Syncovery 11

Manual

This manual focuses on the Windows version of Syncovery, but most features and settings are cross-platform and identical or very similar on macOS and Linux/NAS.

Released on November 16th, 2025 by Super Flexible Software GmbH & Co. KG

Table of Contents

Introduction

Thank you for choosing **Syncovery**, the flexible and powerful solution for file synchronization and backup on Windows, macOS, Linux, and NAS platforms. Syncovery is designed to give you full control over your data: when and how it is copied, synchronized, protected, and monitored. Whether you manage a single workstation or a complex multi-server environment, Syncovery provides the reliability and precision needed to keep your files safe and up to date.

This manual introduces the core concepts, features, and workflows of Syncovery. It explains the available operating modes, scheduling options, versioning and archiving features, real-time monitoring, cloud integrations, encryption, and all advanced settings. Each chapter is structured to help both new and experienced users configure Syncovery efficiently and take advantage of its full capabilities.

Syncovery is trusted by businesses, IT professionals, and home users around the world. We hope this handbook helps you make the most of it.

Chapter 1 — Main Features of Syncovery

This chapter provides a detailed overview of the key features of **Syncovery**. It is intended to give you a solid understanding of what the software can do—from simple backups to complex synchronization and automation scenarios. With this foundation, you can decide which operating modes and settings best match your needs.

1.1 General Use Cases

Syncovery is designed to reliably synchronize or back up data between different platforms, storage systems, and devices. It supports local disks, network shares, NAS systems, cloud services, and remote servers (FTP, SFTP, WebDAV, HTTP, S3-compatible services, and more). You can define multiple "jobs" (profiles) for a wide range of purposes—workstation backups, server synchronization, cloud archiving, exact mirroring, and much more.

Syncovery is flexible enough to serve both home users and professional IT environments.

1.2 Profiles and Task Setup

- Each task is stored as a **profile**, which can be executed manually or automatically.
- Profiles can run on a schedule (using the Scheduler), at system startup, at shutdown, on user logoff, or in response to real-time filesystem events.
- Profiles can also be launched via the command line, allowing seamless integration into scripts or batch processes.
- Two editing modes are available: a **Wizard Mode** for quick and easy setup, and an **Advanced Mode** with full access to all configuration options.

1.3 Synchronization and Backup Modes

1.3.1 One-Way Synchronization

All changes from Source A are copied to Target B. Optionally, deleted files can be removed from the target. This mode is ideal for backups, migrations or "Exact Mirror" configurations.

1.3.2 Two-Way Synchronization

Changes on both sides are detected and exchanged. Suitable when users at different locations work on the same data.

1.3.3 SmartTracking (Intelligent Two-Way Sync)

SmartTracking maintains an internal database and detects file movements, renames, and conflicts.
For example, if you move a file to another folder, SmartTracking applies the same move on the other side, rather than deleting and re-uploading it.

Conflicts—such as both sides modifying a file—can be resolved automatically according to your preferences.

1.3.4 Exact Mirror Mode

Ensures that the target is an exact copy of the source, including removal of files that no longer exist in the source. Useful for vaults, deployment directories, and redundant mirrors.

1.3.5 Unattended Mode

Perfect for automation: all required decisions are defined in advance, ensuring the profile runs without prompting the user.

1.4 Advanced Performance and Filtering Features

- **Multi-threaded copying:** Multiple files can be copied in parallel to significantly speed up jobs.
- **Fast folder scanning:** Optimized directory comparison routines reduce scan times for large datasets.
- **Move detection:** File moves are detected and replicated correctly.
- **Inclusion and exclusion filters:** Fine-grained control over which files and folders are processed (by name, size, date, age, and more).
- **Time zone and DST handling:** Syncovery uses UTC internally and can tolerate specific timestamp differences (e.g., ±1 hour).
- **Resume on full disk:** Jobs can pause automatically if the target runs out of space and resume later—even in the middle of large files.
- **Obsolete Folder Handling:** Prevents the reprocessing of files moved into archive or “obsolete” folders.

1.5 Cloud and Remote Storage Support

Syncovery supports a wide range of remote storage services and protocols, including:

- FTP / FTPS
- SFTP / SSH
- WebDAV
- HTTP / HTTPS

- Amazon S3 and compatible services
- Microsoft Azure
- OneDrive / OneDrive for Business
- Google Drive
- Backblaze B2
- Box.com
- pCloud
- Many NAS platforms (Synology, QNAP, Asustor, UGREEN, etc.)

This makes it possible to synchronize or back up data not only locally or within a LAN, but also across the Internet.

A browser-based GUI for remote control is available on Windows and Linux.

Here's a screenshot of the Internet Protocols and Cloud Services you can choose:

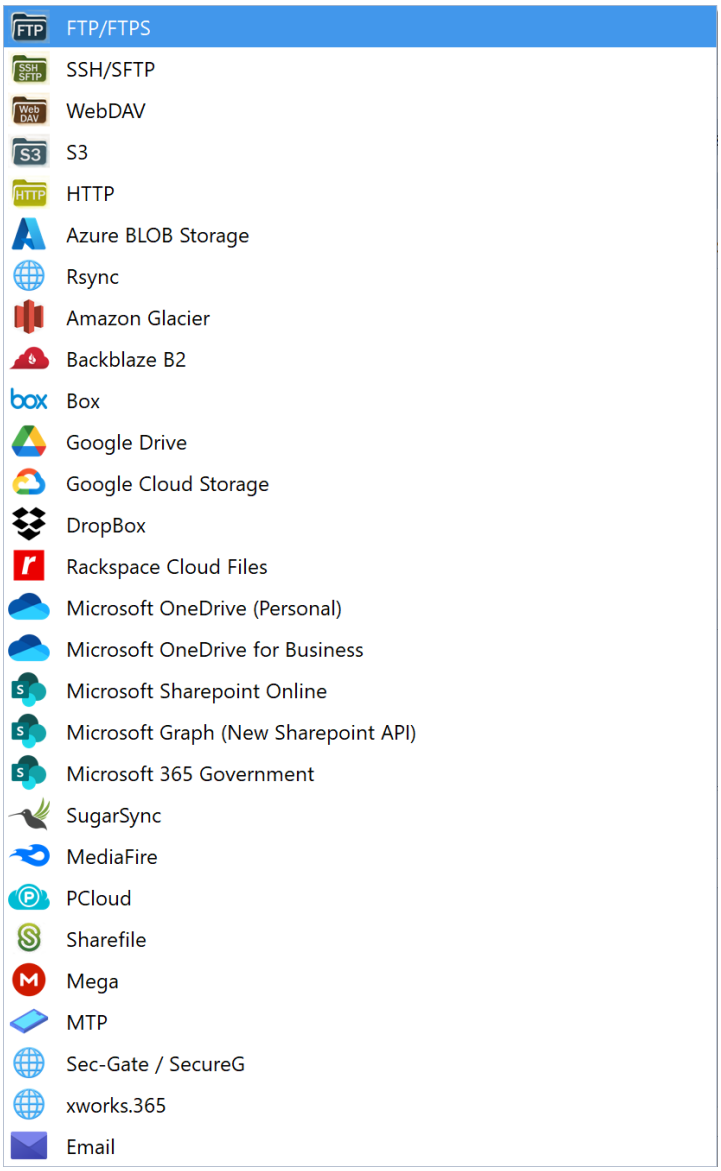

1.6 Versioning, Archiving, and Block-Level Copying

- **File Versioning:** Keep multiple past versions of files to restore older states at any time.
- **Block-Level Copy (Delta Copy):** Only the modified pieces of large files are transferred, dramatically reducing transfer time—especially for databases, PST files, VM images, or container files.
- **Compression & Encryption:**
    - Create ZIP or 7-Zip archives (depending on platform).
    - AES 128/192/256 encryption supported.
    - Optional filename encryption.
    - PGP encryption support in certain editions.

1.7 Automation, Monitoring, and Maintenance

- **Scheduler:** Automate profiles at any chosen interval or trigger.
- **Real-Time Sync:** Monitor folders for changes and react instantly (depending on edition and operating system).
- **Logs and Notifications:** Extensive logging capabilities and optional email notifications on success or failure.
- **Remote Service Tools:** Tools for remotely collecting file lists, unpacking received packages, performing MD5 checks, and monitoring other machines.

1.8 Cross-Platform Operation

Syncovery is available for:

- Windows (including service mode)
- macOS
- Linux (desktop, server, NAS)
- Various NAS firmware platforms

The Scheduler is able to run as a system service so profiles continue running even when no user is logged in.

Advanced NTFS handling, cloud APIs, and a rich GUI make Syncovery suitable across diverse environments.

1.9 Edition-Specific Features

Depending on your license (Standard, Professional, Premium), additional features may be available:

- Real-time monitoring
- SmartTracking
- Remote Service functionality
- Block-level copying
- Web GUI
- PGP encryption
- Priority support
- Specialized cloud connectors
- Detailed performance and filtering options

It may be worthwhile to check the feature matrix if you plan to use advanced features across multiple systems.

1.10 Why Syncovery? — Key Advantages

- **Extreme flexibility:** From simple backup profiles to highly customized configurations.
- **Cross-platform availability:** Works on all major operating systems and storage types.
- **Scalable:** Suitable for home users, small offices, enterprises, and cloud environments.
- **Fully automatable:** Scheduler and service mode allow fully unattended operation.
- **Robust and reliable:** Features like move detection, block-level delta processing, and resume-on-full-disk ensure dependable performance.
- **Professional-grade options:** SmartTracking, multi-threading, advanced filtering, encryption, and remote monitoring.

Chapter 2 — Installation, Updates, and Configuration

This chapter explains how to install Syncovery, how to keep it up to date, and how configuration and licensing work. Syncovery is designed to be simple to install, safe to update, and easy to maintain over long periods—even across major version upgrades. All user settings are preserved unless explicitly changed.

2.1 System Requirements and Preparations

Syncovery is available for Windows, macOS, Linux, and various NAS platforms.

Windows
All versions of Windows since Windows Vista and Windows Server 2008 are fully supported, including Windows 10/11 and Server 2019-2025 and all other versions and editions of Windows that have a graphical user interface. Syncovery is available in 32-bit and 64-bit editions.

Syncovery 9 generally also still runs on Windows XP and Server 2003, but some features may not work (especially cloud features on Server 2003). The 64-bit editions require at least Windows Vista. For older 64-bit Windows versions, please use the 32-bit edition of Syncovery.

For Windows 2000, please download Syncovery version 6 from our additional downloads page.

Macintosh
Syncovery runs on all Intel and Apple Silicon Macs, and on all macOS versions since 10.4 “Tiger” up to including the latest version. Syncovery 8 is available as a 64-bit edition which runs on the newest macOS versions (the 64-bit edition is required starting with macOS 10.15 “Catalina”). A separate version is also available for old PowerPC Macs. This version also requires at least macOS 10.4 “Tiger”.

Linux
The command line and HTML/Web GUI Edition “SyncoveryCL” is available for Linux. It should run on any recent 32-bit or 64-bit Linux distribution for Intel, ARM, AARCH64, PowerPC, PPC64 and PPC64EL processors. See the Linux download page for details. Dedicated installation packages are available for Synology and QNAP NAS devices, as well as Debian and RPM installers.

FreeBSD
The FreeBSD version is very similar to the Linux version, but only available for Intel/AMD 64-bit processors. Separate installer packages are available for FreeBSD 11 and 12, as well as a generic .tar.gz file.

This manual focuses on the Windows version.

No special preparations are required prior to installation. You do not need to uninstall a previous version or stop running profiles; Syncovery’s installer handles updates safely and automatically.

2.2 Installing Syncovery

2.2.1 Downloading the Installer

The most recent Syncovery version is always available from:

www.syncovery.com

Download the Windows setup program and save it to your computer.

2.2.2 Running the Setup Program

To install Syncovery:

1. Double-click the downloaded installer.
2. Follow the on-screen instructions.
3. Choose the installation directory if you want to modify the default location.

By default, Syncovery is installed into:

`C:\Program Files\Syncovery`

All executable components, libraries, and related files are contained within this single directory.

Syncovery does not install DLLs or other components into Windows system folders.

This ensures a clean, portable installation and prevents interference with other applications.

2.2.3 Installation Notes

- Installing Syncovery does not require administrative system changes beyond writing to the application directory.
- The installer can safely overwrite a previous installation of Syncovery.
- Installation does **not** alter or delete configuration files.
- After installation, Syncovery can immediately be launched from the desktop or the Start Menu.

2.2.4 Silent Install and Uninstall

To perform a silent installation, use the command line:
Syncovery64Setup.exe /VERYSILENT

Syncovery uses the InnoSetup tool to create the installer. Additional command line parameters are documented on the InnoSetup web site:
https://jrsoftware.org/ishelp/topic_setupcmdline.htm

To perform a silent uninstall:

- delete Syncovery.exe first
  (this suppresses an interactive question about keeping settings)

- then run unins000.exe /VERYSILENT from Syncovery's program folder.

2.3 Updating Syncovery

Updating Syncovery is intentionally simple. You can update at any time using one of two methods:

2.3.1 Update by Reinstalling

Just download the newest setup program from the Syncovery website and run it.
The installer automatically replaces older program files **in place**.

- There is **no need to uninstall** your existing version.
- You can upgrade even from very old versions directly to the newest one.
- Your configuration is preserved.
-

2.3.2 Update from Within the Application

Syncovery includes a built-in update checker:

Help → Check for Update

If a newer version is available, you can download and install it directly.

2.3.3 Compatibility and Downgrading

- Syncovery is **always compatible with older configuration files**.
- All profiles, settings, schedules, and logs remain untouched during an upgrade.
- **Downgrading** (installing an older version over a newer one) is supported.
    - However, any settings that exist only in the newer version may be ignored or lost when running an older version.

2.4 Configuration Files and Storage Locations

Syncovery keeps configuration files separate from program files to support clean updating and service operation.

2.4.1 Default Configuration Directory

By default, all configuration data—including profiles, schedules, settings, licensing, and logs—is stored in:

`C:\ProgramData\Syncovery`

This location is standard for applications that run in service mode or in shared environments.

2.4.2 Changing the Configuration Location

Users may choose a different configuration directory from within Syncovery's settings. (Changing the configuration path cannot be done during installation; it must be adjusted later inside the application.)

Syncovery will use the selected directory for:

- Profiles (INI or SQL formats, depending on version)
- Global settings
- Databases
- Log files

The paths for Databases and Log Files can be changed on the Program Settings dialog in Syncovery.

Please note that the configuration directory should be on a local hard drive or SDD. Please change the default path only when absolutely necessary, to avoid issues.

2.4.3 Portability Considerations

Because Syncovery keeps code and configuration separate:

- Reinstalling or updating the program does not overwrite profiles.
- Backing up configuration data is as simple as backing up the ProgramData\Syncovery directory.

2.5 Licensing and Activation

2.5.1 Entering Your License

Syncovery uses a straightforward licensing model without online activation or licensing servers.
To activate your license:

1. Open Syncovery.
2. Choose **Help** → **Registration**.
3. Enter your **Registration Name** and **Registration Code** exactly as received.

The program activates immediately without needing an Internet connection.

2.5.2 License Portability

Syncovery's license is based on trust rather than hardware locking.
You may enter the same license code on a new computer **as long as it is no longer used on the old one**, unless your license agreement explicitly allows multiple installations.

Parallel or multi-device usage is permitted in some cases, depending on the purchased license type.

For details, consult the license terms available on the Syncovery website.

2.5.3 Upgrading and Licensing

- Updating Syncovery does not require re-activation.
- License information is stored in the configuration directory and preserved across reinstalls.
- Major version upgrades typically accept previous version licenses unless otherwise stated in the upgrade policy.

2.6 Installed Program Files and Services

Syncovery installs a few executable files along with DLLs, some accessory files, and the WebDocs folder for the optional browser-based Web GUI. Here's a list of the executable files and their usage. Your choices in the Setup program determine if all of these files are installed:

Syncovery.exe
This is the main **GUI program** used to set up your jobs and run them in Attended or Unattended Mode. You can close it when your jobs are set up and the scheduler is running, and you only need to open it when you need it directly. The scheduler and background jobs run in separate processes.

SyncoveryCL.exe
This is the Syncovery **command-line tool**, which also doubles as the **Syncovery Service** when the scheduler runs as a service. See the separate chapters concerning the service and the command line.

SyncoveryService.exe
This executable is the **Background Scheduler** (when the scheduler doesn't run as a service). It is also used to display the **Tray Icon**, and may be running just for that purpose even if the scheduler is running as a service (SyncoveryCL.exe). The process name SyncoveryService.exe is a bit misleading for historic reasons: it used to be the service, but that's a legacy use. The reason the service had to be switch to SyncoveryCL.exe is that SyncoveryService.exe still contains some GUI components, which modern services are not allowed to.

You may see **multiple processes** called SyncoveryCL.exe and SyncoveryService.exe in Task Manager. The reason is that Syncovery launches additional processes to execute profiles, so that they are independent from the scheduler and can't influence the scheduler reliability in any way.

SyncoveryGuardian.exe
The Guardian is an optional service that monitors the main scheduler service and ensures it's running, possibly restarting it if the main service fails (which should never occur). It is rarely needed.

SyncoveryRemoteService.exe
The Remote Service can run on machines that Syncovery wants to synchronize with remotely. It provides various functionalities to assist the main Syncovery program:

- Generate folder and files listings fast
- Unpack incoming compressed packages
- Generate hashes for file verification and block level copying
- Provide SFTP connectivity

SyncoveryFileSystemMonitor.exe
The File System Monitoring Service can be used to enable super-fast Block Level Copying my constantly recording a list of changed blocks on your hard drive. See the related chapter for details.

SyncoveryAuxServicesCPL.exe
The **Auxiliary Services Control Panel** is used to set up, install and configure the Syncovery Remote Service and the File System Monitoring Service.

2.7 Summary

Installing and updating Syncovery is intentionally safe, simple, and non-intrusive:

- Program files are self-contained in C:\Program Files\Syncovery.
- Configuration files are kept in C:\ProgramData\Syncovery.
- Updates are performed in place—no uninstallation needed.
- All profiles and settings are always preserved.
- Licensing is offline, portable, and based on trust.

Chapter 3 – Creating a Profile in Wizard Mode

For your first profile, it is a good idea to use Wizard Mode, which will guide you thru the steps of creating a basic job definition (or profile).

This chapter shows the steps used in Wizard Mode.

From the Welcome Screen, choose “Perform a synchronization or backup now”.

3.1 Step 1 of 6: Specify Folders

A synchronization, mirror or backup always involves two base folders - the left-hand base folder and the right-hand one. The two folders are specified in this step. You can click the "Browse" buttons to pick your folders, or click the "Cloud" buttons to access cloud servers or use Internet Protocols . If you need to create a new folder, you can do that using the Browse buttons, or you can just type the full path including the new folder name.

Folder paths can be local drives, network paths, or cloud / Internet Protocol URLs. Depending on the type of storage location, Syncovery uses standard Windows paths, URLs, or its own internal syntax for cloud connections. Here are some examples:

Local drive:	C:\Users\Steve\Documents
Network drive:	\\SERVERNAME\SHARENAME\FolderName
FTP URL:	ftp://ftp.hidrive.strato.com/users/steve/Documents
Google Drive:	ext://Google Drive/DocumentsMirror
Sharepoint:	ext://superflexible.sharepoint.com:Documents/OpenCases

Cloud Services like Google Drive, DropBox, and Sharepoint require authorization using the OAuth2 method. The authorization process is triggered when you click the Browse button or run the profile. A browser window will open where you can log in and authorize Syncovery.

For network drives, the **SMB** button allows you to specify credentials to connect to the network share. This may or may not be necessary – in many cases, Syncovery can just use your existing connection made by Windows File Explorer. If you are logged on to a Windows domain, the account should automatically give you access to network shares.

3.2 Step 2 of 6: Choose Sync Direction and Subfolders

For a backup or replication, you usually copy Left to Right only (if your original data folder is the left-hand folder). If both sides contain updated files, or files missing on the other side, then you may want to copy in both directions.

Please note: if a file already exists on both sides, the software looks at the "Last Modified" timestamps in order to determine whether it needs to be copied (and in which direction). The file with the older timestamp will usually be replaced by the one with the newer timestamp.

In addition to the copying direction, you can choose which subfolders of your base paths to process. In many cases, it's best to choose **All** subfolders. Exclusions can be configured separately. Howevery, you can also use the **subfolder selection** dialog to choose subfolders.

When making a subfolder selection, it is important to be aware of Syncovery's two selection modes, which determine if new folders in the future are automatically added to the selection. This is chosen by the checkmark:

Automatically Add Future New Folders and Files to Selection

Whether this checkmark is chosen or not, new items in fully selected subfolders will always be included. But future new siblings of currently deselected subfolders are only added to the selection if the checkmark is chosen.

The technical background of this "Automatically add" mode is that the selection is actually stored negatively: Syncovery stores the deselected subfolders instead of the selected ones. That way, it is ensured that only the currently deselected subfolders are excluded from the job when it runs, and future new folders are included.

3.3 Step 3 of 6: Choose Sync Modes

You can choose between various synchronisation modes, which are explained below. All modes will perform an efficient synchronization, and detect if files have been moved into a different subfolder on one side, and aim to perform the same move on the other side.

- **Standard Copying** will copy new, missing, or modified files, but it will not delete any files.

- **Exact Mirror Mode** can be used to make an exact copy of your data. This means that files will be copied in one direction only.

  Warning: files that don't exist on the source side will be deleted from the destination. Files that are newer on the destination may be replaced with older files from the source, because this mode will make sure that the destination is an exact copy of the source. Use with caution.

  You can click the **Configure**... button to fine-tune the Exact Mirror behavior.

- **SmartTracking** is mostly used for two-way synchronization. It can detect deleted and conflicting files and must be configured to your requirements via the **Configure**... button.

- **Move Files to Destination** is a fourth mode, not available in Wizard Mode, but can be selected in Advanced Mode. It will move files from source to destination, removing them from the source side.

3.4 Step 4 of 6: Choose File Types

Syncovery uses Inclusion and Exclusion Masks to determine which files should be copied.

To synchronize all file types, leave these fields unchanged. In order to work only with specific file types, you need to know their file name extensions. For example, a text document is often identified by extensions such as docx, rtf or txt. The file masks to specify would be: *.doc;*.docx;*.rtf; *.txt.

In addition, you can exclude certain file types, for example: *.bak.

You can also use more complex masks, such as
Applications*.xlsx

The Exclusion Masks are also applied to folder names. For example, to exclude a folder named "Archive", you can type:

Archive	excludes folders named Archive anywhere in the folder tree
\Archive	excludes the Archive folder only at the top level of the folder tree
\Steve\Archive	excludes the Archive folder only in the subfolder "Steve"

3.5 Step 5 of 6: Specify Scheduler Settings

Synchronizations can be started manually at any time, or they can be run automatically by the scheduler. On this page, you can specify if the scheduler should run this profile regularly or not.

NOTE: in addition to turning on the scheduling for the profile on this page, you need to start the scheduler from the main window. When the scheduler is off, no profiles can start automatically.

The scheduler can run as a background app or as a service. This is explained separately.

More scheduling options are available in Advanced Mode.

3.6 Step 6 of 6: Save and Run Profile

Congratulations! You have made all the required choices for a synchronization.

In order to be able to use these settings more than once, they are now saved as a profile. Please give it a name.

When you run the synchronization, the software will compare the left and right folders and show the results in a Synchronization Preview window. You should then look at the proposed copying, moving, and/or deleting operations and make sure that they match your intentions. If everything is OK, you will then be able to start the actual copying process.

3.7 The Synchronization Preview

When a job is run in **Attended Mode**, the Sync Preview will show you the list of actions that should be performed. This window shows all files found on both sides that have a proposed action. You can see on which side(s) a particular file currently exists by looking into the **Date** and **Size** columns.

On the left, you have various options to control which kinds of files will be shown. When you hit the **Start** button, only the files shown will be affected, and only if they have an **Action** specified in the Action column (see below the screenshot). To change actions for specific files, select the files and use the Menu button or right-click the file list invoked by the right mouse button.

Possible Actions

➔ : copy to the right side
➔ : copy to the left side

MOVE
File will be moved into a different folder on one side. Use the radio buttons on the lower left to control on which side moved files will be adjusted.

