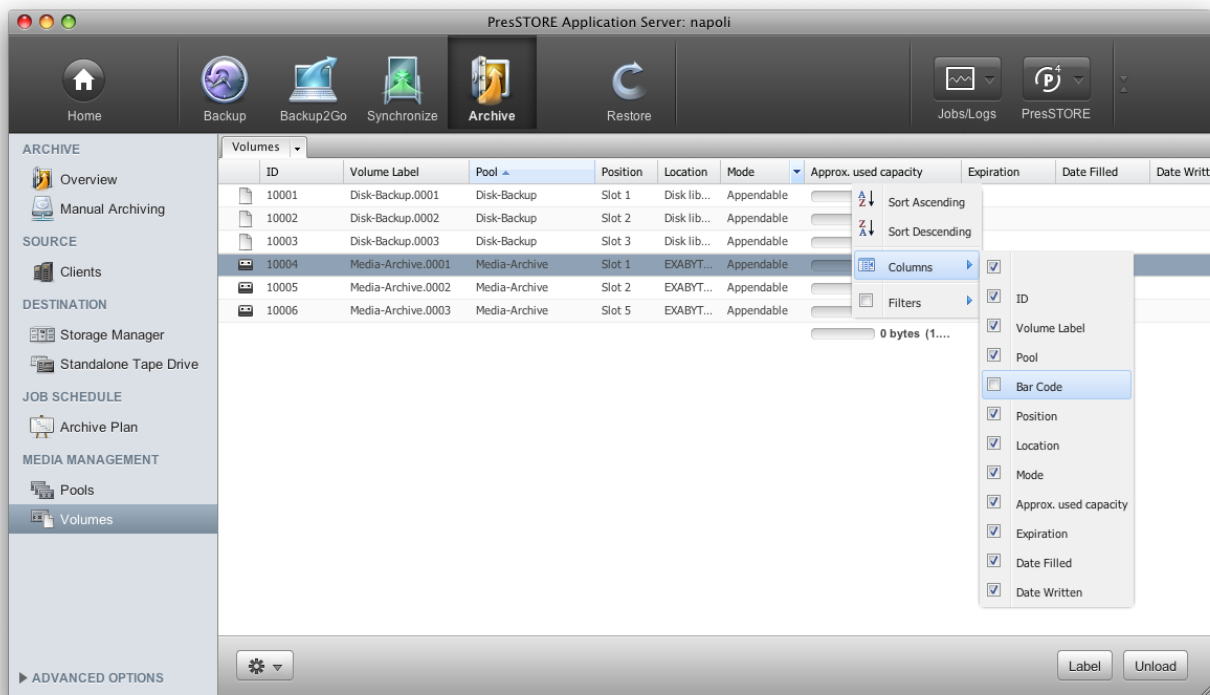
11.3 Customized Views

The table views can be customized. New views can be defined to provide specific lists with a reduced set of entries or columns. Namely in huge lists, for instance when there are hundreds of volumes filling multiple pages in the browser, new views to the lists 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.

Define a new view by clicking on the triangle right to the list tab in the headline of the window. Do not forget to name your view to make clear what it shows.

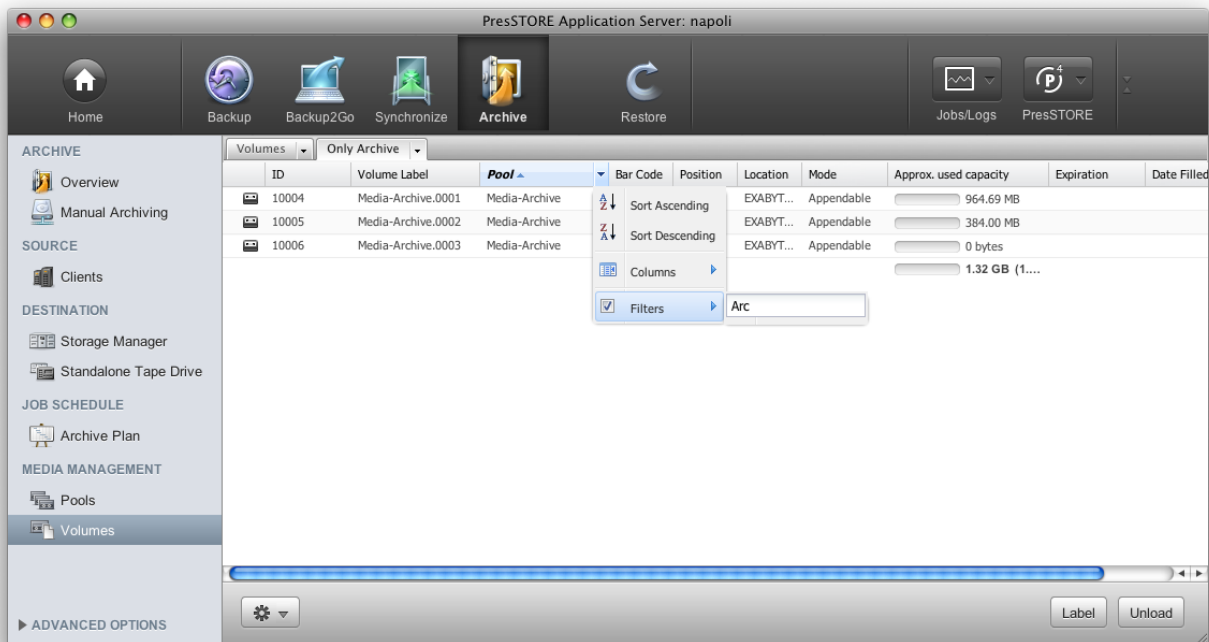
When clicking on the triangle next to a column headline, columns can be disabled:



For instance, in case the volumes are named using their bar code, both the name and the barcode columns contain the same information, so one can be omitted to enhance readability.

Filters can also be applied, for instance to define a separate view that only shows volumes belonging to a specific pool, or only volumes that are not offline.

Columns with an applied filter are marked in the column heading. We suggest to apply filters only in user defined views, otherwise an entry may get obscured because of the defined filter and never thus never noticed or displayed.



Filters are applied to the column they are defined in. The easiest filter is for character string columns, the filter makes P5 show only those entries that contain the filter expression. Date or number fields have operators that allow filtering of ranges.



12 Handling LTFS volumes with P5

The Linear Tape File System is a file system on tape specified by the LTO vendor group for LTO-5 and LTO-6 tapes. This file system, designed for data transport, is supported through a specific driver provided by the LTO vendors. Such a driver must be installed to support LTFS volumes.

Archware provides the MediaLTFS utility on its web site. Media LTFS provides an easy way to install the LTFS driver (based on the HP LTFS driver) on Intel based Mac OSX hosts.

Using LTFS together with Backup or Archive software requires however a media management to control concurrent access. P5 supports LTFS, that means it can label, mount or unmount LTFS tapes and read from or write to LTFS tapes like from or to a disk.

12.1 Creating a LTFS volume

To create a LTFS tape volume, please follow these steps:

1. Create a pool of type "LTFS" as described in section 6.4.1, Creating a Pool
2. Navigate to the Storage manager or standalone drive section and select "Label" as described in sections 6.3.4.1, Label Volumes... and 10.1, Labeling Media

The tape is then formatted to contain a new, empty LTFS file system that can be mounted via LTFS.

12.2 Importing an LTFS volume

In case the LTFS tape already exists and you want to mount the tape with P5 without labeling/reformatting the tape, it can be imported. To do so:

1. Create a pool of type "LTFS" as described in section 6.4.1, Creating a Pool
2. Navigate to the Storage manager or standalone drive section and select "Import" as described in section 6.3.4.4, Import Volume...

12.3 Mounting and Unmounting an LTFS volume

LTFS formatted volumes can be mounted just like normal tapes in PresSTORE, the procedure is described in section 10.2, Mounting Volumes. When P5 mounts a LTFS volume, the filesystem will be made available by mounting it in the PresSTORE subfolder `config/ltfs/volume-IDI`, where *volume-ID* stands for the LTFS volume ID as shown in the PresSTORE volume table (see section 6.5, Volumes. Accordingly, the file system is unmounted together with the tape.



12.4 Accessing an LTFS volume for reading and writing

P5 does not write Backup or Archive data to LTFS formatted tapes. Even if that is technically possible, it has no advantage but rather the disadvantage of reduced tape capacity. When a tape is mounted for use via the LTFS driver, the tape drive is in use and remains unavailable for other jobs until the tape is explicitly unmounted.

It is possible to Backup and Archive data from or restore data to an LTFS tape, however that requires two tape drives: one for the LTFS tape, a second one for the Backup or Archive tape.

When accessing LTFS file systems, please note some limitations one should take care off:

- Never unmount the tape while there are ongoing transactions on the mounted file system
- Avoid parallel read or write operations in LTFS, as that may cause the tape to position back and forth very often which makes the parallel access extremely slow
- Do not copy files on LTFS from one place to another
- Do not edit files on LTFS directly. Copy to disk first and copy the result back to tape.
- Remember that deleting files does not free space on the tape. If the tape is full and you want to erase it, just label it once again.
- Do not rely on the free capacity for the tape too much, the capacity may vary with writing

Please remember, the media is physically still a tape, so it has a strict sequential access. Attempting multiple parallel access will cause delays.