DELETE
File will be deleted. This can only occur in Exact Mirror Mode, or if deletion has been manually selected for a file.

CONFLICT
The items and the left or right side are different, but the software cannot detect the correct action. This can be because on one side, it is a file, and a folder on the other side, or it can be caused by a failed binary comparison of a file. SmartTracking also causes this label when a files has been changed on both sides.

time-> / <-time
The timestamp of the folder should be adjusted.

case-> / <-case
The spelling of the file should be adjusted to match the case of the source side.

3.7.1 Sync Preview Menu

Don't forget to try the **Menu** button or right-click on a folder or file selection in the Sync Preview. There is an extensive context menu with powerful menu items to modify the actions, print, view the preview in HTML format and much more.

Using this context menu, you can manually select the action that should be performed with the selected files and/or folders. In addition, you can remove the items from the list, as well as delete files immediately. The menu also informs you about the keyboard shortcuts that can be used in the File List View.

Note that these menu items can be very useful:

Select All Files In This Folder (F3):
This command extends the selection so that it includes all files in the folder that belongs to the top of the selection. Combining F3 with some of the other commands, you can quickly change the action for or deselect or remove complete folders in the list.

Select All Items In This Subtree (F4):
This command extends the selection so that it includes all files and folders in the folder that belongs to the top of the selection, as well as all sub-folders. Combining F4 with some of the other commands, you can quickly change the action for or deselect or remove complete subtrees in the list.

Chapter 4: Advanced Mode

Advanced Mode consists mainly of the **Profile Overview** (main dialog) and the **Profile Editor** window.

4.1 Profile Overview

This tab sheet shows an overview of all stored profiles and allows you to manage your profiles.

Toolbar buttons allow you to:

- Add New Profile
- Edit Profile
- Rename Profile
- Delete Profile

- Run Selected Profile(s)
- Run Checked Profile(s)
- Pause Profiles
- Stop Profile
- Run the Restore Wizard

- Access additional settings via the Tools and Settings button (Gears icon)
- Access the Program Settings dialog
- Open the Account Manager dialog
- View Log Files
- Open the Profile Group Editor dialog

When editing a profile from the Profile Overview, you will automatically edit in in Advanced Mode and see all available options.

There are two ways to select profiles:

- Click on the rows in the table to select them with a light blue bar. You can select multiple profiles by dragging or ctrl- and shift-clicking. Selected profiles can be run with the single Play button.

- Click on the checkmarks to choose a combination of profiles which is remembered even after you exit the program. These profiles can be run with the double Play button (or fast forward, symbolizing "Start Many").

Also make sure you try the **context menu** which you get by right-clicking on profiles. You have four choices of how to start a profile:

- Run In Attended Mode
  In attended mode, the profile will run in the foreground and you will see all the progress information related to building the file list. You will then see the Synchronization Preview and you can check the copying actions that the program wants to make. You can make changes in the Preview and then continue with the execution or cancel the profile.

- Run In Unattended Mode
  In unattended mode, the profile will also run in the foreground but it will not wait for any user input or offer you to examine the Sync Preview. It will start copying right away, but you can cancel it at any time.

- Run In Background
  When run in the background, the profile will be executed by a separate process which will show in Task Manager as SyncoveryService.exe. You will not see a lot of details, but some progress information is shown in the Profile Overview. You can stop the profile with the Stop button in the toolbar.

- Run in Background With Preview
  Similar to “Run in Background” but you will have the ability to see the Sync Preview before actually starting the copying phase. Ths is quite useful if you need to run multiple jobs simultaneously but want to check the previews.

Running profiles in the background is the only way to run several profile simultaneously. The scheduler can also run profiles in the background and simultaneously, see the option "Start Profiles In Parallel" in the Tools and Settings menu on the Scheduler tab sheet.

4.2 The Scheduler

You have the choice between three schedulers: the Background Scheduler (default), the Foreground Scheduler, and the Service Scheduler (Professional Edition only).

Background Scheduler
This is the default scheduler. It will run in the background in a separate process. All status and progress information can be seen on the Scheduler tab sheet. You can actually make the background scheduler visible by right-clicking on its tray icon and selecting "Reveal Background Windows".

The Background Scheduler usually starts when you log in to Windows. You can select this in the Tools and Settings menu. You can also choose whether to start profiles in parallel (simultaneously).

Foreground Scheduler
While the Foreground Scheduler is running, you cannot access any other parts of the software. For that reason, this scheduler type is only used for special purposes such as debugging.

Service Scheduler
The service scheduler is a bit more difficult to set up. Please note that the word "install" is used in a double sense when speaking about services. First, the service's file must be installed (=copied) onto the computer. Second, the service must be installed (=registered) with the Windows operating system so that it is known as a service and automatically started when Windows boots up.

Install Service / Uninstall Service buttons
These buttons register the Syncovery Service with Windows. A registered/installed service can be started and stopped, and by default it will start automatically when Windows boots up.

Local Service / Remote Computer
You have the choice between administration of the service on your local computer, or on a remote machine. To connect to a remote system, enter that system's network name and click on "Connect".

Start / Stop
When the service has been installed properly, it can be started and stopped. **Start** launches the Syncovery Service process and activates its scheduler. **Stop** cancels whatever the service is currently doing and terminates its process. Please note that several seconds may pass after clicking on either of these buttons before the operation has completely finished.

4.2.1 Running the Scheduler as a Service

This information applies mostly to Microsoft Windows, although the scheduler can run as a service on all platforms.

Running the scheduler as a service is a little more difficult to set up, but it has some advantages. The service will start up automatically with Windows, even while no user has logged on yet. The service runs invisibly in the background without users noticing it. A user can also log out while the service is running a job. Neither will disturb the other. The name of the service is the Syncovery Service.

***Note**: Syncovery comes with several services, most of which are optional. This chapter deals with the **Syncovery Service**, which implements the scheduler functionality. In the most recent Syncovery versions, the executable name is SyncoveryCL.exe (older versions used SyncoveryService.exe).*

Information dialog shown when the Info icon is clicked

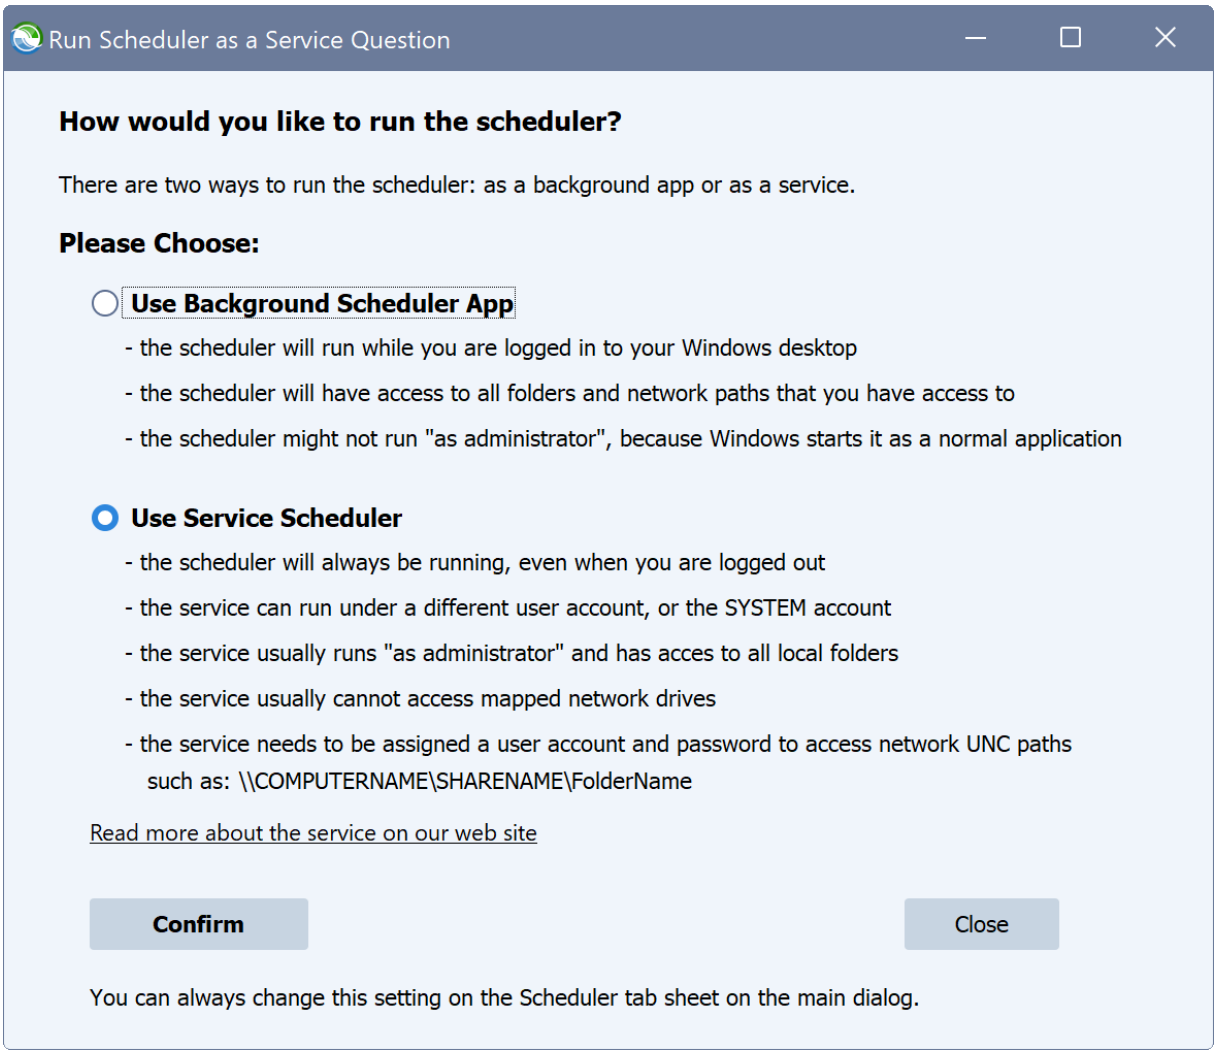

Service Control Panel when the scheduler runs as a service

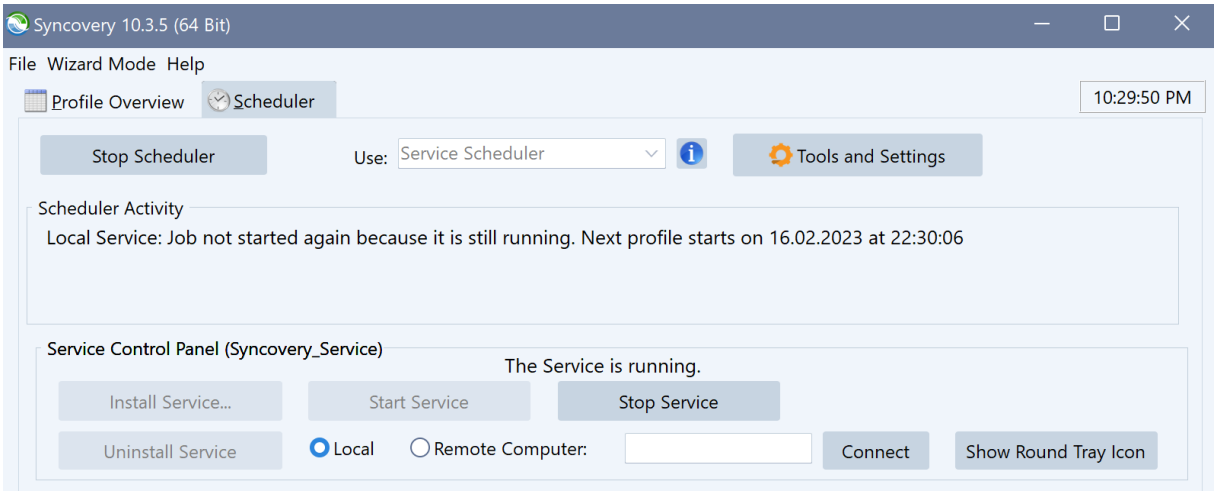

Tips

- The scheduler can run as a Windows Service. This means that scheduled synchronizations take place without users having to log on.

- The service is installed and started from the Scheduler tab sheet of the main application window. If the "Background Scheduler" is running, please stop it and then choose **Use: Service Scheduler**. Then you will see the service-specific install/uninstall, start/stop buttons.

- In order to access network drives, the service must be given a log on account. This is done when clicking on the **Install...** button on the Scheduler tab sheet to install the service. To change the log on account, please uninstall and then re-install the service. Make sure that you specify a user account where you know that Windows Explorer has access to the volumes that you need. Choose an account that has network access without Windows Explorer asking for a password for the network drive. Windows Explorer should have already stored the password.

- Also in order to enable access to a network drive, please try using a UNC path such as \\servername\sharename\foldername rather than a mapped drive letter.

- If this is not sufficient, you can provide a username and password for the network resource in each profile. Use this setting on the Job tab sheet in the profile: **Network Connections...** However, in many cases this is not needed. Rather than specifying the full path for the network connection, you can also try specifying just \\servername

- Deleting to the recycle bin is not supported by the service.

- Even though the service is normally invisible, you can get a pretty good picture of what it's currently doing or planning to do, using the Scheduler tab sheet of the main

application window. You can even connect to a remote computer and control and watch the service running there.

4.3 Authorizing Syncovery With Cloud Services

Syncovery supports a wide range of cloud storage providers, many of which require a secure authorization process based on the **OAuth2** standard. OAuth2 ensures that Syncovery can access your cloud data **without ever storing your password**, using industry-standard authorization tokens instead.

This chapter explains how OAuth2 works within Syncovery, how the authorization process is triggered, and what to expect when configuring cloud profiles.

---

1. Overview of OAuth2 in Syncovery

Services such as **Google Drive**, **Dropbox**, **Microsoft OneDrive / SharePoint**, **Box**, and various enterprise cloud systems use OAuth2 for secure authentication. When Syncovery needs to connect to one of these services, it must first obtain permission from the user through a web-based login and authorization dialog.

Syncovery never handles your plain-text login credentials. Instead:

1. A browser window is opened.
2. You log in on the provider's official website.
3. You grant Syncovery permission to access your files.
4. The provider returns an authorization token to Syncovery.
5. Syncovery saves this token in your configuration so future access is automatic.

Tokens can be refreshed silently when they expire, so repeated logins are usually unnecessary.

---

2. When Authorization Is Required

Syncovery triggers the OAuth2 authorization flow automatically whenever it needs valid credentials. This happens in two situations:

• Clicking the Browse Button

On the cloud settings dialog, most cloud services include a "Browse" button so you can select a folder on the cloud storage. Clicking this button will:

- Open a browser window for OAuth2 login if Syncovery has not yet been authorized, or if the previous authorization has expired.
- After successful login, the folder selection dialog will be displayed.

• Running a Profile for the First Time

If a profile is started and Syncovery does not yet have authorization for the chosen account, it will automatically open the OAuth2 login window before beginning the synchronization.

You do not have to configure OAuth2 manually; it is always initiated on demand.

---

3. Container vs. Folder Authorization (Services with Multiple Entry Points)

Some cloud services provide multiple "containers" or storage roots. Typical examples include:

- **Google Drive:** Personal Drive and Shared Drives (formerly Team Drives)
- **Microsoft SharePoint:** Sites, document libraries, subsites, teams
- **OneDrive for Business:** User drive vs. organization-wide storage
- **Dropbox Business:** Personal workspace vs. team folders

To support these structures, Syncovery offers **two separate Browse buttons** in the cloud settings dialog for such services:

A. Browse (Container)

This button lets you choose the top-level container:

- Google Personal Drive
- A specific Shared Drive
- A SharePoint Site (e.g., "Marketing Site")
- A Document Library within a site

The selected container determines the root of all subsequent browsing.

B. Browse (Folder)

Once the container is selected and authorized, this second browse button allows you to choose the exact folder inside that container.

This separation is necessary because cloud services often treat containers as independent storage roots with their own permissions and metadata. Selecting the correct container ensures that Syncovery communicates with the right part of the cloud storage.

---

4. Managing OAuth2 Tokens

Syncovery manages OAuth2 tokens automatically, so normally no user interaction is required after the initial authorization.

Key points:

- Tokens are securely stored in Syncovery's configuration directory in a separate Vault file (typically in **C:\ProgramData\Syncovery** on Windows).
- Refresh tokens allow Syncovery to renew access silently without user intervention.
- If a token becomes invalid (e.g., after password reset, admin policy changes, revoked access), Syncovery will prompt for re-authorization the next time access is needed.

You can always **Revoke Access** from within your cloud service, or ask Syncovery to **forget the cloud tokens**. This is done in the Syncovery **Account Manager**.

---

5. Browser Requirements

Syncovery uses your system's default browser to display the OAuth2 login page. Any modern browser is suitable.

If your environment restricts opening external browser windows (for example, in a locked-down corporate scenario), you may need to:

- Allow Syncovery to launch a browser, or
- Temporarily disable restrictions during the setup phase.

Once the token has been obtained, Syncovery no longer requires browser access for normal operation.

---

6. Troubleshooting OAuth2 Authorization

If you encounter issues:

- **Browser does not open:**
  Check system policies, antivirus software, or remote session restrictions.
- **Login succeeds, but Syncovery does not receive the authorization:**
  Your browser may be blocking the redirect or popup. Allow popups and ensure JavaScript is enabled.
- **Repeated authorization requests:**
  The administrator may have disabled long-term tokens, or security policies may require frequent re-authorization.
- **Authorization with multiple accounts:**
  Make sure you are logged in with the intended cloud account before granting access.

Summary

Syncovery integrates OAuth2 seamlessly to provide secure, password-free access to major cloud services. Authorization is triggered automatically when needed, typically through a browser-based login. Services with multiple storage containers provide two browse buttons—one to choose the container and one for the folder—ensuring accurate selection of the cloud storage location.

OAuth2 tokens are handled fully by Syncovery and normally require no maintenance. Once authorized, cloud profiles run smoothly, securely, and with minimal user intervention.

4.4 The Account Manager

Syncovery 11 features an Account Manager that allows you to manage all your FTP, SFTP, WebDAV, S3 and cloud accounts in one place, to be used across various profiles.

It is a simple yet powerful dialog that also shows useful information on each account:

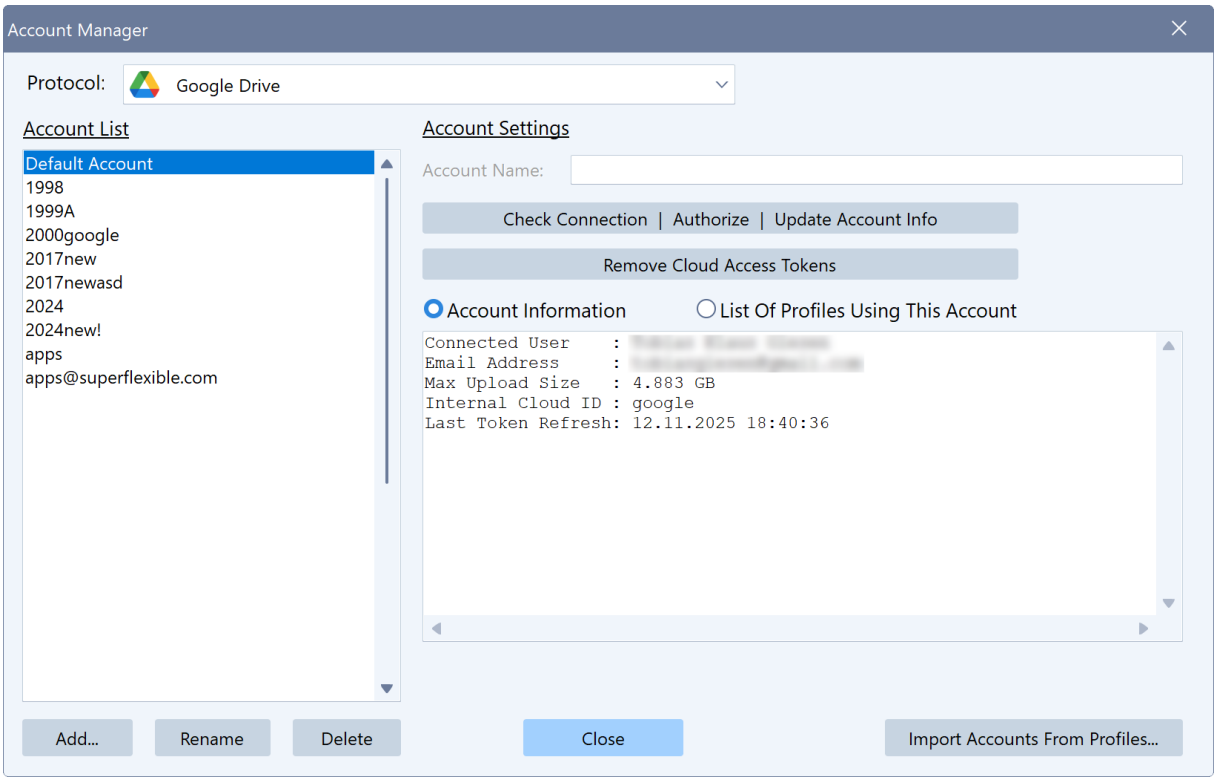

Chapter 5: Profile Editor in Advanced Mode

When you add or edit a profile in Advanced Mode, a separate profile editor window will open, where you will find the following settings.

5.1 Main Profile Settings

The Main Settings define the basic behavior of your sync job. Please choose the copying directions carefully. The Sync Operation Mode will determine Syncovery's sync logic - for example, whether to copy files only or also synchronize file deletions.

The Main Settings ask you to:

- choose the **copying direction**
- choose **subfolders and files**
- specify the **Sync Operation Mode**

Copying Directions: Left To Right / Right To Left
It is very important to select the correct direction(s) for the synchronization. You may want to

allow the synchronization to copy files in both directions, so that the newest version of either location is selected for each file. On the other hand, especially for backup purposes, you may already know that only one copying direction will be needed. Somer other settings are also dependent on the copying directions: some checkboxes will be disabled when only one direction is chosen, while others are only available when both directions are checked.

Include Subfolders
Specify if no, all, or selected subfolders on both sides are to be included in the synchronization process.

Choose Folders and Files...
This option will enable you to select files and directories from the main left/right folders. A new dialog window will come up. This dialog will show the folder trees of both synchronization sides. Please note that in most cases, you should only select files on one side in this dialog, because all folders selected on one side will automatically also be considered on the other.

Sync Operation Mode
The operation modes are also explained briefly on the dialog when you choose them.

Standard Copying
Standard copying will compare the files and folders present on each side, using the filenames, the "Last Modified" timestamps as well as the file sizes for comparison. Newer files and files that only exist on one side will be copied over to the other side. No files are deleted (except in a rare, not recommended case: if you choose Real Time Synchronization with Individual Events mode and allow deletions on the Real Time settings dialog).

SmartTracking
SmartTracking uses a local database in order to track changes that have been made between the various invocations of the profile. That way, the software will know whether a file has been moved on the left side, or on the right side. It can also detect whether a file has been deleted on one side, or whether the file has actually been added on the other. SmartTracking should always be used whenever you want to keep two locations in sync, both of which are being used for work. SmartTracking is not needed when you do a backup or mirror, or any other case where you synchronize in one direction only. SmartTracking has various additional features which you will see in the SmartTracking dialog box (click on the Configure button).

Exact Mirror
This option is available when copying in one direction only. The purpose is to create an exact mirror of the source path. This may even include deleting files, or overwriting newer files with older ones. Therefore, the option should be used with caution. However, before overwriting newer files or deleting files, confirmation will be demanded from the user - unless the process is running unattended or scheduled, and the Unattended options allow these actions to be carried out without user confirmation.

WARNING: if you want to run the profile in unattended or scheduled mode, be sure to go to the Unattended tab sheet and make the appropriate checkmarks there so that the software will indeed delete files etc. - if you want that to happen.

You can click the **Configure...** button to fine-tune the Exact Mirror behavior. In particular, you

can specify a delay for deletions, and you can choose that excluded files are removed from the destination even if they exist on the source side.

Move Files To Destionation
This option will move the files from the source to the destination. The files will be deleted from the source side. You can also move files to two destinations at the same time: one is specified as the Right-Hand Side, and the other one is specified under Files->Deletions: Move Deleted Files Into A Specified Folder.

5.2 Schedule

Syncovery's Scheduler can run your profile automatically and repeatedly. Just specify the desired scheduling below. After saving the profile, please make sure that the Scheduler is running by opening the Scheduler tab sheet in Syncovery's main window. As an alternative to regular scheduling, you can use Real-Time Synchronization. And you can always run jobs manually.

On this tab sheet, you specify whether the profile should run automatically when the Scheduler is turned on.

Schedule This Profile
Place a checkmark here in order to set up your profile for repeated automatic execution. As soon as the schedule is configured on this page, you can invoke the Scheduler by choosing its tab sheet from the main dialog and starting it there.

Run Every Day (Or Specified Weekdays)
To set up the profile to be scheduled once every day, select this option and specify the desired time of day. Also select this option if you want to run it once per day on specified weekdays only, even if it is only one day in the week. The weekdays are selected on a different tab sheet.

Repeat After
Select this radio button and specify the exact interval that you want to have between each profile run. If you want to run it several times a day, make sure the number for days is zero (0).

Specify Next Run
Place a checkmark here to set the time for the next profile run. You can remove this checkmark to have the scheduler choose a time on its own.

Interval specifies the idle time between runs
With this checkbox seleted, the scheduler will wait for the specified amount of time after a profile has completed before starting it again. For example, if the interval is one hour, and the profile takes one hour to complete, then it will be started every two hours. Without this checkbox, it would be started every hour. Or if it takes longer than an hour, it would be started again immediately after having completed.

5.3 Schedule/More

This tab sheet contains additional options for scheduling.

Note that jobs running on logout or shutdown should not take more than one minute because Windows will otherwise abort the logout or shutdown and return to the Windows desktop.

Run Upon Windows Login
This checkbox will cause the profile to be run in the background as soon as you log in to Windows. This is related to the scheduler setting "Auto-Start Background Scheduler With Windows Login" so you must leave that checked too.

Run Upon Windows Logout
With this checkbox selected, the profile will run when you log out of Windows. This will work as long as at least one Syncovery process is running and you see the icon in the system tray. The background scheduler would be sufficient for this to work. The timer running as a service cannot perform this task however.

Run Upon Shutdown or Reboot
This checkbox causes the profile to run when you shut down or reboot the computer.

Run Missed Daily Jobs Immediately When Scheduler Starts
Jobs might be missed due to the computer being turned off or the scheduler not running. This setting causes missed jobs to be started immediately, when the scheduler is started again, rather than waiting for their next scheduled time.

Add a Random Delay
This option is useful if you want to use irregular intervals, or a different time of day for each profile invocation. One reason for this could be many computers trying to back up their data to the same server, so that you want to have the load spread over a period of time.

Warn If Profile Not Run For X Days
Syncovery can remind you if a profile has not been run for a number of days. This feature is actually independent from the other scheduling settings, and it can be used for profiles that you run only manually also. If a profile has not run for the specified number of days, a dialog will pop up allowing you to run it, or to dismiss the warning.

Additional Times
You can specify additional times of day where the profile should run. However make sure to specify a main schedule first. These times will only be used if the profile has a schedule on the "Schedule" tab sheet.

5.4 Weekdays and Time Window

The settings on this tab sheet allow you to limit when a profile can run. You can specify weekdays as well as a time window for the profile.

Weekdays
If you don't want the profile to run each day, please uncheck some of the week days. Specify the time of day or the interval of profile invocation on the Basic or Advanced tab sheets.

Time Window - Run Only Between

Here, you can specify a time window during which the profile may be executed. However, this setting alone will not cause the profile to be run by the scheduler. Specify the time of day or the interval of profile invocation on the Schedule tab sheet.

You can choose to **Ignore this limitation on Saturdays and Sundays**, as well as specify whether running profiles should be **interrupted** at the end of the time window, and whether copying should stop even **in the middle** of a file.

5.5 Monitoring / Real-Time Sync

This tab sheet contains settings related to real-time synchronization and monitoring. Real-time synchronization depends on file system events, which are usually available for local folders and folders in the LAN. In addition, some cloud storages can be polled for changes. It is also recommended to combine real-time sync with a regular schedule to ensure that a full folder comparison is done.

Be sure to read our documentation on Real-Time Synchronization.

This tab sheet serves two different purposes:

- Real-Time Synchronization
- Automatically running a profile when a drive becomes available or is connected

In addition, you will find the checkmark **Use a never-ending profile run** *(formerly called Continuous Sync)*, which may be useful when connecting to servers with Internet Protocols. It causes the syncs to occur as one profile "run" that repeats continuously without logging out from the server connection - as opposed to the normal behavior where Syncovery completes each profile run and logs out, and starts a new run at the next scheduled time, or when the next real-time events occur.

Warning: use the **never-ending profile run** only with an FTP or SFTP server, and avoid it for any other types of syncs.

The setting at the bottom, **Minimum pause between actions** applies to all of these features. It can be used to avoid constant activity and to make sure that the profile isn't run too often.

For all of these features, the Scheduler must be running because it manages the monitoring and profile running.

Run profile as soon as the drives or volumes involved become available
This feature will cause the invocation of the profile to be triggered by the presence of network paths or other devices such as USB drives.

The profile will be executed as soon as a drive becomes available, or as soon as this computer is connected to the LAN.

The presence of a drive is detected by simply trying to access it at regular intervals. Specify the

delay between access attempts here. In addition, you will usually want the profile to pause for a longer time when it has run successfully. You can specify the pause after a successful run by specifying a **Minimum pause between actions**.

5.6 Comparison

Syncovery compares files by their names, sizes, and modification timestamps. There are several options to fine-tune the comparison. Certain time differences can be ignored, as they may be caused by different file systems - but only if the size is identical. In addition, you can choose what to do if the size is different, but the timestamps are identical.

So, If and only if the file sizes are identical:

Ignore 2 Second Time Differences
If this option is checked, timestamps with deviations of two seconds or less are considered identical. This is needed when synchronizing files between FAT and NTFS file systems, because FAT file systems can only store even numbers as seconds.

Ignore Exact 1 Hour Time Differences
This setting may often be useful to compensate for different computers or filesystems and how they treat **daylight savings time**. Sometimes, the timestamps reported by the operating system differ by exactly one hour even though they are really the same. With this option checked, this problem will be solved, and it is irrelevant whether the left or the right side of the comparison is one hour behind.

Ignore Seconds
This setting causes the program to ignore seconds when comparing the timestamps of two files.

Ignore Timestamp Altogether
This setting should be rarely used. Files of the same size are assumed to be identical regardless of the timestamp. You should only use this setting if you know exactly why.

Never Copy, Adjust Timestamp Only
Sometimes you may have a copy of your files where the timestamps do not match. This function will adjust the destination timestamps to match the source. It works only in one-way sync scenarios. By adjusting the timstamps, you can avoid having to re-copy the files just to get the timestamps correct. However, you should only use this feature if you know that all files with the same size actually have the same content. Usually, this option is used only for a single, manual profile run, rather than using it permanently for regularly executed profiles.

When Size Is Different:
Some programs or cloud storages tend to modify the file size but keep the original timestamp when you open a file, even if you don't make any changes or save it. Here you can specify what

to do with such files. The available radio buttons differ depending on whether you are configuring a one-way or two-way sync.

5.7 Comparison/More

This tab sheet contains additional settings to fine-tune how Syncovery compares files and folder properties. The checkmark that is most frequently change is whether to compare and mirror Folder Timestamps. When it is chosen, you may see many folders shown in the Sync Preview, all with the proposed action "time->". This is because Syncovery tries to make the folder timestamps match. If you find this confusing, you can turn the checkmark off.

Compare / Mirror: File Attributes
This setting causes the program to compare file attributes, such as *Read-Only*, *Hidden*, and *Archive*. When files are identical except for the file attribute, the file is not copied again. Only the attribute is changed on the destination side.

Folder Attributes
Folder attributes can be compared/copied also. Most folders don't have special attributes, but some folders do - for example folders with a customized folder icon. The icon will only appear on the copy if this setting is chosen.

Folder Timestamps
If you are synchronizing in one direction, you may want the folders on the destination to have the same timestamps as on the source side. Newly created folders are automatically adjusted in this way, but existing folders may deviate. To have existing folders adjusted on the destination side, use this checkbox.

Adjust Case Spelling At Destination
Filenames on Windows are not case-sensitive. This means that myfile.doc would be the same file as MyFile.DOC. However, you may want your files to look identical. So, if you have a folder where some spellings are different in terms of case, then you may use this setting in order to adjust the case at the destination. You can only use this option when the profile has only one main copying direction.

Strict Case Sensitive Mode
If you are syncing two file systems that are case sensitive and allow duplicate file or folder names that only differ in their case spelling, then you can choose this option to enable Syncovery to sync correctly, and to allow these duplicate names. On Windows, this is not possible with local or UNC paths. It can only be used with FTP, SFTP, and some other case-sensitive protocols (such as the Google Drive API).

Strip Read-Only Attributes
Some storages may contain read-only files, or report all files as read-only (such as CDs or DVDs). When copying, Syncovery can make sure that the destination files will not have the read-only attribute.

Detect Hard Links
On recent versions of Windows, a file can have an entry in several folders. But it is the same

physical file. Normally, the program would copy such files to independent copies in the folders on the destination. To keep the files hard-linked on the destination, this checkmark must be chosen.

Enforce Hard Links For Existing Files
This checkmark can be chosen to update the hard-link status for existing files so that it matches the source side.

Always copy the files, even if timestamp unchanged
The program normally copies files only if they don't exist on the other side yet, or if the "Last Modified" timestamp has changed. However, some programs don't update the files' timestamps even if their content changes. This checkmark causes all files to be copied even if apparently unchanged.

5.8 Files

This tab sheet contains various settings related to files, such as detecting moved and renamed files. The option "Automatically Result" can be important, as it uses temporary file names during transfer. That way, incompletely copied files can never be mistaken for valid files.

Detect Moved Files and Renamed Folders
When files have been moved into different folders, or even when the folder structure has been reorganized, Syncovery can detect this. Rather than creating duplicates by merely copying files between the left and right locations, it will perform the same moves that you have done on one side, on the other side. Select this option to enable the detection of moved files.

A limitation of this feature is that moved files can only be recognized as such if they haven't been edited on one side. If they still carry the same timestamp and they still have the same size on both sides of the synchronization, then this software is able to detect that they are identical, even if they are not in the same folders.

You also need to specify whether the location of moved files should be adjusted on the left-hand side or on the right-hand side. On the side that you choose, the file will be moved in such a way that its new location matches that of the other side. Note that these radio boxes are automatically adjusted when you change the main direction of the synchronization.

If you don't want to specify a fixed side (left or right) where the location of files should be adjusted to match the other side, you can use **SmartTracking**.

Detect Renamed Files (based on timestamp and file size)
This feature can detect renamed files. It works only when file deletions are processed. If it sees a file that is to be deleted, and another one that is to be copied, and they have the same timestamp and file size, it will transform these two actions into a rename action.

Automatically resume (copy with temporary filenames, keep incomplete files when stopped)
This checkmark will enable resuming of interrupted copying operations even when the job is stopped and run again later. During copying, the destination filename will have an additional

extension, flagging it as a temporary file. When the copy is complete, the destination file will be renamed to the actual file name. The checkmark is recommended in many cases, because it can prevent incomplete copies on the destination from replacing good files, and thus avoid data corruption. It is also useful to prevent incomplete files from being processed by other software, or from being copied by other profiles.

Protect Files From Being Replaced With 0-Byte Files
This checkmark prevents empty files from being copied over a non-empty file.

Do not scan destination - copy all files regardless of their existence at destination
When copying only in one direction or moving files, you can select that the program does not check to see which files are currently present on the destination side. The purpose is usually speeding the profile up. This however means that all files will be copied each time unless filters are applied. For backup purposes, you may want to combine this setting with the Archive filter (*Copy Only Files With Archive Flag* and *Clear Archive Flags*).

Number of files to copy in parallel
Copying several files in parallel threads can be a great way to speed up the process. It may make better use of your network bandwidth, but may not be a good idea if you are copying extremely large files.

Copying Speed Limit (Bandwidth)
If you want to limit the bandwidth that the copying process will use, you can specify it here. The limit is specified in MegaBytes per second. As an example, if you want to limit to 1 MegaBit/sec, you need to type 0.1 here. Divide a MegaBit value by 8-10. You can use the **Scheduled Limits...** checkmark to specify different limits for different hours and days. If you need to use a really small bandwidth, you may also want to reduce the number of files being copied in parallel (on the Files tab sheet).

Skip Files Whose Size Is Changing
This setting helps prevent copying incomplete files which are just coming in or being created on the source side. Syncovery will wait for a few seconds to see if the file size is still changing. If so, it will skip the file. The profile should be scheduled in such a way that the file will be copied in the next run soon after. For example, such profiles can use a small interval such as every minute, or use real-time copying.

5.9 Block Level Copying

Block Level Copying copies only the changed blocks in a file. It can be a powerful way to save time and bandwidth, but it works only with eligible file types, such as databases, Outlook PST files, and disk images. It is not always beneficial since it incurs additional CPU usage and may result in a large sync database and additional processing time. It works mainly with local destinations, network folders, and SSH/SFTP servers.

You will find all details in a separate chapter in this manual.

There are several ways that changed blocks can be determined. Please choose whether to use Checksums, File System Monitoring, and whether checksumming should use the Remote Service for the right-hand side. If the Remote Service is not used, a local database will remember the checksums for the destination side. If the Remote Service or File System Monitoring is used, block level copying can work with a remotely stored source file too. Otherwise, it is recommended to install Syncovery on the same machine where the data is stored, so that block level changes can be detected quickly.

Use Remote Service to Copy Files for Block Level Copying + Versioning
When combining block level copying with versioning, Syncovery can either upload the changes in separate compressed files every time (= Synthetic Backup), or it can use the Remote Service to make a copy of the file, so the older version can be kept separately, and the newest version can be updated on a block level. So this checkmark helps combine block level copying and versioning if you don't want to use compression, encryption, and Synthetic Backup. The Syncovery Remote Service must be installed and configured on the destination system.

5.10 Deletions

Syncovery offers various ways to treat deleted files. On many storages, they can be moved to the Recycle Bin. As an alternative, deleted files can be moved to a specified folder. This feature is compatible with all storages and Internet protocols, and Syncovery can remember the deletion date by encoding it into the file name.

Use Recycle Bin
These options cause Syncovery to try to use the recycle bin for overwritten files and/or files to be deleted. That way, you will have a chance of retrieving older versions that were replaced or deleted, if necessary. Please note that the recycle bin is not always available, especially not on network drives or FTP servers.

Move Deleted Files Into A Specified Folder...
This setting provides a more reliable alternative to the recycle bin. Rather than deleting files, they will be moved into a special folder which you specify. This setting works even with network drives or FTP servers.

Remember Deletion Time and Date
When Syncovery moves deleted files into the "folder for deleted files", it usually doesn't remember when the deletion occurred. When this checkmark is chosen, Syncovery will encode the deletion time and date into the filename in the folder for deleted files. That way, when restoring files, you can restore exactly the files that existed at any fiven point in the past.

Never Delete Any Files (Only Folders)
This checkmark allows Exact Mirror to be used to mirror the folder structure but prevent any files from being deleted from the destination.

Double-Check the Non-Existence of Files Before Deleting them on the Other Side
This checkmark causes the program perform a double-check of files before deleting them from the other side. Deletions can occur in Exact Mirror and SmartTracking mode if files have been

removed from one side. However to rule out any uncertainty the program can double-check if such files are really no longer there. This might be necessary if unstable network connections might cause files not to be seen even though they are still there.

Delete Older Versions Permanently
This setting causes older versions of deleted files to be removed from the destination.

Execute the File Deletions Phase Before the Copying Phase
To free space on the destination, file deletions normally take place before new and updated files are copied. However to protect the integrity of the data on the destination better, you can uncheck this option so that the deletions are done last.

5.11 Files/More

This tab sheet contains additional options for fulfilling less common requirements.

Use Windows API Copying Function
This option causes the program to use a file copying algorithm provided by Windows, rather than using its own copying routine. This affects only the copying process of each individual file, not the general behavior of the software when processing a profile. The Windows API Copying function is sometimes more compatible with some devices, so this option is automatically checked for removeable devices, network drives, and drives connected via USB. However, the Windows API Copying function may sometimes be slower than the internal copying routine.

Copy only X files per run
When this option is selected, the profile will stop after copying the specified number of files. The remaining files, or another X files, will be copied during the next profile run.

Copy only X MegaBytes per run
This setting allows you to limit the total amount of data to be copied during a single profile run. The amount is specified in MegaBytes, so for 1GB you would type 1024, or for 100GB it would be 102400.

Never Replace Any Files
A rarely used option. This setting causes missing files to be copied over, but existing files won't be updated even if newer versions exist on the other side of the synchronization.

Don't Add Any Files
Use this setting if you want to update existing files only and not copy any new files from one side to the other. This is also rarely used.

Ignore Global and Group Speed Limits
This causes the global speed limits to be ignored. The global speed limit can be specified on the Program Settings dialog. Group Speed Limits can be specified by right-clicking a profile on the Profile Overview screen. This option is enabled only if the Global Speed Limit Mechanism is

activated on the Program Settings dialog, tab sheet Advanced.

Always Append Smaller Destination Files (Use Only For Log Files)
This setting causes destination files to be resumed when the source file is larger. The program does not verify if the file is still the same, so this option is not usually employed in order to resume transfers. For general resuming, use the Auto-Resume feature from the Files tab sheet. A good example of where the Auto-Resume feature cannot be used so that you must use the setting "Always Append..." would be log files which get bigger each day, and you only want to copy the new information each time. In some cases, you may need to combine this option with the following one so that the program uses the suitable copying direction.

Always Consider Larger Files Newer (Ignore Timstamp)
This settings causes the timestamp to be ignored, and files to be compared only by their size. This is not normally recommended but may sometimes be necessary.

Check Destination File (Again) During Copying Phase
Normally, the program checks the existence of the destination file just prior to copying each file, in order to see if it is going to be overwritten. However, some NAS devices respond slowly to these checks if the folder contains many files. Therefore, you can disable this setting in order to speed up the process. Similarly, you can choose if the check should be done **Via Internet Protocols Too**, as it can be a little slower depending on the protocol.

And Compare File Details (Again)
This checkmark also causes the Last Modification timestamp of the files to be compared again just prior to copying. If the situation has changed and the file no longer needs to be copied, it is skipped.

Copied Files Receive Current System Time as Modification Timestamp
This option is rarely used. Normally, files copied to the destination are given the same timestamp that they had on the source side. However, if you would like to see the time of the copying operation as modification timestamp on the destination, this option will do it.

Preserve Last Access Timestamp On Source
By default, Windows may update the Last Access Timestamp of the source files when they are accessed for copying. Using this checkmark can prevent that.

Create Links to Source Files Instead of Copying Files
This interesting but rarely used feature will cause Syncovery to generate Windows links (lnk files) on the destination side, which point to the source file on the other side. No files are copied.

Bypass File Buffering
In some rare cases, the file buffering done by Windows can lead to problems, such as excessive memory usage for caching. Bypassing the buffer can sometimes solve such problems, and can speed up the copying in some cases. In other cases, it might slow it down. Whether this option improves the performance needs to be tried out. The default setting is not to bypass the file buffering. Use with care and only if absolutely necessary.

Use Temporary Files On Local Drive When Copying From Cloud To Cloud
Since version 11, Syncovery can avoid using temporary files from copying from one Internet storage to another (using an Internet Protocol on both the left and right side in the profile). In many cases, this is quite beneficial and faster than with temporary files. However, sometimes you may prefer using temporary files, and you can choose it here.

5.12 File Access

Volume Shadowing
Volume Shadowing is a feature available since Windows XP which allows the program to copy locked files (files which are in use by another program). You can specify if volume shadowing should be used, and when it should be used.

Since volume shadowing takes some time and causes additional CPU and disk usage, the default setting is to use it only to copy locked files. If there are no locked files, it is not used at all.

However, in many cases it is better to use the option "Use For All, Create Shadow Before Building File List". This ensures that a consistent snapshot of the folders is copied. It avoids changes or discrepancies between the building of the file list and the copying phase.

Using volume shadowing for all files also avoids any problems with applications trying to write to a file while it is being copied.

Volume shadowing is only available for local drives, not for network drives. This means that the software must be installed on the computer where the data is stored, rather than another computer on the network.

Dabase-Safe Copying: Enforce Exclusive Access
The database-safe mode makes sure that no other application can have a file open with write access while it is being copied. If an application is currently working with the file, the program will try to use volume shadowing, or it can be set up to wait until the file is free. If volume shadowing is not used while it is being copied, an application trying to access the file in read/write mode will fail and issue an error message. As soon as the copying is finished, the file can be used nomally again.

Take Administrative Ownership Temporarily If Needed To Replace (Update) Files
This feature works when the program is running "as administrator", or the scheduler is installed as a service and it has administrative privileges. It can help copying files which are normally not copyable by the account that the program is running under.

Prior to copying, verify that all files can be copied by opening them
This setting is rarely used. It ensures (as much as possible) that the copying is only done if all files can actually be copied. That way, you will either get all files copied successfully, or none. To make this possible, a separate step (checking all files) is added before the actual copying starts. If any file is not accessible on the source side or is not writeable on the destination, then the

profile will be stopped. Usually this setting is used when profiles are run via scheduler, so that another attempt to copy the files will be made at the next scheduled time.

5.13 Folders

Create Empty Folders
Here, you can specify whether empty folders are replicated. The setting can also be used to avoid creating a empty folders which are shown as empty only because of your filename masks (inclusions/exclusions).

Remove Folders That Were Emptied
Rarely Used! This is not necessary to create an Exact Mirror. In Exact Mirror mode, non-existent folders are deleted from the destination by default. In SmartTracking mode, deleted folders are deleted from the other side without this setting, too. This option is something different (see below).

If you would like folders that are emptied during the profile run to be removed also, activate this setting. One use case for the setting is the operating mode "Move Files To Destination" where you can choose to delete emptied folders on the source side with this checkmark.

On the right side, create a new folder each time
If you are using the profile for backup purposes, you may want to direct your regular backup into a new folder each time. This folder will be created, and it will contain the current date of the backup in the folder name. You must combine this setting with the two "Archive" flag settings on the Filters tab sheet. The "Archive" flags will signal to the software that a file needs to be backed up. Because older backups are contained in various separate folders, the software cannot check whether a file is already contained in an older backup. Therefore, it must depend on the "Archive" flag if you are directing each backup into a separate folder.

Flatten the Right Side to Max Subfolder Levels
Sometimes you may encounter a requirement where the destination side should have no subfolders, or have a hierarchy that is only 1 or 2 levels deep. If you would like to move or copy files from a folder tree on the left side into a flat folder without subfolders on the right side, then use this setting. Files from any subfolder on the left-hand side will be copied into the base folder on the right-hand side. In most cases, you will specify the maximum subfolder level as 0.

5.13 Folders/More

Never Delete Any Folders (Only Files)
In Exact Mirror or SmartTracking modes, this setting will cause empty folders to remain even if the logic would normally cause them to be deleted.

After Copying and Other Actions, Touch Affected Parent Folders
On the
Left-Hand Side / Right-Hand Side

This feature was introduced in Syncovery 9. It is automatically activated when doing (near)-Real-Time synchronization over Internet Protocols such as FTP, SFTP, or WebDAV. To use it in other scenarios, the feature has been made available on this tab sheet.

When this functionality is chosen, Syncovery will update folder timestamps across the complete hierarchy so that they reflect recent changes. The result is that another Syncovery job (on another machine) can use these timestamps to quickly find the folders with changes, and perform a quick scan rather than a complete folder scan.

This feature is necessary because most file systems do not update all parent folder timestamps when a change has occurred in a subfolder.

Scan Destination Subfolders Even If Not Present On Source Side (To Detect Moved Files)
This behavior is only optional in Standard Copying mode. The checkmark determines whether Syncovery will scan all subfolders on the destination, or only those that also exist on the source side. Traversing into all subfolders can help detecting moved files to avoid re-copying them. In the SmartTracking and Exact Mirror modes, Syncovery will always look into all subfolders on both sides (except deselected or excluded ones).