13 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 bureaucratic support in order to reduce down times 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:

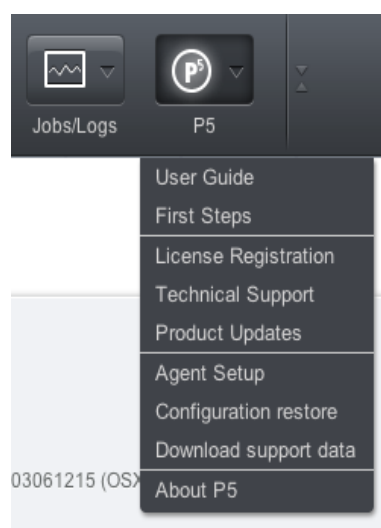
- downloading support data.
- restoring the configuration.
- disaster recovery.

13.1 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, through their many years of experience, to quickly locate and solve your problem.

To contact the support, please submit a ticket on the Archiware support page www.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 6.11, 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. Please send that log file if asked by the support staff.

13.2 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 path to the drive holding the volume from which to recover the configuration.



The following table shows some examples for the drive's OS Path:

Operating System	Path
Mac OS X	/usr/local/aw/conf/dev/awst0 (This file is generated by P5)
Solaris	/dev/rmt/0cn
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 <code>Scsi1</code> stands for the first adapter, <code>T2</code> stands for SCSI ID 2 and <code>L0</code> means LUN 3.
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.



Please note that restoring the configuration will overwrite the existing one.

13.2.1 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 8.1, 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 13.2, 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 6.14.2.2, Check.



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, an archive index, which could not be corrected by the check index function, cannot be recovered automatically, since it does not automatically get saved. 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 8.1, Restoring Data From Backup).

4. After completing this step, the new P5 server will be completely recovered and ready to start routine operations again.

13.3 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, it is therefore not possible to resolve every error that may come up.

There are two main groups of problems: either the external tape device is not found or cannot be set up. The first part handles those. The second part is about errors during operation when the device was already operable.

13.3.1 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 **Mac OS X**, click the **Apple menu** → **About this Mac** → **More Info**

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 chance to access them in P5. In addition to each tape drive (usually identified by their serial number and device vendor), the media changer (if present) must be visible as a separate device.



In case the devices are not or not all 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 **Mac OS X**, rebooting is the only option

On **Linux**, call in a terminal

```
modprobe -r sg
```

```
modprobe sg
```

to reload the generic SCSI driver

On **Solaris**, unload and load the generic SCSI driver by calling

```
rem_drv sgen
```

```
add_drv -v sgen
```

Note that you will always need to rescan after the device was powered down, disconnected or reconfigured.

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.

13.3.2 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.

13.4 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.

13.4.1 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.

13.4.2 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 an 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 excepts 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.



13.4.3 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).
- In Backup2Go, the workstations connect to the server's port 8000.



14 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 13) 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.

14.1 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 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 safety copy, but to store data that is no longer required for online access on offline media. 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 6.2.1, *Creating a Stand-Alone Tape Device*, a media changer or library, see section 6.3.1, *Creating a Jukebox*, or an emulation of a media changer on a hard disk, see section 6.3.2, *Create a Virtual Jukebox*.
2. Define a set of media to write to. In P5, this is done by defining a media pool, see section 6.4.1, *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 6.6.1) or backup plan (see section 6.7.1).

14.2 Backup Scenarios

This section illustrates some standard scenarios to create backups. In order to extend these scenarios to a backup to disk to tape setup or as an alternative setup for small business, see also section 14.3, Using Synchronize as Backup to Disk.

14.2.1 Standard backup for small business

A simple and easy to understand traditional setup for the backup 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 then can 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 realize 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 less 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.

14.2.2 Progressive Backup for huge amounts of data

P5 introduces a new paradigm useful when the time window available to perform a full backup is insufficient. Progressive Backup is rather easy to set up:

- just define an incremental level backup schedule

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 remains always 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.

14.2.3 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 reacts on the tapes that are made available and saves the data missing in that 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 6.7.1, 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 .

14.2.4 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 need 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 folder 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 6.6.1, 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 6.4.1, *Creating a Pool*. The recovery of a tape from its clone is described in section 6.5.3.2, *Regenerate Volume*

14.3 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.

14.3.1 A fast accessible backup for small businesses

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 exacter 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-sun backup would do.

Find more about how to setup synchronize in section 6.8.1.

14.3.2 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 14.3.1. 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.