Ensure Folder Timestamps Reflect the Most Recently Modified File in the Subtree
This option offers a solution to the following problem: in Windows, folder timestamps are updated when a file in the folder changes or a new file is added. However, only the direct parent folder of changed files are updated. Folders further up in the hierarchy will still show an older modification date. For this reason, folder timestamps cannot safely be used as a filter because an older date in a folder doesn't mean that there aren't newer files in any subfolder.

When making a copy of a folder hierarchy, Syncovery can make sure that on the destination side, the folder timestamps do really reflect the most recently modified file in the complete subtree. This folder timestamp adjustment is only made on the destination side. The use case for this is that if you use additional profiles to distribute your files further from this profile's destination, then you will be able to use folder timestamps for filtering.

Use Intermediate Copying Location
This setting will copy the files to an intermediate location first. When this part of the copying is complete, the files will then be moved to their final destination. This is rarely used.

Create Folder Symlinks from Destination to Source - Don't Copy Any Files
In this very special mode, Syncovery won't copy any files and won't create any folders. Instead, it will scan only the topmost level (the profile's base folders) and will create symbolic links on the destination side for any subfolders, and each symbolic link will point to a subfolder on the source side.

Split Job Based on Folder Mask...
This option makes it possible to save memory by running jobs in smaller subsets. Please use this only based on instructions from our tech support. It is usually not necessary because Syncovery can split jobs after X million files. The setting for this splitting is on the Advanced tab sheet of the Program Settings dialog.

5.14 Job-Related Settings

This tab sheet provides important settings, allowing you to make network connections, executing custom scripts and configuring additional job-related behaviors.

Network Connections...
Click here to specify network connections that should be made when the profile is run. This can be necessary to ensure that the scheduler can access network shares.

Check Free Space Before Copying
This checkmark allows you to turn off the free-space check in case of problems.

Execute command or script before / after...
This option allows you to execute programs or script at the beginning and/or at the end of the profile.

Pascal Script...
Allows to customize the job using the Pascal programming language. Please see the documentatation on our web site.

Override Email Settings...
A separate dialog box will appear which allows you to override some of the email notification settings in each profile. For example, you can specify whether emails are sent for this profile, and override the recipients list.

Access Local Files As User...
If you need to run this profile under a different user account in order to be able to access specific files, then you can specify the credentials here. Use this only to gain access to local files! This is not related to accessing files on a network share. Please use this feature rarely, and only if absolutely necessary. In most cases, it is better to run the scheduler as a service and give the service an appropriate account, or let it access local files with the SYSTEM account.

Verify right-hand side volume name
This setting is used when you make sure that the profile is only run with the correct backup medium inserted. When you click on this checkbox, the current volume label is read from the drive containing the right path. In the future, the Syncovery will verify that the disk inserted has the same volume label before running this profile. You may, however, use multiple disks with the same volume label. If the volume label does not match, the program will check if any other drive letter matches and will replace the drive letter automatically. As an alternative to this checkbox, you can simply write the volume label in the left or right base path for the synchronization, for example:
USBSTICK:\My Files *(where USBSTICK is the volume label of the desired drive to use)*

External copying tool...
This option allows you to set up external command-line tools which will do the copying instead of the internal copying algorithms.

File List Threads
Specify how many parallel threads should be used when scanning the folders. The default

numbers are specified on the Program Settings dialog, tab sheet *Performance*. A higher number of threads may make the "building file list" step faster, but will also use more CPU. The number should not be too high. In most cases, between 3 and 10 seems reasonable, but you can go much higher. Some settings and cloud servers may cause only one thread to be used, so the setting may be ignored in some cases.

Ignore Internet Connectivity Check
On the Program Settings dialog, tab sheet *Startup*, you can choose to run jobs only if a specific Internet connection is active. The feature allows you to avoid making cloud backups while on a mobile connection, for example. This checkmark here allows you to skip the Internet connectivity check for this profile, which you may want to do for local copying.

Show Checkboxes In Preview
This setting causes checkboxes to be shown in the Sync Preview to allow you to easily choose the files to be copied. It is mostly used by the Restore Wizard.

Shutdown after this job
Specify in which situations the computer should be shut down when the job completes.

5.15 Inclusion Masks

Filename Mask(s)
Specify which files to process. Normally, **(all)** or ***** means all files. Alternatively, you can specify filename masks such as *.doc or *.xl?. You can also specify complete file names here, but no folders.

The masks, or the file list to copy, can also be read from a text file. In that case, start the line in the Inclusion Masks edit field with a colon : and type the full path of the text file directly after it.

If the text file contains only subfolder selections, start the line with :+ like this:
:+C:\subfolders.txt

The subfolder text file can contain lines such as:
\folderA
\folderB

Specify Folder Masks...
A separate dialog will come up allowing you to specify folder name masks. For example, you can use this in order to find all folders with the word ARCHIVE in their names and have them processed.

Restrictions...
Rarely used. Allows you to specify separate file masks for one of the two possible copying directions.

Include backup files (*.$?$.*)
When using the setting to keep multiple backup versions of the same file, Syncovery may create backup files such as MyFile.1.doc. In order to avoid these files being copied back from the

backup location to the main data location, these backup files are normally not included in any synchronization. However, if you want to override this hard-wired exclusion, you can place a checkmark next to this setting.

5.16 Exclusion Masks

Exclusions (files & folders)
You can exclude certain files or folders such as *.tmp or *.bak. You can also specify complete file or folder names here. You can enter only the name to exclude, or full paths or a part of the name combined with wildcards. Wildcards to use are * and ?.

Trailing backslashes are not allowed. Folders are specified in the same way as files. Wildcards are only allowed in the name itself. So if you specify a full path, only the last part (the name) can contain wildcards.

The recommended way to enter full paths is to use them relative to the profile's base paths. So you start with a backslash and then type only those folders in the path that are relative to the base.

Examples:
*.tmp
\TEMP
\Users\All Users
.*
\My Documents*.jpg

The exclusion masks can also be read from a text file. In that case, start the line in the Masks edit field with a colon : and type the full path of the text file directly after it.

Don't Copy versus Ignore Totally
When the Exclusion Masks are processed in **Don't Copy** mode, this means that such files won't be copied but they will still be processed for deletion if they exist at the destination and not at the source, and if you are using Exact Mirror Mode. In contrast, the **Ignore Totally** mode will not process files that match the exclusion masks in any way. This can sometimes lead to a situation where such files are left at the destination and folders that should be deleted can't be deleted because they still contain these files. Therefore, **Don't Copy** is the default setting.

Use Global Exclusion Masks Also
The Global Exclusion Masks are specified on the Program Settings dialog and they are applied to all profiles unless this checkmark is deselected. They contain masks that avoid copying of some temporary files as well as the Recycle Bin contents and other system folders that cannot and should not normally be copied.

5.17 General Filters

Process Hidden Files
This option determines whether hidden files are seen and handled by the program.

Search Hidden Folders
This option determines whether hidden folders are seen and examined by the program.

Copy Only Files With Archive Flag
This option is used for data backup purposes. Whenever a file is changed, Windows will normally set the Archive flag for that file. This signals that the file needs to be backed up. For normal synchronizations, the flag is not needed because the software looks at the left and right sides and decides for itself which files have been updated or added. However, in a data backup context you may sometimes use the software in such a way that the backup medium doesn't always contain all the files from previous backups. In that case, the software would start to back up old files again. To avoid that, you can use this setting. Use it in conjunction with the following setting:

Clear Archive Flags
This setting is used in conjunction with the previous one, also related to the Archive flag. For the preceding option to be useful, the last backup job for any given data folder should clear the Archive flag so that in order to signal that the file has been backed up, and doesn't need to be backed up again unless it is modified. If you have only one backup job for your data, and you are using the filter "Copy Only Files With Archive Flag", then you should check this setting too. If you have several backup jobs for the same data, please make sure that the last backup job that runs clears the Archive flag. Again, you only need to do this if you are working with Archive flags at all (and the preceding option is checked).

Restore Deleted Items
When restoring (in direction right to left), this will cause Syncovery to find files from 'Older' versioning folders if the main file version was deleted. So you can restore deleted files, if they still have a version in the 'Older' folders.

Skip Offline Files
This checkmark will skip files from your OneDrive folder if the file isn't physically present on your hard drive. It may work with other cloud storages too that have local folders representing the cloud content.

Copy Pinned Files Only
This checkmark will make Syncovery copy only those files from your OneDrive folder which have been "pinned", i.e. set for permanent local availability on your hard drive.

Use .$symlink files to save Symbolic Links
This feature allows you to back up symbolic links to storages that don't natively support them, or over any Internet Protocol (as Syncovery cannot create real symbolic links over Internet Protocols). Instead of creating an actual symbolic link, Syncovery will save the link information in a file with a .$symlink filename extension. When restoring, Syncovery will read these files and recreate the original symbolic links on the local side.

Analyze Reparse Points
This checkmark determines whether Reparse Points / Junction Points such as Symbolic Links are processed or ignored. You may want to ignore them if you get errors and don't need them copied, which happens sometimes with some of the default links that Windows creates, such as "My Music". The detection and processing of reparse points may depend on whether the job is being run under an Administrator account or not. For this reason, you may be getting different results in manual profile runs from those runs initiated by the scheduler when it is running as a service, because the service is not limited by Windows UAC (User Account Control).

Follow Junction Points / Symbolic Links to Files
This checkmark determines whether symbolic links (soft links) to files are followed (i.e. the content is copied even if it is actually saved in a different location on the source side). If it is not followed, Syncovery tries to recreate the link as a link on the destination side.

Follow Junction Points / Symbolic Links to Folders
This checkmark determines whether symbolic links (soft links) to folders are followed (i.e. the linked subtree is scanned as if it were a normal folder). If it is not followed, Syncovery tries to recreate the link as a link on the destination side.

Copy Other Reparse Points
Specify whether Syncovery should try to copy other reparse points as they are. These can be custom reparse points from other applications, as an example.

5.18 File Age and Size Filters

File Sizes Must Be Within
This option enables you to copy only files within a certain range of file sizes. Sample file sizes are: 512, 2.12k, 5M, 2G.

File Dates Must Be Withing
Using these settings, you can limit the synchronizing to files last modified or created within a certain period of time.

File Age
Rather than specifying a fixed date range, you may want to choose a certain age for the files that you want to synchronize. The age that you specify is always seen relative to the current date when the profile is executed.

Filter By: Last Modification / Creation
By default, the above timestamp filters take into account the date and time of the last modification of each file. In some cases, you may want to filter by creation date, so you can choose it here.

Apply To: Files | Folders | Both
To get the most reliable operation of the file date filter, please use the default setting of **Apply To: Files** only.

Folder timestamps under Windows are not consistent in such a way that they would be reliable

for filtering. Applying timetamp filters to folders will often result in undesired exclusions of files.

However, there are some use cases where you may want to apply the filter to both files and folders, or even to folders only. For example, if you are distributing files from a copy of your folders made by another Syncovery profile, then you may filter by folder timestamps if the other profile has the option "Ensure Folder Timestamps Reflect the Most Recently Modified File in the Subtree" chosen on the Folders tab sheet, or the new options to **touch parent folders**.

Target Date for Restore
This option is mainly used by the Restore Wizard. By specifying a date in the past, the program can automatically choose the appropriate older versions of files when restoring (i.e. copying from the right-hand side to the left-hand side). This works for backups created with the Versioning settings.

Scan Only Subfolders Whose Timestamp Has Changed Since the Previous Profile Run
Rarely used because file systems do not normally update folder timestamps correcly!
This feature was introduced in Syncovery 9. It is automatically activated when doing (near)-Real-Time synchronization over Internet Protocols such as FTP, SFTP, or WebDAV. Because only folders with new modification timestamps are read, the synchronization will be much quicker. To use it in other scenarios, the feature has been made available on this tab sheet.

This really works only when it is ensured that folder timestamps are all updated. File systems and cloud services will not usually do that. But Syncovery itself can ensure that folder timestamp are updated, with the option **After Copying and Other Actions, Touch Affected Parent Folders** from the **Folders** tab sheet.This could enable multiple Syncovery clients to perform 2-way syncs with a mutual storage location, where each client sets the folder timestamps to reflect the changes it has made.

5.19 Retries

Sometimes, errors can occur: files could be locked and in use, servers might be temporarily unavailable, there could be Internet connection issues, or files no longer exist when they are about to be copied.

Syncovery can wait until files become accessible, and it can re-run the entire job if there are errors.

Wait For File Access
In both Database-Safe and normal mode, it can occur that other applications prevent the program from accessing a file that needs to be copied. When that happens, the program doesn't just abort or skip the file. Instead it will continue copying the remaining files, and then (depending on this option setting) start retrying the files that couldn't be processed yet. You can specify for how long the program should retry the files. It will use only minimal resources during this period, so don't be afraid to set the retry-time to a high number.

Wait And Retry If Transfer Or Reading Problem
This checkmark can make the copying extremely fault tolerant. If there are any connectivity

problems, the program will retry for as long as you specify.

Re-run the profile if an error occurs

- while the building file list
- while running the profile

These settings allow you to have the whole profile re-run automatically if any error has occured. You can specify whether to re-run only once, or to re-run until it has completely succeeded. Also specify the number of seconds to wait before each re-run.

5.20 Safety/Attended Mode

These options specify which warnings Syncovery is going to generate in certain situations encountered during file synchronization. These options are only in effect when the program is running in attended mode (i.e. not scheduled, and not otherwise invoked in unattended mode).

Warn before running a profile that moved files and deletes from source
This warning will come up when a profile with the "Move Files" mode is started. The warning can be confirmed and disabled for the future when it appears, so you do not need to remove the checkmark on this tabsheet (it will be removed when the warning is permanently confirmed).

When not keeping multiple backup versions:

Warn before overwriting read-only files
In Windows Explorer, when you right-click a file and choose "Properties", you can choose to protect a file by choosing its *read-only* attribute. Some applications also protect important files by marking them read-only. This option determines whether read-only files are replaced with or without requiring confirmation from the user. Note that this warning only occurs when the program is not configured to keep multiple backup versions of each replaced file.

Warn before overwriting larger files with smaller ones
This is the warning that is most frequently turned off. When editing documents, the files don't always become larger. It often happens that the most current version is smaller than a previous one. If you don't want to have to manually confirm such cases, you can uncheck this option. Like the previous one, this warning only occurs when the program is not configured to keep multiple backup versions of each replaced file.

In Exact Mirror Mode, or when manually specified:

Warn before overwriting newer files with older ones
In Exact Mirror Mode, the program will try to make the mirror match the model exactly. This can mean that sometimes, files need to be replaced by ones with an older modification date. In that case, the software will normally require confirmation from the user. To avoid having to confirm this replacement, uncheck this option.

Warn before deleting files

Another thing that can occur in Exact Mirror Mode is that files need to be deleted from the mirror because they no longer exist on the "model" side. In that case, the software will normally require confirmation from the user. To avoid having to confirm these deletions, uncheck this option.

5.21 Special Safety

Syncovery includes additional safety checks to avoid accidental deletion of files. It is important to be aware of these safety checks because they can prevent intended deletions from being carried out if you ignore them. The safety checks must be individually configured or turned off in each profile.

To turn the safety checks on for Unattended Mode, please see Profile Settings - Safety/Unattended Mode.

When a profile is executed and more files are going to be deleted than allowed by the new safety checks, then the user must confirm the deletions even in unattended mode. Without the confirmation, the profile will be executed without the deletions.

The safety checks have been designed to catch the following situations:

- files have been accidentally deleted on one side of a synchronization
- incomplete or empty FTP directories are received due to bad communication

Especially when incompletely received FTP directory listings could cause the profile to delete files, it is recommended to turn these safety checks on (but you can change the pecentages that trigger the warnings).

Again, to turn the safety checks on for Unattended Mode, please see Profile Settings - Safety/Unattended Mode.

When Running in Background With Preview, Warn If Deleting or Replacing Newer Files
This option was added in Syncovery 9 because the "run in background with preview" does not show any other warnings. A warning dialog will appears when files need to be deleted or newer files are to be replaced with older ones. This dialog will enable you to confirm the deletions and replacements, or run the profile without them.

5.22 Safety/Unattended Mode

These options specify how Syncovery behaves in unattended mode, and at which times unattended mode is in effect.

Read-Only Files May be Overwritten When the Job Runs Unattended or Scheduled
When this option is checked, the software will automatically overwrite read-only files when running in unattended mode. In attended mode, it always demands a confirmation from the user. When this option is unchecked and the profile is run in unattended mode, then read-only files are skipped without any notification whatsoever.

Larger Files May be Replaced With Smaller Ones if the Smaller One is Newer

When this option is checked, the software will automatically overwrite larger files with smaller ones when running in unattended mode - if the files' modification dates require it. In attended mode, it always demands a confirmation from the user in this case. When this option is unchecked and the profile is run in unattended mode, then such files are skipped.

Exact Mirror Mode May Replace Newer Files With Older Ones
In order to create an exact mirror, the Exact Mirror Mode may sometimes have to overwrite newer files with older ones. Normally, it would ask the user to confirm this. However, when run unattended, this option must be checked or those files will not be replaced.

Unattended File Deletion Allowed
In order to create an exact mirror, the Exact Mirror Mode may sometimes have to delete files. Normally, it would ask the user to confirm this. However, when run unattended, this option must be checked or those files will not be deleted. In addition, please specify a percentage of files that may be deleted. For example, if you leave the figure at 90%, the profile would run **without deletions** if more than 90% of the files currently present on either side would be deleted.

In addition to the percentage, you also need to specify a maximum number of files that may be deleted in a single job run.

The program will usually **not delete all files** from either side of the synchronization, even if you specify 100% here. If deletion of all files can occur and you really want it to happen, then you need to create a dummy (text) file which always stays on both sides and which disables this security mechanism, simply by being there. If this is not feasible and you really need the profile to delete all files regularly, please contact support@syncovery.com for a solution.

Ransomware Protection: Replace Maximum % of Files on One Side
This feature allows you to specify a maximum number of files to replace in a single job execution. If there are more modified files than allowed, the job will stop.

Enable Special Safety Checks
These are new safety checks to prevent accidental deletion of files. **They may prevent your profile from being executed at all**. Please see the Attended tab sheet for details.

5.23 Special Features

Cache Destination File List...
The destination cache feature is a great way to speed up your sync profiles, since it makes the file list building phase very quick. However it can only be used for one-way sync jobs.

The file list of the destination side can be buffered in a cache, which is stored as a database on your computer. When this database is used, building the file list can be much quicker. However, you need to make sure that no other program or persons modifies the destination side, because otherwise it would get out of sync with the cache.

This setting will speed up running a profile by using a database as cache for the destination file list. It can be used in the following case:

- you are copying only from source to destination (not in both directions)
- nobody else, and no other profile, is modifying the files on the destination
- you want to speed up building the file list, and it needs the speed-up because the destination is relatively slow (such as FTP or WebDAV)

There are many other ways to speed up building the file list. Please see the documentation on our web site:
https://www.syncovery.com/documentation/faq/fastlist/

Process Security And Shares...
Need to preserve folder and file permissions, or re-create folder shares on the destination?

Just use this checkmark, which will open a dialog with the relevant settings.

A separate dialog will appear where you can specify how the program copied file security settings and/or folder shares.

Copying Order
Choose in which order files will be copied:
- Standard (Alphanumerical)
- Smallest First
- Largest First
- Oldest First
- Newest First

Left/Right side listing uses Remote Service
The Syncovery Remote Service can speed up building the file list if you are synchronizing across a network or even over the Internet. It is installed on the other end and generates the file list so that the main program can download it. These checkmarks are used in the main program to activate this feature. See also www.syncovery.com/remoteservice.

5.24 Special/More

Use Cache Database For Source
This setting is only used by the Restore Wizard. It enables you to use the database cache for the source side rather than the destination. This is not usually reasonable for profiles that run regularly. The typical application is a restore operation.

If the Destination Machine Modifies Received Files, Changing Their Sizes Upon Reception, Then Copy Such Updates Files Back
In some cases, when uploading documents to a server such as a Sharepoint Server, the server may add some additional information to the file. The file size then no longer matches. To fix the mismatch, this checkbox will cause Syncovery to copy the updated file back from the server shortly after having uploaded it.

Set Target Volume Label
This function modifies the volume label of the destination drive (when copying only in one

direction).

Double-Check Each File's Destination Timestamp After Copying
This option may help ensuring that modification timestamps are set correctly when files are copied to local hard drives, mounted images, or network drives.

Detect Changed Files Via File System Monitoring Even if Timestamps Unchanged
Some files, such as VM images or TrueCrypt containers, may keep their timestamps unchanged for some time even though the contents changes. Syncovery can now detect changes using the File System Monitoring Service. This is usually combined with the Block Level Copying and File System Monitoring checkmarks from the same tab sheet. This option should never be used for normal documents or image files or other normal files.

Spawn Separate Sub-Jobs For Each Subfolder
This feature will make Syncovery treat subfolders as separate jobs, by simultaneously spawning separate sub-jobs that handle each subtree. You can specify the subfolder level on which the separation into sub-jobs should take place. This feature may help speeding up the synchronization of large folder hierarchies with many folders and files. However, it should rarely be used, since the progress display and the log files can become confusing.

Alternate Data Streams
Alternate Data Streams are normally copied along with the file without any special setting. If you need special handling of Alternate Data Streams, you will find some options via this button. Note that Alternate Data Streams are not copied when zipping or encrypting files, or when copying via Internet Protocol.

5.25 Database

These settings enable the program to share the same database between several profiles. Normally, you don't need to edit any of the fields on this tab sheet. When saving a profile under a new name, you will be asked if you would like to continue using the same database.

Sharing the same database may be useful or necessary only in special cases, for example if you have more than one profile that work on the same folders but with different file masks or subfolder selections.

In most cases, you will want to avoid two jobs from using the same database. You will find a checkmark on the *Database* tab sheet of the Program Settings dialog to "never allow two profiles to share the same database". If two profiles do share the same database, only one of them can run at the same time.

Open Database Read-Only
This setting is used by the Restore Wizard when using the database as a file list cache for the source side to ensure that the restore operation can't modify the database.

Fast Mode: Don't Check All DB Entries
When using SmartTracking or *Cache Destination File List* or Exact Mirror with Delayed Deletions, the program depends on the database containing a complete, up-to-date list of all files with

their details and properties. For this reason, the whole database is checked against the actual file list when a job is run in attended mode. To save time, you can choose this **Fast Mode** checkmark.

5.26 Verification

There are two main ways to verify files and ensure that their content is 100% identical on both sides: Verification of files directly after copying, and verification of all existing files. Verifying copied files is recommended and while it makes the copying slower, the impact is still reasonable. On the other hand, verifying all files which already exist on both sides should usually not be done every time the profile runs. The "Remember Results" feature helps avoiding the re-verification of files already verified previously.
Checksums (or hashes) can be used to speed the comparison up. Syncovery will use them automatically if a cloud service provides them. In addition, you can install the Syncovery Remote Service on the other machine to generate checksums remotely, which greatly speeds the process up.

Verify copied files (can double execution time)
This checkmark will cause the program to verify the file contents after copying. The verification is done by reading the complete source and destination files once more and comparing them. So this will potentially double the execution time, drive access and network traffic. To avoid this, it is possible to use a remotely generated MD5 checksum for verification. The checksum can be calculated by the Syncovery Remote Service running on the other end, or by an FTP server supporting the MD5 command. The MD5 comparison option must be chosen under Comparison/More.

Binary Comparison Of Existing Files While Building the File List
Use this setting to verify that the file contents of existing files is 100% identical on both sides. Only files which would normally be considered identical by looking at the file size and timestamp will be verified. If the binary comparison shows that the files are different, they will be marked for copying in the profile's main copying direction. If the profile has both directions checked, then the Synchronization Preview will show such files with the word CONFLICT and it will not do any copying by default. You can then set the desired copying direction in the Synchronization Preview.

This checkmark should not usually be used for jobs that run regularly, because it causes a lot of disk usage or network traffic. Moreover, this checkmark does not cause newly copied files to be verified. To verify files after copying, use the checkmark *Verify copied files* on the Files tab sheet.

However, in conjunction with the new **Remember Results** feature, it is actually feasible to leave this option checked for repeated runs. Syncovery will remember which files have verified successfully, and not repeat the Binary Comparison for them in future runs. This also makes it possible to interrupt jobs during the comparison phase. When the job is started again, it will skip the binary comparisons that were previously successful.

Use Remotely Generated MD5 Checksums For Comparisons
This checkmark will make the program use MD5 checksums for verification rather than a complete binary verification. To make this work, the Syncovery Remote Service must be running

on the other end. Do not choose this checkmark if you want MD5s to be used with some cloud services that support it. That will be done automatically.

Verify Sync Statistics After Completion
This checkmark should be used only rarely! It will cause the left and right folder trees to be scanned completely again after the synchronization has completed. This verification stage will end in detailed statistics which are used to verify that the backup or mirror is complete. Any files or folders that still don't exist on both sides are listed.

5.27 File Integrity

Syncovery 11 is aware of a large number of file formats, and can analyze most of your files to verify the integrity of the internal file structure. The verification depth varies depending on the file types.

The results of the integrity checks are shown in the log file, as well as the result summary. Unknown file types will be skipped. Sometimes, an integrity check may fail even though the file is still usable.

This new feature does not use AI and does local processing only (no cloud service).

5.28 Versioning

Versioning allows you to keep multiple versions of each file on the destination side, or even on both sides. Older versions are usually renamed so that multiple versions of a file can reside in the same versioning folder.

A related special feature is "Filename Encoding", which will add the modification timestamp to all files on the right-hand side, not just the older versions.

Keep Older Versions When Replacing
When files on one side are replaced by newer versions from the other, Syncovery can keep the older versions by renaming and/or moving them. You can specify how many backup versions to keep.

You can choose between various renaming ways:

Add .1., .2. etc.
For example, a file named letter.doc would be renamed to letter.1.doc. The next older version would be letter.2.doc, and so forth.

Add Timestamp
This way of renaming adds the timestamp to the filename.

Filename Encoding: On Right Side, Put Timestamp Into All File Names
(Also Preserves Timestamp on Servers That Normally Lose It)
Filename Encoding is the recommended way to do a versioning backup as well as preserve the timestamps on sites that don't accept timestamps.

Putting the timestamp into the file names can serve two purposes when backing up data:

- Some cloud servers don't accept the original file timestamps natively. To serve as a fully transparent side for the synchronization, the timestamp has to be retained somehow. Filename Encoding can be the solution because the time stamp will be shown fully correct in future runs of the profile. The encoded filename will be decoded by Syncovery and the original filename will be shown with the correct timestamp.

- To keep multiple versions of the same file on the destination.

Filename Encoding is also the way of renaming used when using Synthetic Backup.

See also Versioning and Filename Encoding on our web site.

5.29 Synthetic Backups

Synthetic Backup is the combination of Partial File Updating with Versioning, Filename Encoding and Zipping. It can be used with any backup storage location, including FTP and other Internet Protocols. The zip files can be encrypted. Synthetic Backup can be described as Differential Versioning.

Synthetic Backup combines versioning, compression, and block level copying. It allows you to back up larger files with incremental backups containing only the changed blocks. You can always restore the latest or previous versions of the file. As opposed to pure block level copying, Synthetic Backup works with all types of backup storage..

See also the documentation on our web site.

Checkbox to choose:
- tab sheet Versioning->Synthetic Backups: "Use Synthetic Backups"

The following dependent options are then checked automatically:
- tab sheet Special: "Partial File Updating" (without Remote Service)
- tab sheet Zipping: "Zip Each File Individually"
- tab sheet Versioning: "Filename Encoding"

All incremental versions must be kept in the same folder.

Be sure to keep all zip files on the destination as they may all be needed to reassemble the file upon restore. However, the number of older incremental parts that need to remain on the backup storage can be limited thru the new "Checkpoint" feature. You will find its settings on the new tab sheet Versioning->Synthetic Backup.

A checkpoint is an incremental backup just like the daily backups. However, it is a little larger because it includes the changed blocks of a longer period of time, so that many preceding incrementals become unnecessary.

Any older version can be restored by choosing the desired target date in the Restore Wizard, or by right-clicking the file in the Sync Preview and choosing the desired version. The option "Keep multiple files" is implied and does not need to be checked. It cannot be used to limit the number of older versions. However, if this option is specified on the Versioning tab sheet, it will be used as a minimum number of versions to keep even if the Checkpoints feature would allow to keep fewer versions.

A restore of the synthetic backup files can be done independently from the original job, to any destination. No database is needed for the restore - just the files.

Block-level copying is applied to files which are at least 400,000 bytes in size. The granularity of the differential backup is determined dynamically for each file, based on its size. The smallest block size used will be 2048 bytes.

Each partial backup includes an MD5 checksum for each file, so that restored files can be verified to be 100% correct. You can see this in the file MD5.TXT that the partial zips contain.

The Zip file sizes can now be limited too so that large files can be split (only in conjunction with Partial File Updating). The transfer can be stopped any time and zip parts already uploaded won't be lost.

5.30 Versioning/More

This tab sheet contains various additional settings related to versioning.

The first two checkmarks are the most important ones. During normal operation, Syncovery decodes certain filenames on the right-hand side. For example, it deduces the original file name from a compressed archive, if the archive name has been created and encoded by Syncovery. However, if you want to make another 1:1 mirror of such encoded files, then that mirror job should not decode any file names, and you need the first two checkmarks to be ticked.

Do Not Decode Left/Right-Hand Filenames When Building The File List
Filenames where the program has encoded additional information into the filename are normally shown in their original, unencoded form. However, if you would like to make a 1:1 copy of a backup location which contains mangled filenames, then you need to use this setting so that files are copied without decoding their filenames.

Timestamp Encoding Format
Choose the timestamp format to use for versioning. This is the format that is added to filenames if the versioning settings need such filename encoding.
- Syncovery Native (.dYYYYMMDD-uHHMMSS)
- Windows 10 File History (YYYY_MM_DD HH_MM_SS UTC)
Both encodings use UTC times (timezone independent world times).

Decode Timestamp Formats
Choose whether Syncovery should recognize and process only the chosen encoding format, or all known formats. Filenames encoded with timestamps are expected to exist on the right-hand side of a profile only, and are decoded for comparison against the left-hand side, which should

contain the actual, original, unencoded filenames.

Clean up identical duplicates of files with differently encoded Unicode umlauts
Unicode allows accented characters and some symbols to be encoded in different ways - for example, as one accented character, or as the base character followed by the accent as a separate character. Due to the different encoding, files whose names look identical can exist in the same folder. This option will attempt to clean them up.

Allow merging duplicate folders on source side
This will enable merging duplicate folders on the source side, even in a one-way sync. This is rarely needed - only if two folders with identical name exist as duplicates.

Remove parenthesized version numbers before extension on the right side: file(2).ext
Rarely used. This option will remove the version numbers in parentheses, like the example, from files on the right side, when copying right to left. Note that this checkmark is unrelated to any other versioning settings. The parenthesized version numbers are not something that Syncovery uses itself, it can only remove them.

Remove the versioning tags -1 and -2 and rename such files immediately while building the file list
Rarely used. This option helps cleaning up versioned files. These versioning tags may have been added by the conflict detection in SmartTracking mode.

Clean up all older versions based on the specified number to keep
If you change your versioning settings, you can use this checkmark to clean up the files and delete older versions which are no longer supposed to be kept. If unchecked, the limits are enforced only for files being copied.

Files Backup Up With V4 Used UTC/GMT (World Time) For Encoding (Like V5)
Up to version 4, filename encoding could use either the local time or world time. Because this was ambiguous, the filename encoding was changed to always use world time in version 5. If you still have older files, this checkmark is used to tell the program how to interpret the encoded timestamps.

5.31 Compression

There are **two ways** of using file compression. You can either

- **Compress Each File Individually**. One zip or sz file will be created for each file that is copied from the left-hand side to the right-hand side. The advantage of this method is that in future synchronizations, the compressed files are fully transparent and they will be shown in the Synchronization Preview as normal files. Since future synchronizations will see the compressed files as normal files, they will know exactly if they are still identical to the left-hand side. When copying these compressed files back (from right to left), they are unzipped automatically.

- **Use Compressed Packages** (Many files per archive). Putting many files in a single backup archive is also possible. However, this is not the native way Syncovery handles files. In future runs, the files inside the Compressed Packages are not seen by the software. Therefore, you may need to resort to the "Archive" flag in order to determine which files are new or changed and need to be backed up. Compressed Packages can also be used in order to speed up file transfer over a small or medium bandwidth network connection. If used that way, the Syncovery Remote Service must be installed on the destination machine and be configured properly.

There are also **three file formats** for file compression.

- **Zip format:** a standardized file format to ensure compatibility with many other applications. It offers reasonable compression capabilities. When uploading to a cloud server, or using an Internet protocol such as FTP, Syncovery creates a local temporary zip file first. If you back up large files, you need to make sure that enough temporary file space is available.

- **Sz format:** a new, proprietary file format. This is currently supported only by Syncovery. It is a streaming format that can be used to back up to cloud servers without using any temporary space on your local hard drive. It also features the two Ultrafast and Maximum compression modes which are both better than the zip format. Since Syncovery 9, the Sz archives can contain more than one file, so they are now compatible with the "Compressed Packages" feature.

- **.7z format** (Windows only): an efficient format known from the 7-zip compression utility. Needs temporary files for uploading just like the Zip format.

All file formats support the same level of encryption, using the AES 256 algorithm.

Compress Each File Individually (when copying left to right)
This method compreses each file individually as described above. The zip or sz file will have the same name, plus a few digits encoding the uncompressed size of the file, plus the .zip or .sz extension. The uncompressed size needs to be encoded in the file name in order to make the compressed file transparent for future synchronizations.

Use Compressed Packages (many files per compressed archive)
This settings enables the use of Compressed Packages as described above. A separate dialog with some related options will appear when this option is chosen.

Compress directly to destination
This applies to ZIP files only. Normally, compressed ZIP files are generated in a temporary folder and then copied to the destination. This setting changes that and generates the ZIP files directly in the destination folder. This setting cannot be used for FTP transfers.

Unpack all Zip or Sz files which are copied to the destination
Normally, the program unpacks only Zip or Sz files when restoring from the right side to the left

side, and only those "individually compressed" ones that it has generated itself during a previous run. To unzip any Zip or Sz file after copying it to the other side, use this setting. Both the compressed file and the unpacked files will remain in the destination folder.

Limit Zip File Size
This option allows you to specify a maximum file size for the zip files. When the size is reached, the program will begin a new split zip file. This works with both the individual zipping mode and the ZIP Package mode.

5.32 Encryption

File Content Encryption

Encryption can be performed when files are copied from the left-hand side to the right-hand side. Copying in the opposite direction or using the Restore Wizard will decrypt the files.

Encryption is always done along with compression, although you can use the .zip or .sz file formats without compression too. The recommended encryption method is AES encryption (256 bit). This encryption can be decrypted with some popular zipping tools, such as WinZip, WinRAR, and PowerArchiver (when using the zip format).

The password phrase should be longer than normal passwords since it is used to generate an encryption key. There is no longer any limit in the allowed pass phrase length.

PGP Encryption / Decryption

Syncovery can use PGP public and private keys to encrypt and decrypt files, as well as sign files and verify signatures. See our separate documentation on exchanging files with PGP.

Filename Encryption
You can also encrypt the file and folder names on the destination side. Please note that the encrypted file names cannot be deciphered by a human and will make it harder to handle your backup folders. Use this only if you are absolutely sure that you want to hide and encrypt your file names. The filename encryption algorithm is proprietary. Your file names cannot be recovered with any other tool.

5.33 Error Handling

Ignore These Errors:
Some types of errors can be ignored, because they are bound to occur regularly and are not significant. By ignoring these errors, you can avoid emails and result summaries that contain errors. Use only if you need to ignore repeated error conditions. An example of an error that occurs frequently is "Files That No Longer Exist". This can occur if a file is deleted by somebody after Syncovery has built the file list, and before it starts to copy the file.

Chapter 6: Additional Topics

6.1 The SmartTracking Sync Operation Mode

SmartTracking is used when synchronizing two locations, and new or updated files can be expected on both sides. SmartTracking can detect on which side a file has been changed, moved, or deleted.

SmartTracking uses a local database in order to track changes that have been made between the various invocations of the profile. That way, the software will know whether a file has been moved on the left side, or on the right side. It can also detect whether a file has been deleted on one side, or whether the file has actually been added on the other. SmartTracking should always be used whenever you want to keep two locations in sync, both of which are being used for work. SmartTracking is not needed when you do a backup or mirror, or any other case where you synchronize in one direction only. SmartTracking has various additional features which you will see in the SmartTracking dialog box.

SmartTracking keeps information in a database linked to the profile name. On the first run, SmartTracking will record the state of each file. Starting with the second run, the SmartTracking settings will become active. Do not rename the profile, because then it will not be able to find its original database. Do not use more than one profile for synchronization of the same folders.

The SmartTracking dialog has four tab sheets:

Detection of Moved Files
"Moved files" refers to files that have been moved into a different folder on one side, but not on the other. By looking at these files only, it is not possible to determine on which side the file has been moved - in other words, it is hard to tell which of the folders the desired new location for the file is. SmartTracking solves this problem by comparing the current locations against the previous locations stored in its database. That way, the software can find out on which side the move has been performed, and it can then perform the same move on the other side as part of the synchronization process.

Files Deleted On One Side
In many cases, when the synchronization is started, there will be files that exist only on one side of the synchronization. These files are usually new files that should be copied over to the other side. But sometimes, this situation occurs because a file has been deleted on one side. The user will normally not want to have this file copied back from the other side. The desired action is often to have it deleted on the other side too.

SmartTracking can detect these cases, but ultimately you are responsible for what the program does, so please verify the actions in the Synchronization Preview before starting it. To ensure that no important files can get lost accidentally, the software will not actually delete these files. It will move them into a special folder which you need to specify as the last setting in this dialog box. That way, you can always retrieve files that have been deleted but for some reason shouldn't have.

Files Changed On Both Sides
Traditionally, when synchronizing bidirectionally, this synchronization software would always

copy the file with the latest timestamp to the other side, overwriting the older file that is already there. This is usually fine, but in some cases it may be a problem: what if the file has been edited by two different persons on both sides of the synchronization? One of the files may have a later timestamp, but it is not clear that this is the (only) version that should continue to exist on both sides.

So, you can choose to have such files be **labeled as CONFLICT**. Then, you will need to look closely at the Synchronization Preview window and resolve these conflicts manually by specifying a copying direction for each file. Or you can specify that both versions of the file, from the left and right sides, should be kept. To keep both versions, one of them is renamed and then both are copied.

Consider different files with the same name, not present in database, a conflict
This covers a very specific case, where a file exists on both sides, but it is not identical, and it's not in the database either. Normally the program would simply copy the newer version of this file and overwrite the older version on the other side with it. To prevent that, you can use this checkbox. It will make the program consider these files a conflict and treat the conflict according to the settings in the upper part of this tab sheet.

OPTIONS

Detect Unchanged Files
Some storage locations, such as FTP or WebDAV servers, always give uploaded files the current system time as Last Modified timestamp. This is generally a problem for synchronization, because you can never get the timestamps to match on both sides. If you are synchronizing in both directions, to and from the problematic server, then you may want to have SmartTracking detect whether files with new timestamps are really changed or not. It does so by querying the timestamp of freshly uploaded files immediately after the upload. This timestamp is stored in the SmartTracking database. On the next run, if the timestamp is still the same, the program will know that the file has not changed.

6.2 Real-Time Synchronization

Syncovery allows you to synchronize files in real time to ensure that your backup is always up to date. Whether for individuals or businesses, the method continuously monitors folders and copies new or edited files to the selected storage device immediately after they are changed. This eliminates the need to scan folders individually and significantly speeds up the synchronization process.

When using Real Time Synchronization, Syncovery will monitor folders for changes and copy new or modified files with a very short delay after the change has occurred. Because the folders are monitored, they don’t have to be scanned – the “Building File List” step does not occur. Real Time Synchronization can be very convenient and speeds up the synchronization process considerably. However, it will only copy new and modified files and it may not be appropriate in all situations.

Real Time Synchronization is chosen in each profile, on the right-most tab sheet among the Schedule settings. The scheduler must be started in a separate step to start the monitoring.

Because Real-Time Scheduling does not compare the source and destination folders completely, it is recommended to also schedule a complete profile run, for example once a day. This can be done on the “Schedule” tab sheet.

6.2.1 Prerequisites and Limitations for Real-Time Synchronization

This feature relies on real-time change notifications that it receives from the operating system. Therefore, it works only when the storage that you are copying from fully supports such change notifications. For example, local hard drives are usually fully supported. Most file systems also send change notifications over the LAN. But many NAS devices do not send change notifications properly. And sometimes even if the notifications are sent, the operating system may ignore them (such as Mac OS). In addition, change notifications are not sent by servers using an Internet Protocol such as FTP, SSH, WebDAV etc.

However, Syncovery can regularly check Internet based folders for changes. This works best with cloud storages that use a changes-based listing method: Google Drive, Sharepoint, OneDrive, DropBox, and Box. On other storages, Syncovery can check the base folder for changes regularly, and scan subfolders if they have new modification timestamps. However, most servers will not update folder modification timestamps correctly across the whole hierarchy, so that a regular full scan would still be necessary. To avoid this, Syncovery itself can ensure that parent folder timestamps are updated when it uploads to an Internet server. It will do so in real-time two-way profiles, or if the option “Touch parent folders” is chosen on the Folders tab sheet in the profile. This technique might not work with some types of servers (such as FTP, Amazon S3, Azure – but SFTP is fine).
The new behavior can be chosen on the Real-Time Settings dialog, on the new tab sheet “Internet Folders (Polling for Changes)”. The new features are also available as separate options, when you scroll down on the “Folders” and “File Age and Size” tab sheets.

When you see that real-time events are not received and the copying is not triggered as expected, you need to switch to regular synchronization, such as every 5 minutes.

6.2.2 Combining Real-Time Monitoring With a Time Window

You can also combine real-time synchronization with a time window.
Syncovery can monitor the folders for changes all day long but do the copying only within the specified time window. To take advantage of this functionality the following settings need to be combined:

- Schedule: no schedule on the Schedule tab sheet
- Real Time Sync selected on the Monitoring/Realtime tab sheet
- Real Time Settings:
    - the checkbox “Run profile entirely once” must be deselected
    - “Process Complete Folders” must be chosen (on Windows only)
- Weekdays and Time Window:
  Specify the time window when it should do the copying.

6.3 Block Level Copying

Syncovery can detect which parts of a larger file have changed, and copy only the changed blocks rather than the complete file. This feature works similar to Rsync (but it's not the same).

Copying only the changed blocks can save bandwidth and time, especially over a slow connection. When copying between local disks or in a LAN environment, it can save bandwidth too, but may not always save much copying time, because the source file has to be read in its entirety every time in order to determine the changed blocks (except when using the File System Monitoring Service).

Only block-oriented file types are eligible for block-level copying. These include database files such as SQL or Outlook PST, as well as drive images and virtual hard disk images (VMs). Stream-based files, on the other hand, will usually cause all blocks to be changed whenever they are modified (for example text documents, spreadsheets, zip files, and photos). Thus block-level copying won't be able to save much bandwidth with stream-based files.

In Syncovery, block-level copying is sometimes also called "Partial File Updating". In many cases, you need to choose only the checkmark "Block Level Copying", which is on the Special settings category in the profiles (in Advanced Mode).

The program needs to have fast access to at least one of the sides of the synchronization (except for Mode 0). The other side may be a low-bandwidth connection. If you are using an Internet Protocol, please note that only SSH/SFTP supports block-level updating directly. The other protocols can only be used with Synthetic Backup (see further down below).

Please note: Block level copying with SSH/SFTP has only been implemented for uploads, not downloads.

Block Level Copying can work in four different ways:

Mode 0: With File System Monitoring Service (new in Syncovery 8):
– the Syncovery File System Monitoring Service must be running and monitoring the source folders
– if the source is on a local hard disk, the File System Monitoring service is automatically configured when you save the profile.
– extremely fast, no full scanning of the source file necessary and no MD5 checksums needed
– Destination files must not be modified by any other profile, person, or tool
– Destination must be accessed via LAN, VPN, or SSH/SFTP
– if the source is a UNC path, a special set-up is necessary (see at end of page)
– cannot be combined with Mode 3 (Synthetic Backup) at this time

Mode 1: With Database:
– Source access must be fast
– Destination may be slow
– MD5 checksums are stored in database
– Destination files must not be modified by any other profile, person, or tool
– Destination must be accessed via LAN, VPN, or SSH/SFTP

Mode 2: With Remote Service:
– Syncovery Remote Service computes MD5 checksums on remote computer
– The "slow" side can be both source or destination

– MD5 checksums are newly calculated each time
– Files on both sides can be modified by other profiles, persons, or tools
– One side must be local or LAN/VPN, the other can be LAN, VPN, or SSH/SFTP

Mode 3: Synthetic Backup:
– similar to Mode 1, plus:
– adds Zip compression, versioning, and filename encoding
– can work locally or with any Internet Protocol for the destination side
– the changed blocks are uploaded in a new, separate zip file every time
– all older zip files must stay on the backup storage, but can be thinned out
– all connection types and Internet Protocols are supported

Mode 1: Slow Destination
In this mode, the speed-up is available when you copy files from a location to which you have fast access (preferably your own hard disk). The destination can be a slow connection, but it must be a normal file system (either LAN or VPN) or SSH/SFTP. For other connections, you can use Synthetic Backup.

Instructions for Mode 1
In your profile, make the following checkmark: **Block Level Copying**, which is on the Special tab sheet when editing the profile in Advanced Mode. The next time you run the profile, a database is created on your hard drive where information needed for the speed-up is stored. The second time you run the profile, you should notice the speed-up.

Mode 2: With Remote Service
This mode can speed up updating large files in both directions. The remote computer can be both source and/or destination. This is achieved by running a small service application on the remote computer, which will create the necessary checksums on the fly, when requested by the main application running on a different machine. See more information on the Syncovery Remote Service here.

The other (local) computer, where the main Syncovery program is running, needs to have normal file system access to the remote computer (LAN or VPN), or it can use SSH/SFTP. It needs to have write access to the remote computer so that it can save the checksum request file there. The MD5 checksums are created when needed, so that no database is being used.

Instructions for Mode 2
On the remote system, run the Setup program and install the Syncovery Remote Service along with its control panel. Start the control panel from the Syncovery group in the Start menu. On the tab sheet Configure Checksummer, enter the base folders that will be used for synchronization. Click Apply. On the tab sheet Service Configuration, click on Install Service and Start. The service will be using the Windows System account by default. If this account doesn't have sufficient access privileges, you may have to change the account in Windows Control Panel -> Administrative Tools -> Services.

On the local system, you are running the main Syncovery program. In your profile, the right-hand side must be the remote system. Specify one of the folders which you have specified for the remote service to monitor. The left side should be your local folders, or a network drive with

relatively fast access. On the Specials tab sheet in Advanced Mode, make the following checkmarks: Block Level Copying and Right side uses Remote Service.

Mode 3: Synthetic Backup
This feature is intended for backing up from a local storage to any type of backup storage. The backup can be local or online.

Choose “Synthetic Backup” on the tab sheet Versioning->Synthetic Backup. This will automatically place these additional checkmarks:
– Block Level Copying (under Special)
– Filename Encoding (under Versioning)
– Compress Each File Individually (under Compression/Encryption)

Find out more on Synthetic Backup on this page.

Setting up the Mode 0 (File System Monitoring Service) with a UNC path

The file system monitoring service can be used to enable block level copying from a UNC path, in addition to copying from local drives.
To make this work, the monitoring service must be installed and configured on all machines that write to the UNC path. The Auxiliary Services Control Panel must be used to configure the File System Monitoring Service. You need to type the UNC path that is the source path for the Syncovery job into the configuration field.
In addition, the File System Monitoring Service must be given a user account that has access to the UNC path. This is done via Windows Control Panel->Admin Tools->Services by editing its properties. Then you can start the service. You should see a hidden .Syncovery folder in the UNC path, and a file such as SyncoveryFSMonitor.MACHINENAME.active in it. If you have problems setting this up, please contact support@syncovery.com.

6.4 The Syncovery Command Line

In addition to the user-friendly GUI, Syncovery is also a flexible command-line sync tool. Syncovery's command line interface offers a powerful, flexible way for users to automate and control their file synchronization and backup operations directly from the terminal, scripts, or other programs. Whether you're looking to automate your backup routines, integrate Syncovery operations into scripts, or manage your tasks in a more hands-on manner, Syncovery's command line provides the versatility and control needed to implement your exact requirements.

You can run, create, edit, rename and delete profiles via command line.

Windows
You can invoke Syncovery.exe, SyncoveryCL.exe and SyncoveryService.exe with the command line parameters shown on this page.

Command Line Examples for Running a Job on Windows:

- Run a job with the command line tool SyncoveryCL:
  SyncoveryCL.exe /RUNX="Profile Name"
- Run a job in the GUI but minimized:
  Syncovery.exe /RUNX="Profile Name" /S /M
- Run a job invisible but with progress reporting to the GUI:
  SyncoveryService.exe /RUNX="Profile Name" /S /M /P
- You can use multiple /RUNX= parameters to run several jobs.

Useful command line parameters to control how a job is run:
/P = with progress reporting to the GUI
/T = run multiple jobs in threads; the process will also accept additional jobs from the scheduler
/S = silent (not relevant for SyncoveryCL)
/M = minimized (not relevant for SyncoveryCL)

Macintosh
On Mac, you can use the command line parameters in a shell script like in this example:
open ./Syncovery.app --args /RUNX="Profil Name" /U

In addition, you can use the new SyncoveryCL command line tool. It is contained in the Contents/MacOS folder inside the Syncovery.app bundle. You could use it there directly, or copy it to a more convenient location. It takes the same parameters that you see on this page. For additional SyncoveryCL command line examples, please see the Linux page.

Parameters for Running a Job

- **/RUN="Profile Name" and /RUNX="Profile Name"**
  With these options, you can invoke a specific profile from the command line for immediate execution. /RUN= starts the profile and leaves the program running. /RUNX= starts the profile and exits the program after the profile has been executed.
  You can use the asterisk wildcard at the end of this option, for example:
  /RUN=* or /RUNX=* or /RUNX=Office*
  You can also add this to cause a shutdown when done by adding /SHUTDOWN

- **/O** (show only results)
  Runs the job without Sync Preview but shows a dialog with the results.
- **/U** (unattended mode)
  The profile is started and executed without further user interaction. When this option is not specified, the file list is being shown, but the actual synchronization must be started manually. In unattended mode, the program will only show an error message if one of the synchronization paths cannot be accessed.
- **/S** (silent mode)
  Like /U but there will not be any error messages whatsoever.
- **/M** (minimized mode)
  The program is minimized to the system tray while executing the profile. This option implies /U.
- **/P** (progress reporting – SyncoveryService.exe and SyncoveryCL.exe only)
  SyncoveryService.exe will report progress information to the main program, in the same way the scheduler starts background jobs.
- **/T** (run in threads – SyncoveryService.exe and SyncoveryCL.exe only)
  SyncoveryService.exe will run the job in a separate thread rather than in the main process thread, in the same way the scheduler starts background jobs when the option "Start profiles in parallel" is chosen.
- **/LEFT="Folder Path"**
  allows you to specify a left path that is different from the one stored with the profile.
- **/RIGHT="Folder Path"**
  allows you to override the profile's right path.
- **/MASK="File1.txt;*.doc"**
  allows you to override the profile's inclusion masks, or simply specify one or more files to copy.
- **/EXCL="*.bak"**
  allows you to override the profile's exclusion masks.
- **/SHUTDOWN**
  Shutdown after running the profile(s) specified with /RUNX="Profile Name".
- **/NOLOG**
  Do not generate a log file for this run.
- **/CHOOSESUBFOLDERS**
  The subfolder selection dialog is shown prior to running the job. The selection that is made is not saved permanently unless you also specify the /SAVE parameter. Alternatively, you can add /SAVEONLY in order to save the new selection without running the job. In all cases, the profile name should be specified with /RUNX, even if you use /SAVEONLY.

Managing Jobs

- **CHANGE "Profile Name" /Disabled**
  Disables a job. If it's currently running, the run will complete normally.
- **CHANGE "Profile Name" /Disabled=0**
  Enables a job.
- **/PAUSEJOB="Profile Name"**
  Pauses a running job. This only applies to jobs which are currently scanning folders or copying/deleting files. These activities can be paused. A paused job must later be resumed or canceled.
- **/RESUMEJOB="Profile Name"**
  Resumes a running job.
- **/STOPJOB="Profile Name"**
  Stops / cancels a running job.

General-Purpose Parameters

- **/INI="Path to configuration file"**
  Specifies the configuration file to be used. For example, /INI="C:\ProgramData\Syncovery\Syncovery.ini"
- **/IMPORT="Path to XML or CSV file"**
  Import profiles from XML or CSV. To find out how to make such files, please create a sample profile and export it by right-clicking it on the Profile Overview.
- **/EXPORTONEXML="Profile Name"**
  Export a profile in XML format. The xml file will be written to the current directory.
- **/EXPORTALLXML="XMLFileName.xml"**
  Export all profiles in XML format. The xml file name can be a complete path with file name.
- **/EXPORTWP**
  When added to the previous export parameters, the profile will be excluded with passwords. Passwords are encrypted with a portable, hardcoded key and can be imported on any other machine that runs Syncovery.
- **/EXPORTPW="Password Phrase"**
  When added to the previous export parameters, the given password phrase is used to encrypt the exported passwords, so that a leaked export file cannot be imported without knowing the password.

Generating Sync Preview Only

- **/PREVIEWONLY**
  Perform a dry run, generating the Sync Preview and then stopping. By default, the preview will use a simple text output format to the console.
- **/PREVIEWFILE="path\to\preview.xml"**
  Output the preview to a file. Recognized filename extensions are xml, html, csv, tsv, ssv. Other extensions will output a simple text format.

- **/PREVIEWFORMAT=XML|HTML|CSV|TSV|SSV**
  Specify the preview format, if it can't be guessed from the filename. CSV=comma separated values, TSV=tab separated values, SSV=semicolon separated values

Example: output the preview as a CSV file with semicolon as separator.
SyncoveryCL /RUNX="Profile Name" /NOLOG /PREVIEWONLY
/PREVIEWFILE=C:\Tests\Preview.CSV /PREVIEWFORMAT=SSV

Command Line Parameters for Managing Profiles

The command line to create any given profile can be seen by creating it manually first, then going to Information->Show Profile Details... in the profile editor.
Here's an example:
Syncovery.exe ADD /Name="Documents Backup" /Left="C:\Users\Me\Documents"
/Right="D:\Backup\Docs" /L2R /Deletes /ReplaceNewer /MaxParallelCopiers=2

Alternatives to ADD

RUN (creates & runs a job without creating a permanent profile)
DELETE "Profile Name" (deletes a profile)
CHANGE "Profile Name" (changes an existing profile, keeping any existing settings which are not set on the command line)
RENAME "Old Profile Name" "New Profile Name"

ADD can be used to completely replace an existing profile with new settings.

6.5 Variables to be used in Profile Paths

These variables can be used in the sync paths, in addition to standard Environment variables such as %TEMP%.

$PROGDRIVE	=	drive specification where the software is installed
$DATE	=	current date
$TIME	=	current time
$WEEKDAY	=	day of week
$2WEEKSDAYS	=	like $WEEKDAY but for two weeks, i.e. A Monday ... A Sunday, then B Monday ... B Sunday, then starts over
$DWMQ	=	rotation with separate folders for each day of the week, week, month, quarter (see example). Keeps one copy for each quarter, then one copy for each month in the last quarter, one copy for each week in the last month, and one copy for each day in the last week.

$MYDOCUMENTS	=	current user's My Documents folder

$ALLUSERSPROFILE	=	All Users folder, which is usually C:\ProgramData
$USERPROFILE	=	Current user's profile folder
$APPDATA	=	Current user's Application Data folder

$DAYOFMONTH or **$MONTHDAY**	=	today's day number in the month (01-31)
$DAYOFMONTHSD	=	today's day number single digit (1-31)
$WEEKDAYNUM	=	today's day number in the week (1=Monday)
$WEEKOFYEAR	=	The week number (European).
$YEARFORWEEKOFYEAR	=	The year that matches the week number (European).
$USWEEKOFYEAR	=	The week number (U.S.).
$YEARFORUSWEEKOFYEAR	=	The year that matches the week number (U.S.).
$MONTH	=	this month's name
$MONTHNUM	=	this month's number (01-12)
$MONTHNUMSD	=	single digit month (1-12)
$YEAR	=	current four-digit year number
$HOUR	=	Hour the profile was started (two-digit)
$MINUTE	=	Minute the profile was started (two-digit)
$SECOND	=	Second the profile was started (two-digit)
$HHMMSS	=	Time the profile was started

$YESTERDAY	=	yesterday (01-31)
$YESTERDAYSD	=	yesterday single digit (1-31)
$YDWEEKDAY	=	yesterday's name
$YDMONTHNUM	=	yesterday's month (01-12)
$YDMONTHNUMSD	=	yesterday's month single digit (1-12)
$YDYEAR	=	yesterday's year

$TOMORROW	=	tomorrow (01-31)
$TOMORROWSD	=	tomorrow single digit (1-31)

$TOMWEEKDAY	= tomorrow's name
$TOMMONTHNUM	= tomorrow's month (01-12)
$TOMMONTHNUMSD	= tomorrow's month single digit (1-12)
$TOMYEAR	= tomorrow's year

$LEFTPATH	= Adds the left-hand path (for use in right path only)
$RIGHTPATH	= Adds the right-hand path (for use in left path only)
$PROFILENAME	= The name of the profile.

Variables that can be used in the Execute before/after paths

$LEFT	= left base path of synchronization
$RIGHT	= right base path of synchronization
$LEFTDRIVE	= Drive letter or UNC file share of the left-hand side
$RIGHTDRIVE	= Drive letter or UNC file share of the right-hand side
$LEFTMACHINE, **$RIGHTMACHINE**	*Windows only (since v8.27c):* = first part of UNC path with machine name but no share name
$MSG	= Short result summary
$RESULT	= GOOD or BAD (if an error has occurred)
$PROFILE	= Profile name
$LOGFILE	= Path and name of log file

6.6 The Syncovery Remote Service

The Remote Service is a useful tool to enable some special features which are somewhat less frequently used. There are two kinds of services, the Syncovery Service, which is the scheduler and runs profiles in the background, and the Syncovery Remote Service, which is not needed to run any profiles.

The Syncovery Remote Service is installed on the "other end". For example, if you use Syncovery on a client which exchanges files with a server, then Syncovery Remote Service can be installed on the server to perform one or more of the following tasks:

- Unzip incoming Zip or Sz Packages. This feature is used in conjunction with the Compressed Packages capability in order to speed up the transfer of files.
- Generate a Remote Listing. This feature is used to speed up the "Building File List" phase.
- Generate MD5 Checksums for Block Level Copying and for file verification (hash comparison).
- Includes an optional SFTP Server that can be used for Peer-to-Peer file synchronization.

Installation of the Syncovery Remote Service

On Windows, the Remote Service is installed using the same Setup program from our download page that you use to install Syncovery. When running the Setup, make sure you include the Syncovery Remote Service and the Auxiliary Services Control Panel. On other platforms, the Remote Service is a separate download.

You will find the Auxiliary Services Control Panel in the Syncovery program group in your Windows Start menu. Please run it and install the service with it.

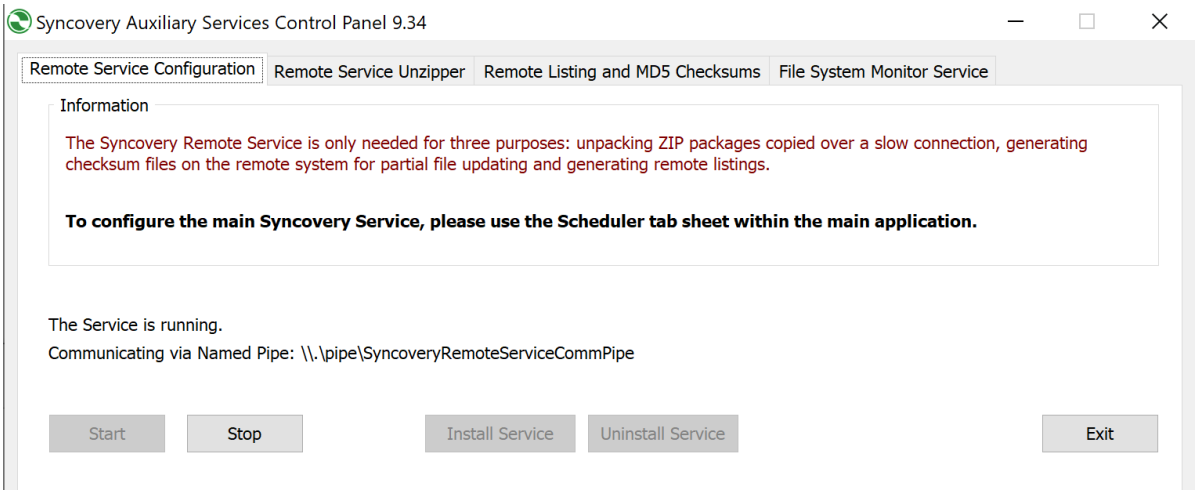

Configure Unzipper

To use the unzipping capability, you need to specify the "folders to watch" on the tab sheet "Remote Service Unzipper". Each folder must be the same one that you specify as destination folder in the Syncovery profile on the other computer. Except that here, you specify local folders, but they must be the same ones. If you have several profiles with different destination folders, each one must be specified here on a separate line, even if they are all on the same drive.

The Syncovery profile on the sending maching needs to be set up to “Use Compressed Packages” and on the Compressed Package Configuration dialog, you need to use the setting “Syncovery Remote Service will be used”.

The Syncovery main program will put the compressed packages into a subfolder called “Incoming Syncovery Packages”. This is actually the folder that the Remote Service’s unzipper will watch, but you don’t specify it. You specify the same folder that is used in the profile.

Configure Checksummer and Remote Lister

The checksummer and the remote lister are configured on the third tab sheet. Like the unzipper, they watch certain folders for incoming files. The incoming files are special request files sent by the synchronizer on the other machine. So, the Synchronizer does not communicate with the Remote Service using any kind of network protocol. Instead, it puts a “request file” with a special filename into the folder on the server, and the Remote Service sees the file, reads it and then knows what to do and deletes the request file.

You need to specify the “folders to watch” for incoming checksum or remote listing request files. Each folder must be the same one that you specify in the Syncovery profile on the other computer. Except that here, you specify local folders, but they must be the same ones. If you have several profiles accessing different folders on this machine, each one must be specified here on a separate line, even if they are all on the same drive.

Using the Remote Listing Feature

Once the Remote Service has been set up correctly, you can speed up the “Building File List” phase considerably by using the checkmark “Left or Right side listing uses Remote Service”, which can be found in the profiles on the “Specials” tab sheet, when editing them in Advanced Mode. In most cases, only one of the two remote listing checkmarks are used (only the one for the non-local side).

Using the Checksummer Feature for Block Level Copying

You will find more information on Block Level Copying on this page.

6.7 Copying Security / Permissions and Shares

Syncovery can copy NTFS security settings along with files.

On the tab sheet Special in the profiles, you will find a checkmark called “Process Security and Shares”. On the dialog that pops up, you need at least three checkmarks on the tab sheet “File and Folder Security”: Copy Owner/Group/Permissions.

If the source and destination base folders do not have the same permissions yet, you need to also choose “Process Base Folder Too”. This feature will make sure that the destination base folder has the same permissions as the source folder, so that all subfolders and files can inherit from that. If necessary, Syncovery will break the inheritance on the destination base folder so it can make the permissions match. It will not break inheritance on any subfolders, unless their inheritance is also broken on the source side.

In addition, the program needs to be running “As Administrator”. To ensure it runs as admin, you need to right-click its icon and choose
“Run As Administrator”, unless you have Windows versions older than Vista, or you have disabled UAC.

If using the scheduler, it probably needs to run as a service. Please see
Installing the scheduler as a service

Chapter 7: How-to Guides

This chapter contains a few details guides from our web site.

7.1 Windows File Server Migration Guide

Copying File Shares from One Windows Server to Another Using Syncovery

Migrating file shares from one Windows server to another involves careful planning and execution to ensure data integrity, minimize downtime, and preserve NTFS security settings. This guide will walk you through the steps of file server migration using Syncovery, including copying files, preserving NTFS permissions, creating the file shares on the new server, and planning and implementing a possible transition phase as well as the final cutover.

File Server Migration Prerequisites

Be sure that the following requirements are met to be able to perform the server migration:

- Access to both source and destination Windows servers with administrative privileges
- Syncovery installed on either the source or destination server (installing on a third machine is possible but less efficient and not recommended)
- Adequate storage space on the destination server.
- Knowledge of the current shares and NTFS permissions setup.
- Servers should be in the same Windows domain (if not, special steps may be necessary)
- You can start off with our free demo version of Syncovery and order a Syncovery Premium Edition license if you decide to use it.
- Determine if there will be a transition period where files need to be updated on the destination server. Or, in a more difficult case, there could be a period where your co-workers or clients might be working on both the old and the new server. Although this should be avoided, it is possible to set Syncovery up to handle such a case.
- Be aware of your timing requirements and the final cutover for your file server migration.

Key Decisions Before Starting the Migration

Before the migration can start, you should make the following decisions:

Installing Syncovery on the source or destination server

You can install Syncovery on either the source or destination server, or even on a third machine. For efficiency reasons, to increase speed and use less network bandwidth, it is recommended to install Syncovery on either the old or the new server, rather than in independent third machine.

Advantages of installing Syncovery on the source server:

- Initial folder scanning will be extremely fast
- Real-Time mirroring during a transition period may be more efficient and reliable
- Folder scanning performance for repeated synchronizations can be increased using the option “Cache Destination File List”

Advantages of installing Syncovery on the destination server:

- Setting file metadata and NTFS security settings on the destination folders can be faster
- The new server usually has more CPU power and memory (RAM), allowing Syncovery to run faster

Installing Syncovery on an independent third server, such as a VM, would make it necessary to connect to both source and destination over the network (LAN). You would not benefit from any of the aforementioned advantages. Network traffic would be doubled, as all data needs to be transferred over the network to the machine that Syncovery runs on, and then again be copied to the destination machine. While not using any TEMP space on the Syncovery machine, it is still considered inefficient and not recommended (though entirely possible).

Copying Individual File Shares or Entire Drives

If you have multiple file shares on a server hard drive, you can decide between copying the whole drive (or selected folders on the whole drive), or copying individual file shares. Copying the whole drive has the advantage of needing fewer profiles, but the potential disadvantage is that you need to use a file share that exposes the whole drive (such as an administrative file share like D$).

If you have many file shares, Syncovery can create them on the destination server for you. The setting to create file shares is on the same “Process Security and Shares” dialog where you choose to copy NTFS permissions.

Here’s an example for the profile’s base paths when copying a whole drive:
\\SOURCESERVER\$D → D:
or
D:\ → \\DESTSERVER\D$

Copying file shares separately may be necessary if the drive layout on the new server will be different from the old one. When copying file shares separately, it could look like this:

\\SOURCESERVER\Data → D:\Data
\\SOURCESERVER\Archive → D:\Archive
\\SOURCESERVER\Documentation → D:\Documentation

or
D:\Data → \\DESTSERVER\Data
etc.

A third possibility is to use the base paths shown below. Specifying simply “\\SOURCESERVER” on the left-hand side will allow you to select the shares to copy using the subfolder selection dialog in Syncovery:
\\SOURCESERVER → D:

Creating and Running the Syncovery Profiles

After determining which source and destination paths you want to use, you can proceed with the profile creation in Syncovery. You may need to create one or several profiles. Make sure you are using Syncovery in Advanced Mode and choose the following settings for a file server migration. Not many of the default settings need to be changed.

Sync Operation Mode: Exact Mirror

Files → Number of Files to Copy In Parallel: The default of 3 copying threads should be increased if you have many smaller files. You can easily increase it to 10 or 20 if you have a fast network connection between the servers and aren't mostly copying huge files (such as video). The number you choose also depends on how much network and server load you can use and how fast you have to complete the file server migration.

Job → File List Threads: The number of File List Threads (= folder scanning threads) can be set to between 10 and 30 if you have fast drives, servers, and a fast and stable network connection. If you need to be careful with CPU usage and general server load, and your folder and file count isn't extreme, you can keep the default setting.

Special → Process Security and Shares: If NTFS Security Settings (aka Permissions) need to be copied, please choose the security copying settings according to this guide.

Selective Synchronization: Subfolders, Masks, and Filters

By default, Syncovery copies all folders and files except for the Global Exclusion Masks. You will find these on the Program Settings dialog, tab sheet "Types, Limits". The Global Exclusion Masks are pre-filled with some default masks when Syncovery runs for the first time. For example, the files named Thumbs.db are excluded by default. Please take a look at the masks and decide if you want to keep them. You can also set individual profiles to ignore the Global Exclusion Masks. This is done under "Masks & Filters"->"Exclusions" in each profile.

If you need to do a selective synchronization, you can use the subfolder selection dialog as well as masks and filters. For example, you could exclude folders named Cache;Caches;Temp by simply typing this into the Exclusion Masks. You can use filters to copy only files from a specific date range, or copy only files within a range of file sizes.

Folder and File Security Settings (aka Permissions)

In the previous step, you chose whether to copy folder and file security settings. These include the **Owner, Group, and Access Control Lists (ACLs)** with separate permissions for additional users or groups. If the two servers are in the same domain, copying the permissions should be straightforward. If they are not, and you do want to copy NTFS permissions, please make sure that the all users and groups exist on the destination server and choose "Translate Security IDs" on the Advanced tab sheet of the **Security and Shares** dialog box in the Syncovery profile. If the user names on the new server are not identical to the old one, Syncovery can translate user names. Please contact support for details on how to set up this rarely used feature.

Note that in addition to folder and file security, the network shares have ACLs, too. The permissions of a folder and its share are combined according to the **Least Privilege Principle**: The effective permission for a user is the most restrictive of the combined NTFS and share permissions. This means that if either the NTFS permission or the share permission is more restrictive, that will be the effective permission.

Data Verification

When copying between two servers in a stable and modern wired LAN, you can rely on the copies matching the source files exactly and don't need to worry about data corruption. However, if you want Syncovery to verify each file after copying it, you can choose **Files → Verify copied files**. This will read the whole file back after copying it and compare it against the source byte-by-byte. Hashes are not used by default, but can be if you install the Syncovery Remote Service on the other server(s).

As an alternative to verifying each file after copying it, Syncovery can also compare the contents of all files that already exist on both sides. This is done during the folder scanning phase. To activate this verification, choose **Comparison → More → Binary Comparison of Existing Files While Building Files List**. Because this will slow down the folder scan, you will want to run it only once (or occasionally), and then remove the checkmark again. You can also use the checkmark "Remember Results", causing Syncovery to remember which files it has already verified and skip future verifications for those.

Copying With Temporary Filenames

Syncovery can optionally use temporary filenames on the destination side while a file is being copied. When all bytes have been transferred, the interim filename is renamed to the actual original filename. The special naming of interim filenames allows Syncovery to resume broken transfers even if the profile is stopped and later started again. It also ensures that broken transfers can never be mistaken for valid files. In general, this feature is not needed for a file server migration, if the two servers are connected with a stable LAN. Nonetheless, it is something to be aware of and to consider. The option is on the "Files" tab sheet in the profile and the checkmark is called "Automatically resume (copy with temporary filenames, keep incomplete files when stopped)".

Running the Migration Job

Note that Syncovery runs large jobs in parts by default, splitting the profile run after reaching 2 to 5 million files during the scanning phase. If your server has ample RAM, you can increase the splitting limit or remove it. Please look at the Program Settings dialog, tab sheet "Memory" before starting the profile.

Next, you can run the initial copying task for your file server migration. Make sure you are logged on as a domain Administrator and open Syncovery by (shift-)right-clicking its icon and choosing "Run As Administrator". You can start the first profile execution in Attended Mode if you want to see the Sync Preview before starting the copying. If this is not necessary, you can run the job in Unattended Mode or in the background by right-clicking the profile and choosing the desired mode of operation.

Continuous Sync During Transition Phase

If you have the requirement to update the new server continuously during a transition phase, you can use Syncovery's scheduler to run the job regularly. You can schedule regular syncs to occur multiple times per day, or once during the night. You can also choose Real-Time Synchronization to avoid complete folder scans and update only the folders with modifications in near real-time. Make sure to install the scheduler as a service so that it runs with Admin rights, and also runs when you log off. The scheduler must be set up and started on the "Scheduler" tab sheet.

To allow Syncovery to make the network connection to UNC paths automatically, it may be necessary to specify the network credentials in the profile via **Job → Network Connections...**. This is not necessary if Syncovery runs under a domain account.

To speed up the folder scanning for regular syncs, please install the Syncovery Remote Service on the other server so that it can build the file list quickly. Also see our page about speeding up folder scanning, which includes additional tips.

Ideally, nobody would be working on the new server yet, during the transition phase. This will ensure a smooth and reliable file server migration. For special cases, where both the old and the new server have to be used in parallel, please contact Syncovery support for additional instructions.

Synchronizing Deletions During Transition Phase

If there is a transition phase where files need to be updated on the destination server, you also need to decide if files should be deleted from the destination if they are deleted (or moved/renamed) from the source side. **Exact Mirror** mode was recommended above, which will process deletions. If this is not desired, you would choose **Standard Copying** instead. Please note: if the job is run in unattended mode or via scheduler or in real-time, you need to allow unattended deletions according to this documentation page.

Preparing For the Final Cutover

Even if you regularly keep the destination server updated, you should perform one final full comparison / sync on the cutover day. Ideally, this would be done during a few hours where nobody works with the files, either on the old or on the new server. Remember to increase the scanning threads and use the Remote Service to speed up the comparison phase. If you don't have a sufficient period of scheduled downtime, you can switch to Standard Copying mode to prevent Syncovery from deleting any files from the destination, and also avoid newer versions of files being replaced with older ones from the old server. Run the final sync in Attended Mode and check the Sync Preview to see if any remaining copying actions need to be done. You can remove files from the Sync Preview that don't need to be copied or deleted.

If you previously chose "Cache Destination File List", you should turn it off for the final sync. The destination cache would only be beneficial if Syncovery is installed on the source system or another machine (not the destination itself). But it should be off for the final folder comparison.

Additional Thoughts and Questions

If you are going to merge the contents of several file servers into one, you may need to use some different settings. For example, you may not be able to use Exact Mirror mode, because that mode would possibly delete files which have been copied from one of the other servers.
A PascalScript is available that would rename duplicate folders, if necessary.

Sometimes you'll want to copy the smallest files first, or the most recent ones first. Different copying orders can be chosen on the "Special" tab sheet in the Syncovery profile.

If there are any remaining questions concerning your file server migration, feel free to write to support@syncovery.com or visit our support forum.

7.2 Working With Sharepoint, Microsoft 365 and OneDrive

This chapter focuses on the Sharepoint-specific topics, and especially on how to connect Syncovery to your Sharepoint site or to your OneDrive. First, let's clarify some of Microsoft's products and APIs.

7.2.1 OneDrive versus OneDrive for Business

OneDrive and OneDrive for Business are both cloud storage services provided by Microsoft, catering to different audiences with distinct features.

OneDrive (Personal)

- **Target Audience:** Individual users for personal storage.
- **Use Case:** Ideal for storing personal documents, photos, and files accessible from any device.
- **Account Association:** Linked to a Microsoft account used for personal services like Outlook.com, Xbox Live, or Skype.

OneDrive for Business

- **Target Audience:** Businesses and organizations, part of the Office 365 or SharePoint Server subscription.
- **Use Case:** Designed for storing, sharing, and collaborating on work documents within and across organizational boundaries.
- **Account Association:** Linked to an Office 365 or Microsoft 365 business account, managed by the organization's IT department.

How to Know Which One You Have

1. **Account Type:** If you log in with a personal email address (e.g., @outlook.com, @hotmail.com), it's likely OneDrive personal. For OneDrive for Business, you'll use your work or school email address.
2. **Web Interface:** When logged in, the URL can give a hint. Personal OneDrive usually has a URL pattern like onedrive.live.com, whereas OneDrive for Business will be accessed through the Office 365 portal or a direct URL that includes the organization's name (e.g., yourcompany.sharepoint.com).

7.2.2 Connecting Syncovery to your OneDrive

To authorize Syncovery to sync with your OneDrive, please click on the Internet button for one side in your Syncovery profile. Change the protocol from FTP to "OneDrvNew" for personal OneDrives, or "OneDrive for Business" for the business OneDrive. Then click the Browse button and you will be asked to authorize Syncovery in your web browser. When that is done, you can choose the folder to synchronize with.

7.2.3 Connecting Syncovery to Sharepoint Sites

For Sharepoint Sites, there are additional steps needed to achieve the connection, due to the bigger number of choices you have with Sharepoint. A corporate site frequently has subsites, groups, and multiple document libraries within a single site.

Please note that you cannot just copy a Sharepoint URL from your web browser into Syncovery. Instead, you need to pick your document library from a list that Syncovery will show you.

7.2.4 Choosing the Protocol for a Sharepoint Site

Microsoft allows accessing Sharepoint sites via two distinct APIs. In most cases, the Graph protocol is recommended. However, in some organizations, you may get an error message saying that this action requires Admin consent. In that case, you can try the Sharepoint protocol instead of Graph, or contact your Microsoft 365 administrator so they can allow Syncovery to be used in your organization.

1. **Graph**: this is the newer protocol from Microsoft, which features greater flexibility and may be able to show you a more complete list of all the sites, subsites, groups, and document libraries/drives in your organization. To connect Syncovery via the Microsoft Graph, please choose “Graph” from the protocol list on the Internet Protocol Setting dialog, and make sure that the field Domain/Site is empty. Then click the upper one of the two Browse buttons, and authorize Syncovery via your web browser. Next, Syncovery will enumerate all the sites and document libraries that you have access to. Choose the library you need, and then proceed setting up the Syncovery profile.
2. **Sharepoint**: the original Sharepoint protocol also still works very well. Authorization differs slightly, in that you need to enter your Sharepoint domain into the field “Domain/Site”, and then click the Browse button next to it. Enter only the domain without any https:// prefix and without any slashes at all. For example: contoso.sharepoint.com. The Sharepoint protocol can be used whenever there are authorization issues with Graph.

Choosing the Document Library

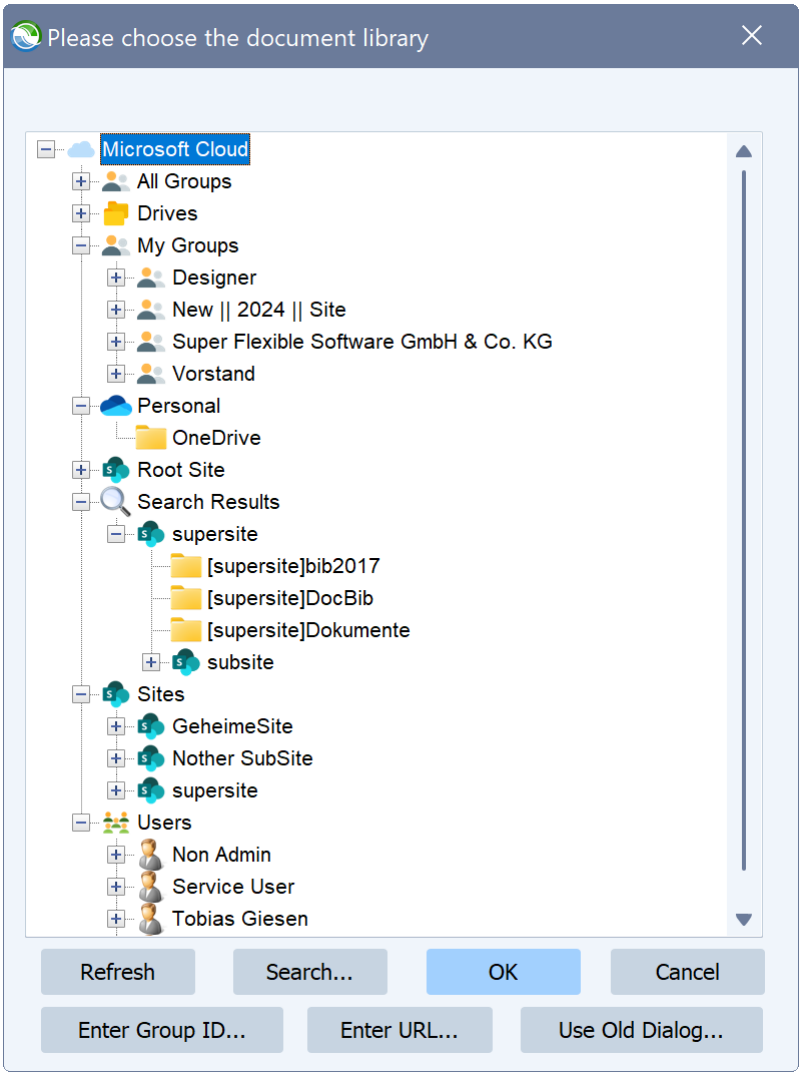

Sharepoint Document Library Selection Dialog Window

When you click on the upper of the two Browse button on the Internet Protocol Settings dialog, the document library selection dialog will appear. Since Syncovery 10.13.0, the dialog looks like this.

You can expand the various nodes in the tree to find your library. Before clicking the OK button, you need to select an item that has a yellow folder icon. These items are the actual document libraries (or drives) where folders and files are stored. You cannot actually select a folder within the library on this dialog. The dialog's only purpose is to navigate to the document library and select it.

The quickest way is often to use the Search button and type the group or site name, which will then appear in the tree under "Search Results".

Sometimes you may find it hard to find the site you need in the document library selection dialog, if the user account that you are connecting with has limited access. You can temporarily authorize with an Admin account just to browse the sites. To do this, please right-click in the selection dialog and choose "Switch to Admin Account...". Syncovery will make a separate

authorization using the app name “Permission Granter for Syncovery”, which you can revoke in the Azure Portal when no longer needed.

After choosing the document library, you can use the second Browse button to choose a folder within the library.

7.2.5 Downloading or Uploading with Shared Folders

Since version 10.15.6, the library selection dialog includes an item named **Shared** which allows you to work with shared folders that appear in your OneDrive for Business under “Shared”. This feature is only available when the Graph protocol is chosen. Syncovery can upload to and download from shared folders, depending on the level of access that was granted to you. You can synchronize all your shared folders or use the second Browse button on the Internet Protocol Settings dialog to choose a specific shared folder. If additional folders are shared with you, they will be added automatically. You can also use the subfolder selection dialog to choose specific folders.

7.2.6 Need Admin Approval?

Some corporate sites won’t allow you to authorize third party apps unless you are an M365 Administrator. If you get an error like “Need Admin Approval” or “Admin Consent Required”, you need to contact the M365 Administrator and ask them to allow either “User consent for apps” or “Admin consent requests”.

Allowing **User Consent For Apps** can be done from the following URL: https://portal.azure.com/#view/Microsoft_AAD_IAM/ConsentPoliciesMenuBlade/~/UserSettings

Here’s a **screenshot** of how to allow users connect third parts apps:

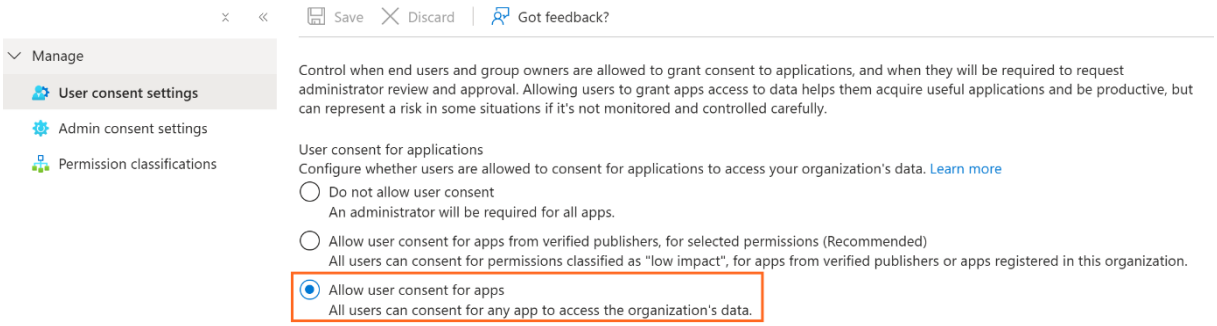

Allow user consent to apps.

If you can’t just broadly allow user consent for apps, you can instead allow users to submit Admin Consent Requests. Allowing **Admin Consent Requests** does not reduce the security of the corporate M365 site. An administrator must visit the following page in the Azure Portal: https://portal.azure.com/#view/Microsoft_AAD_IAM/ConsentPoliciesMenuBlade/~/AdminConsentSettings

Here’s a **screenshot** of the required setting:

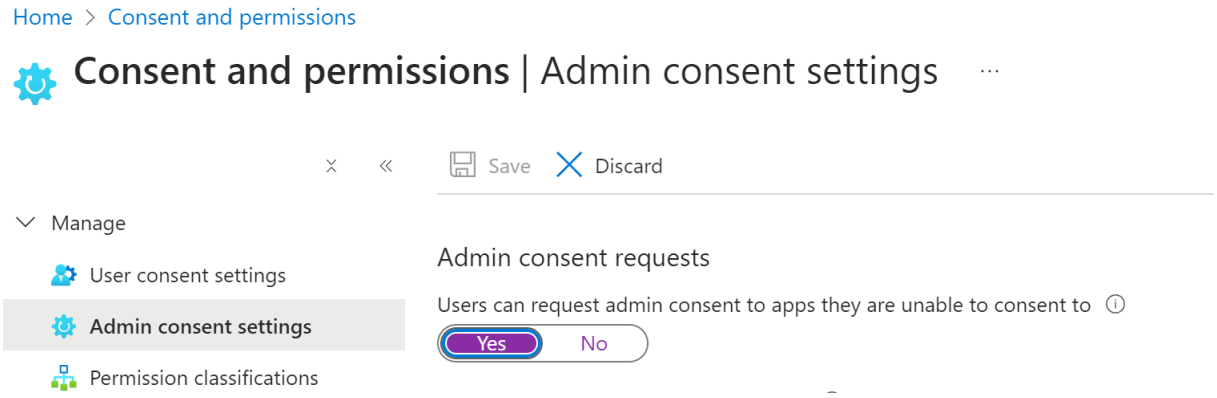

Once Admin Consent Requests have been granted, you can try authorizing Syncovery again, using the Graph protocol. This time, you will be able to send a consent request. After you have sent the request, the authorization process in Syncovery will fail once more. But now, the admin can visit the following page in the Azure portal to grant your request:

https://portal.azure.com/#view/Microsoft_AAD_IAM/StartboardApplicationsMenuBlade/~/AccessRequests

Once the admin has granted your request, you can use Syncovery with the corporate M365 site.

7.2.7 Limiting Access to Specific Sites

To access Sharepoint sites, Syncovery needs to request the permission scope "Sites.ReadWrite.All". Some security departments may find this scary and won't allow it. But what they need to know is that "Sites.ReadWrite.All" is actually limited by the user account that is used to authorize Syncovery. Syncovery will **not** get access to all sites – it will only get access to the sites that the user has access to.

If you need to limit Syncovery's access to only a few sites, please create a new Sharepoint Online user that can access only the necessary sites. The user doesn't normally need any Office licenses.

When authorizing Syncovery with this user, Syncovery's access rights will be very limited.

7.3 Automatic PGP File Exchange (Encryption/Decryption)

Need to exchange files via PGP encryption and SFTP or other Internet Protocols? Syncovery is the ideal tool for such requirements.

If you regularly need to receive a PGP encrypted file, download it and decrypt it, you can set up a simple Syncovery profile to perform the task. It can do the opposite as well: encrypt a file with PGP and send it using Internet protocols like SFTP and others.

As a reminder, you need a Private Key to decrypt PGP files, and a Public Key to encrypt them.

Here's how to set up the Syncovery profile. On the left side, you will usually specify your local folder, or a UNC path that corresponds to a folder in your network. On the right side, you would set up the Internet protocol. Click the Internet button and change the protocol from FTP to the one you need (such a SSH/SFTP), and specify the credentials.

For a profile that PGP-encrypts files and sends them via SFTP, the top of the Syncovery profile dialog will look similar to this:

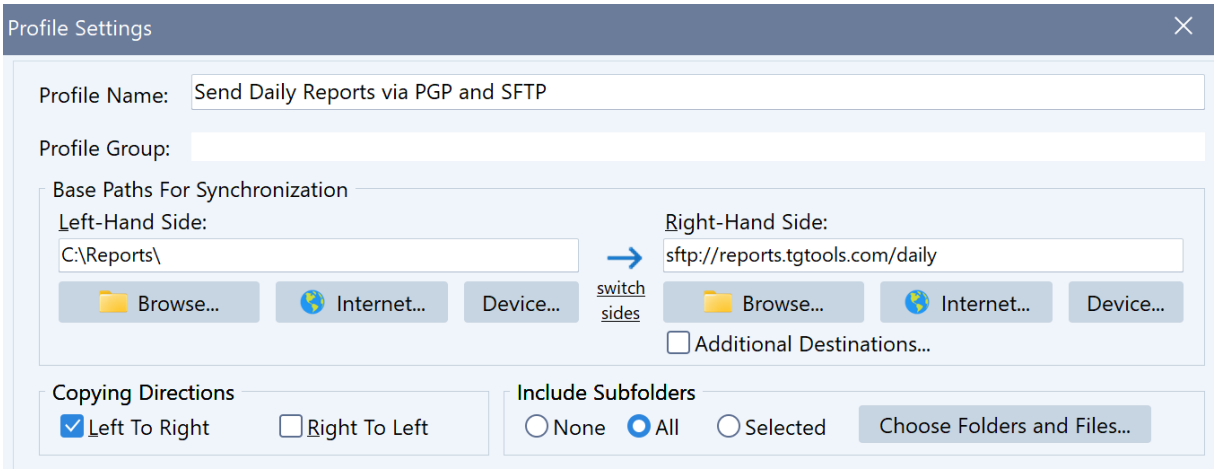

Encryption with PGP

The PGP encryption is chosen in the Advanced Settings category "Compress/Encrypt", on the tab sheet "Encryption":

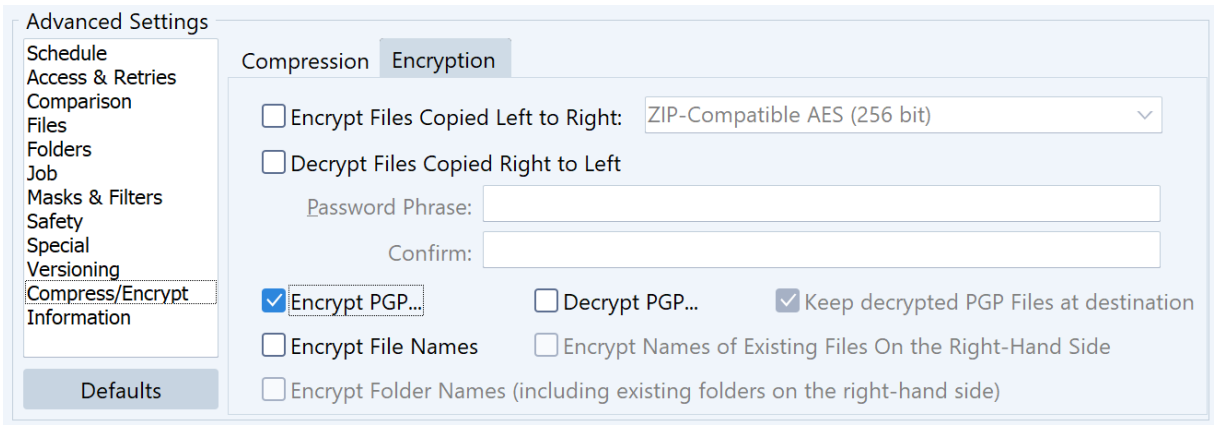

When you click on the checkbox "Encrypt PGP...", a dialog box will appear where you are asked to paste the Public Key that should be used for encryption. You should have received the public key from the person or company you will be sending files to. It will look similar to this:

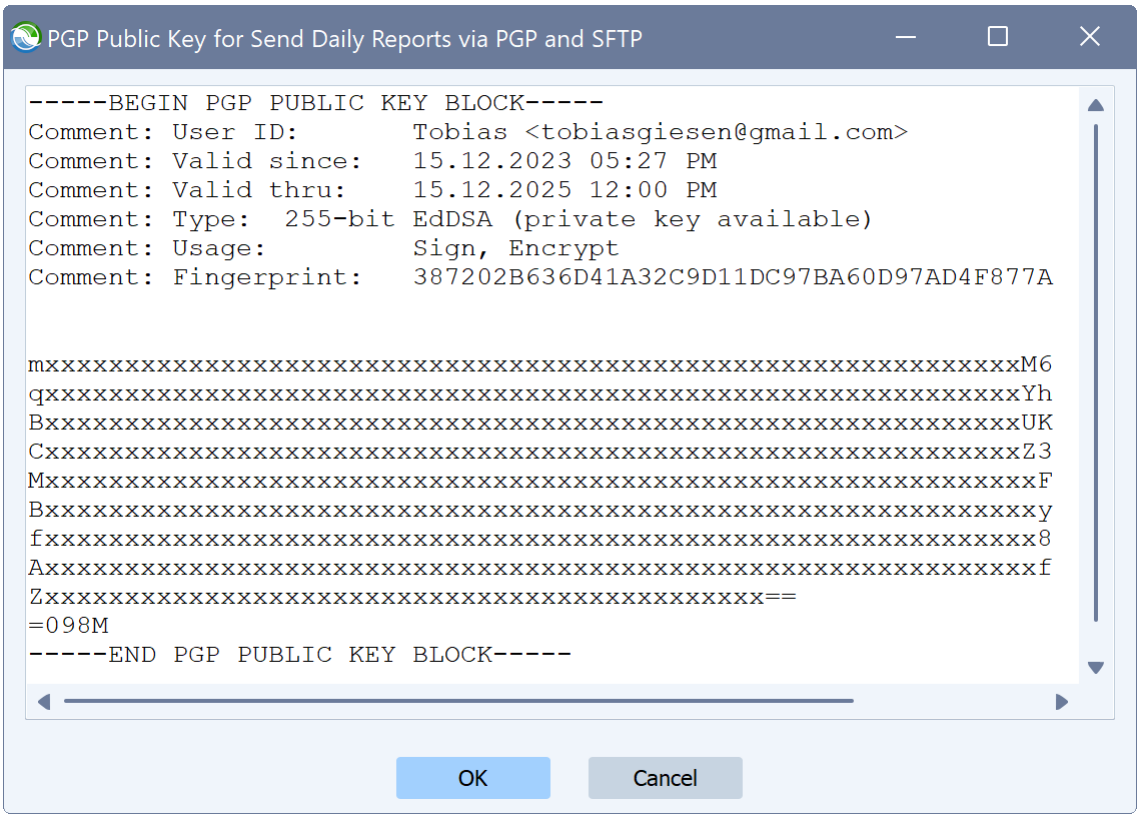

Decryption with PGP

Conversely, if you are going to receive PGP encrypted files, you would choose “Decrypt PGP...” and paste your Private Key into the dialog box. If your Private Key is password-protected, please enter the password into the fields “Password Phrase:” and “Confirm:”. You may need to choose the checkmark “Decrypt Files Copied Right to Left” for the password field to be enabled. However, PGP decryption is not limited to right-to-left sync jobs.

When decrypting PGP files, you can choose if you want to keep the PGP file itself too, or just the decrypted file. The main reason for keeping the PGP file, too, is to prevent Syncovery from re-downloading the same file again the next time when the job is executed.

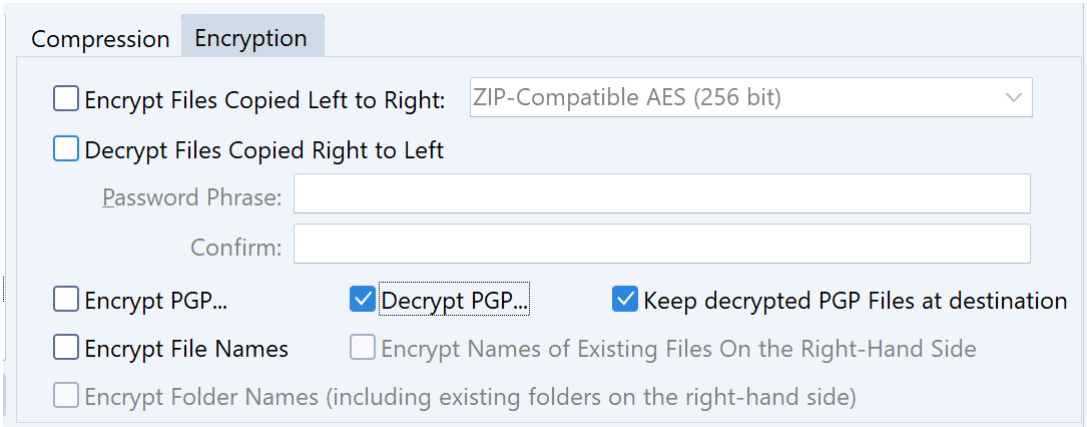

It may not be necessary to keep the encrypted PGP file, for example:

- if you want to download the same file every time
- if you download once per day, and there’s a different file on the server every day

- if you use a date/time filter such as File Age: less than 24 hours old
- if you use SmartTracking to remember which files have already been downloaded
- or, if you copy from a local or UNC path rather than an Internet Protocol, you can use the two Archive Flag filter checkboxes on the General Filters tab sheet to mark files as processed and avoid re-copying them

If the decrypted file is a zip file, you can also let Syncovery unpack it. The checkbox for unzipping can be found at the bottom of the tab sheet "Compression". Again you have the choice to keep the zip file after unpacking. In this case there is no reason to keep it really:

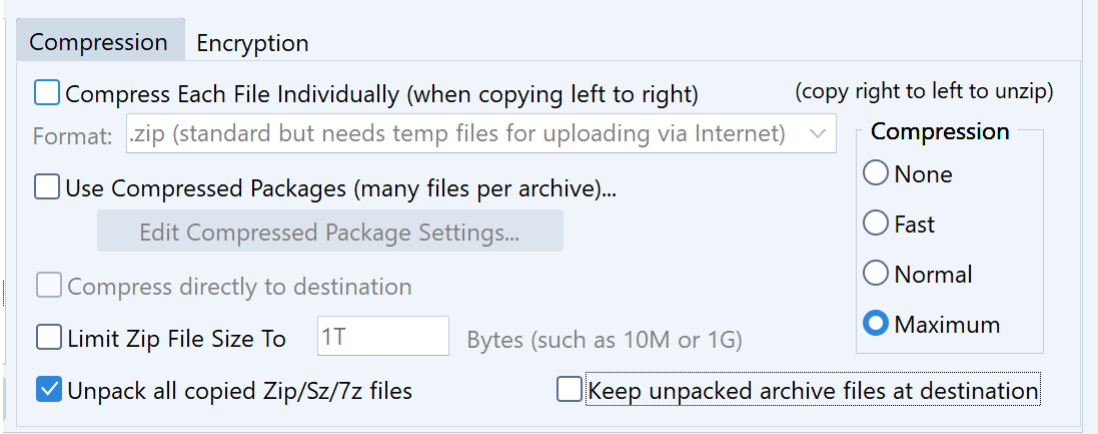

7.4 Speeding up Building the File List

Before proceeding to copy files, Syncovery must build the file list. It normally scans all subfolders and builds a complete folder listing, rather than starting to copy right away. One of the reasons for this behavior is to be able to detect moved files.

There are various ways to speed up building the file listing.

- Use recent Syncovery versions, which can generate a fast, multi-threaded file listing. If necessary, you can increase the number of folder scanning threads on the Performance tab sheet on the Program Settings dialog, or on the Job tab sheet in the profile.
- Make sure you haven't chosen "Binary Comparison Of Existing Files" on the *Verification* settings category in the profile.
- Make sure you don't use any unncessary logging. Especially "With Timing Info" or "File List Building Details" or "Internet Protocol Logging" dramatically slow down the listing process.
- If you are using "Process Security and Shares" on the Special tab sheet, consider whether you need the checkboxes "Update existing files" because that will be slow. If you must update existing items, make sure the setting is "Update Existing Items: Folders (and files will inherit)".
- If you are using FTP, you can choose the recursive FTP Listing Command LIST -alR on the second tab sheet of the Internet dialog. If you are using SSH/SFTP and the server is a Linux/Unix type of server, you can try the "Recursive Listing" checkmark.

- You can install the Syncovery Remote Service on the other computer to generate the file list remotely. The Remote Service is available for Windows, Linux, FreeBSD and Mac.
- If you have a one-way sync and you are just mirroring or backing up, you may be able to use the option “Cache Destination File List” from the Special tab sheet. However this option means that any changes on the destination by another person or program are not seen by the software, because it always remembers the last state in its cache and never looks at the destination folders again.
- You can use Real Time Synchronization which simply copies new and changed files rather than comparing the two folder structures. However, it is recommended to also schedule a full run regularly to catch any files that may have been missed in real time. Also, Real Time changes are only detected on local drives and via LAN or VPN. Changes are not detected via FTP, WebDAV and so forth (except see the new polling feature in Syncovery 9). Some computers or network devices may not be sending real-time notifications over the LAN.
- If you only need to copy new and modified files to the destination and never delete any files from the destination, you could turn off scanning the destination completely (on the Files tab sheet), and under General Filters, use the two Archive Flag checkmarks (or on macOS/Linux: “use extended attributes to mark files as copied”). On the first run, this will mean copying all files because the Archive flags are still set for all files. You can avoid that by adding a fixed date/time filter such as Date later than XX/YY/ZZZZ.
- If the Archive Flags cannot be used, you could still turn off scanning the destination and use a File Age filter such as “less than 1 days” old.

7.5 Fixing a “cannot access left / right path” error

- This usually affects network paths. Please make sure you use UNC paths rather than drive letters (such as \\Servername\Sharename\Foldername)
- Test your UNC path in Windows Explorer and let it remember the network credentials (if any).
- If the scheduler is running as a service, it does not have network access without extra steps. There are two requirements for network access:
- First, if running as a service, it must be given a log on account. This is done when clicking on the **Install...** button on the Scheduler tab sheet to install the service. To change the log on account, please uninstall and then re-install the service.
- Test the account that you assign to the service, by logging in to Windows desktop as a person with the same account. Test that Windows Explorer has access to the volumes that you need. Make sure Windows Explorer remembers the credentials for any network paths.
- Again, remember that network volumes must be specified using a UNC path such as \\servername\sharename\foldername rather than a mapped drive letter.
- If this is not sufficient, you can provide a username and password for the network resource in each profile. Use this setting on the Job tab sheet in the profile: **Network Connections...** However, in many cases this is not needed. Rather than specifying the full path for the network connection, you can also try specifying just \\servername.

- If problems remain, please compare a working and a non-working log file. Take a look at the top, where you can see “Running as” which shows the user account that the jobs run under. They might be different. Log in to the Windows desktop as a person, with the same account that is not working. Test access in Windows Explorer.
- To change the user account for the service, please stop and uninstall the service and then click “Install” again to specify a different account.
- For more details, please see Running the scheduler as a service.

7.6 Peer to Peer File Transfer With Syncovery 11

Introduction to Peer to Peer File Transfer

Peer-to-peer file transfer is a powerful solution for sharing and syncing data directly between two devices without relying on a cloud service. In this guide, we'll walk you through setting up an SFTP server on one end, configuring your firewall and router to allow seamless communication, and using Syncovery to manage the file transfers. Whether you're looking to sync files, create backups, or simply transfer data securely, this step-by-step guide will help you configure everything from start to finish, ensuring smooth and secure file transfers between your systems.

First Steps

The steps outlined in this guide assume you will be transferring data, or performing a folder comparison or synchronization between two Windows PCs over the Internet. If you are synchronizing within a LAN or VPN, it will be much simpler, since you don't need to configure your router. But most of the steps shown below can be useful for synchronization with a LAN too – particular if you don't want to use traditional Windows network shares (SMB).

The first step is to decide which PC will host the SFTP server and which will run the main Syncovery program. The SFTP server runs in the background and has only a minimal GUI, while Syncovery on the other PC will be used to set up profiles/jobs to perform folder comparisons, backups and synchronizations.

7.6.1 Installing and Configuring the SFTP Server

Download the latest version of Syncovery 11 from your Download Page. Run the Setup program on the PC that will act as the server. Be sure to include the two setup choices shown in the screenshot below:

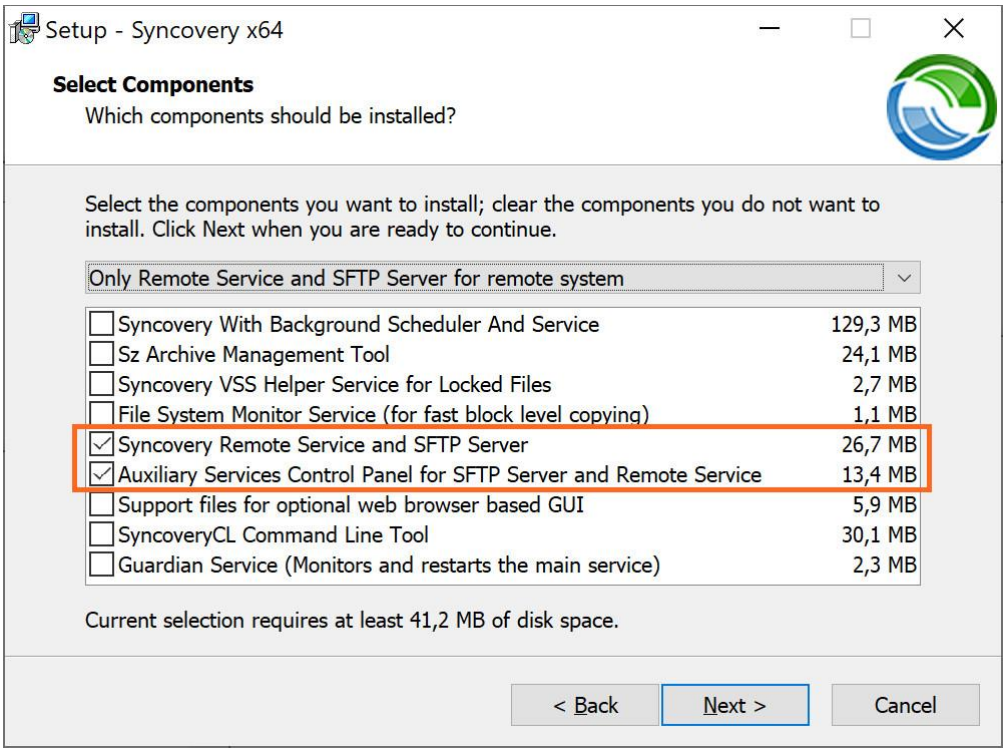

When the software installation is complete, start the Syncovery **Auxiliary Services Control Panel** by typing "Aux" in the Windows Start Menu search field:

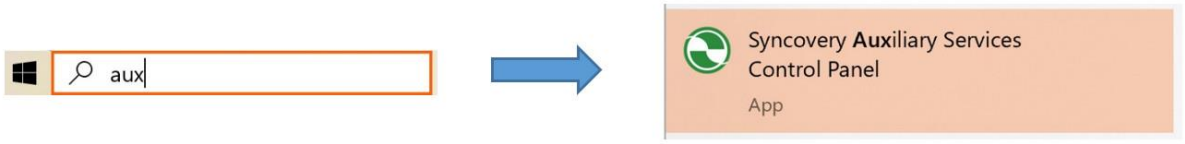

Next, we will configure the **Syncovery Remote Service**, which helps us generate folder listings fast. Just type the paths you want to work with on the third tab sheet as shown below:

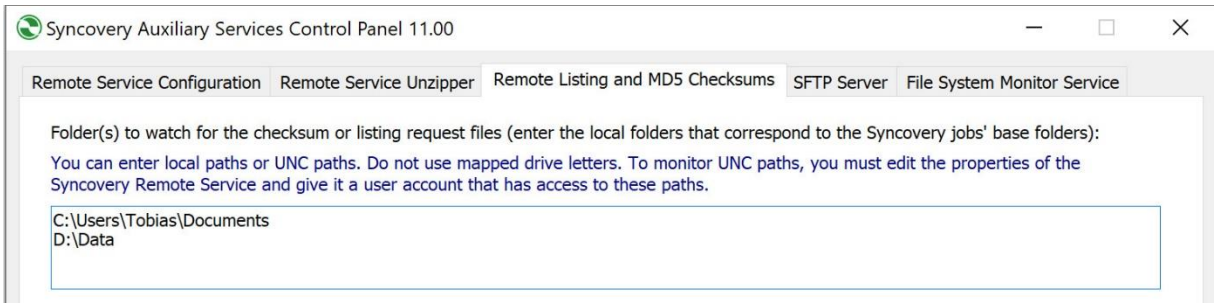

Now it's time to set up the **SFTP Server** feature. Choose the checkmark **Activate Syncovery SFTP Server** and click "Add" to add users:

You need to set up **at least one user account**. The account must be given a **strong password**, including numbers, letters, and special characters. It should be at least **twelve characters long**. Remember that your server can be reached over the Internet. Some other people might be going to attempt to log in, too, and we really need to make sure they can't get in. If you don't need the permanent availability for peer to peer transfer, you may even want to stop the SFTP server when not needed.

If you have folders on different drives, you need a separate user for each drive. In addition, you should not use C:\ as the SFTP base folder for a user, so that can also be a reason you need more than one user. Here's two examples:

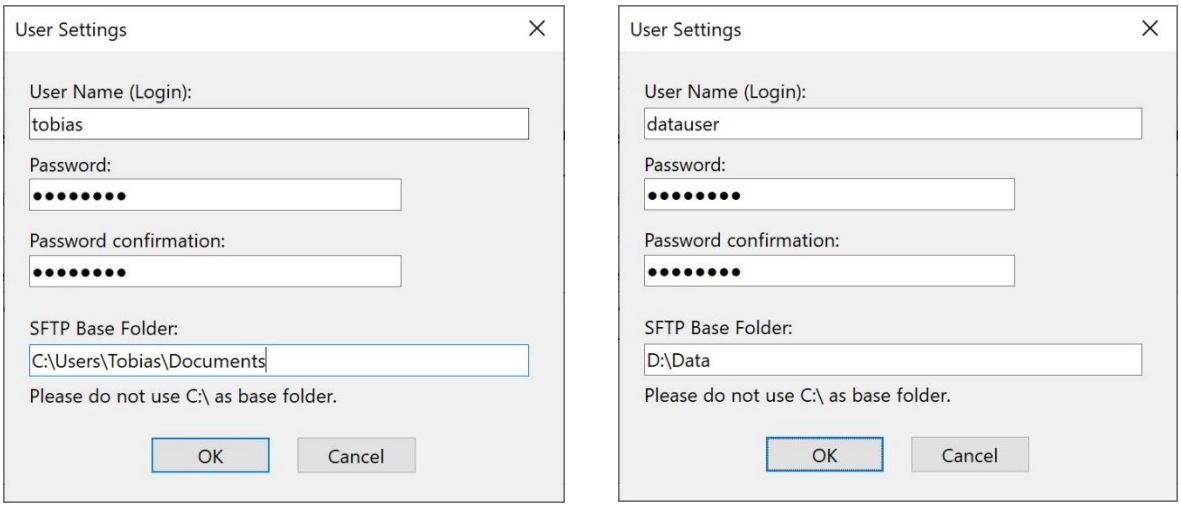

7.6.2 Finalizing the SFTP Server Setup

The SFTP Server tab sheet will now look like the following screenshot. Please choose the IP address of the LAN adapter that connects your PC to the Internet. Note the port number, which can be changed. The standard SFTP port is actually 22, but we do not want to use that for security reasons. Click the Apply button, and then "Run As App To Test Firewall". This button will do a test run of the SFTP server in a command prompt window, hopefully triggering a prompt from the Windows Firewall, which will allow you to let the SFTP server to be accessed from outside.

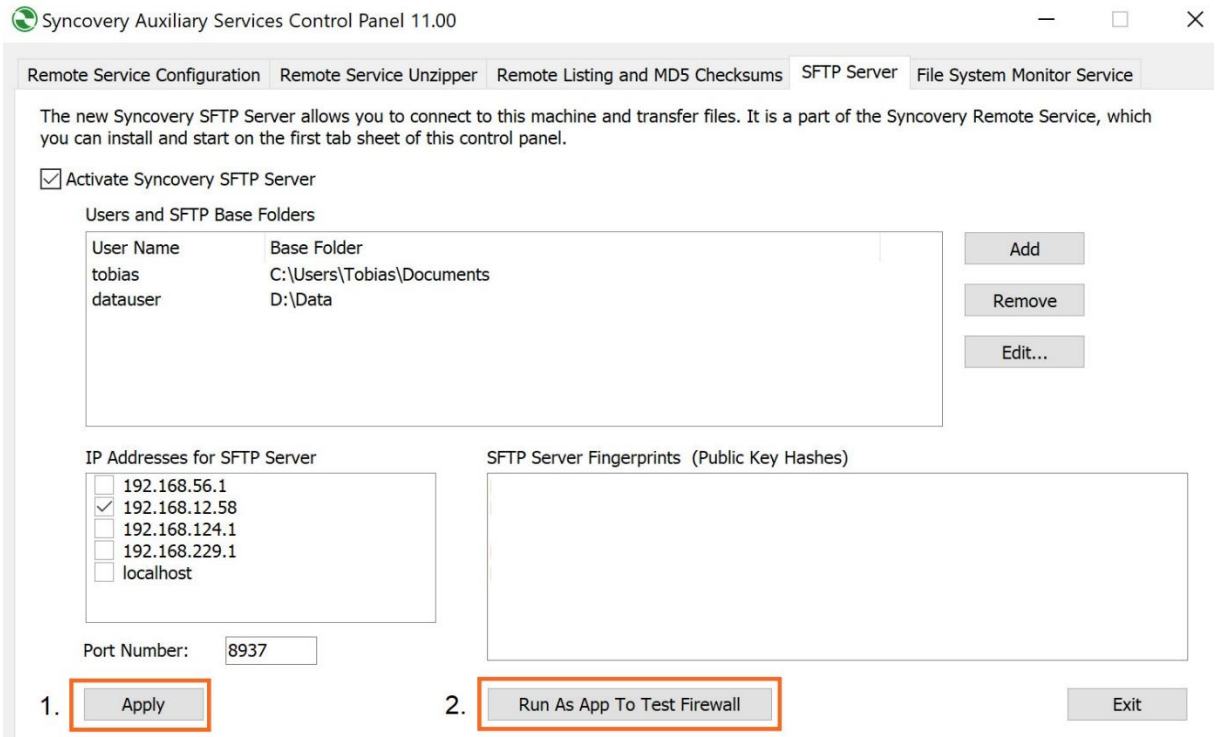

If you get a **Firewall prompt** like the following, please confirm it. If not, we may have to set up a Firewall rule manually. In either case, please close the command prompt window that you see.

Next, you can **install and start the Syncovery Remote Service**, which includes the SFTP Server. Make sure that the testing command prompt window is closed.

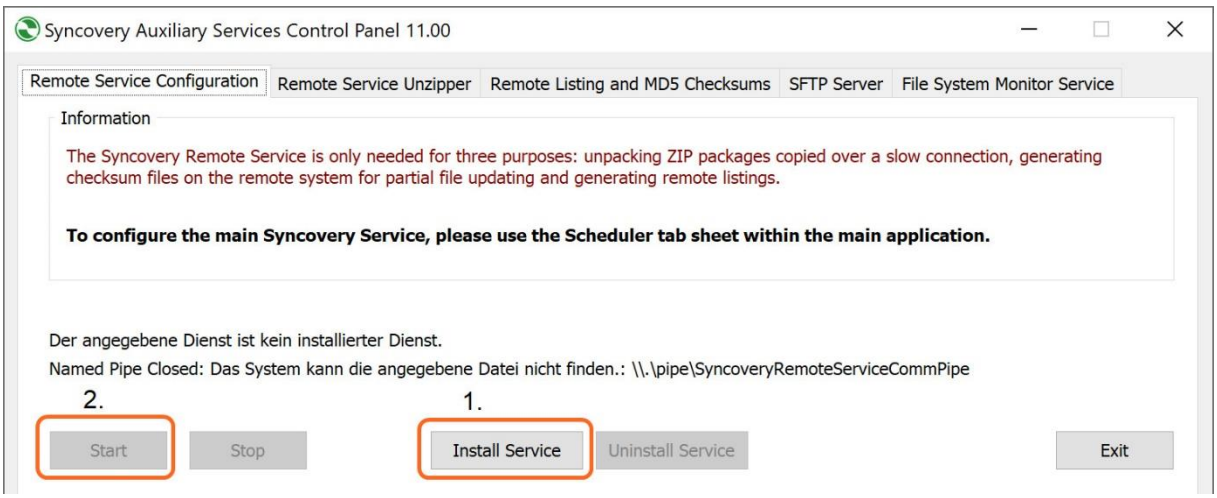

7.6.3 Setting up a Firewall Rule

If you didn't get a Firewall prompt, or if you later find out that the Firewall is still blocking the SFTP Server port, you can set up a rule by opening the **Windows Firewall settings** according to the following screenshots. Type "Fire" into the Windows Start Menu search field to get to the Firewall settings:

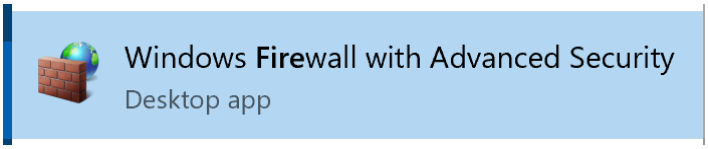

Click on **Inbound Rules** to the left and on **New Rule...** under Actions on the right side of the window:

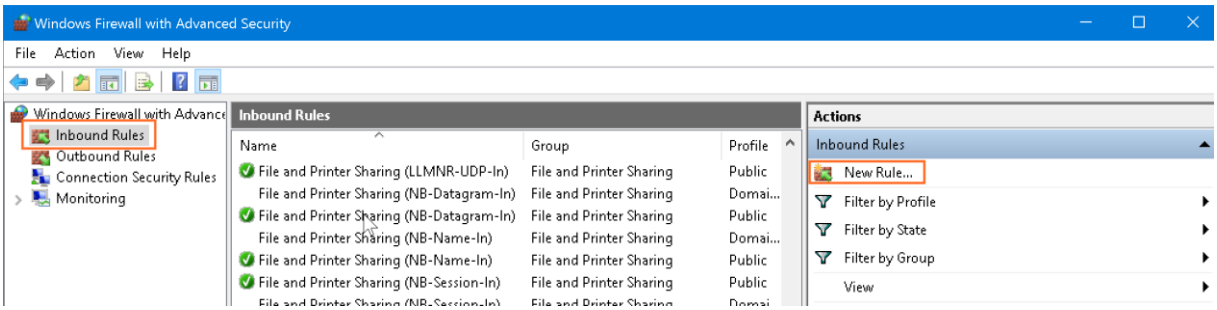

A wizard will ask you some questions about the rule. First, make sure you create a rule for a **Port**:

Second, specify the **port number**. If you kept Syncovery's default SFTP port, it's 8937:

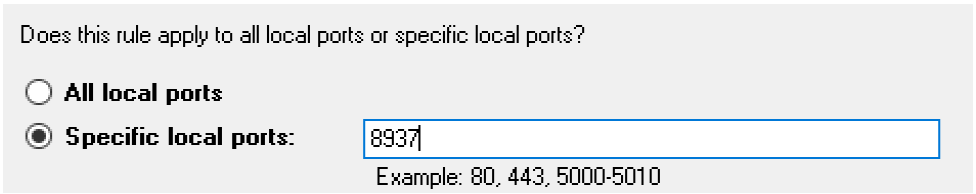

Next, please choose **Allow the connection**:

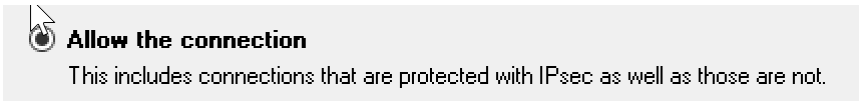

You can safely keep all the next choices chosen, as shown in the following screenshot:

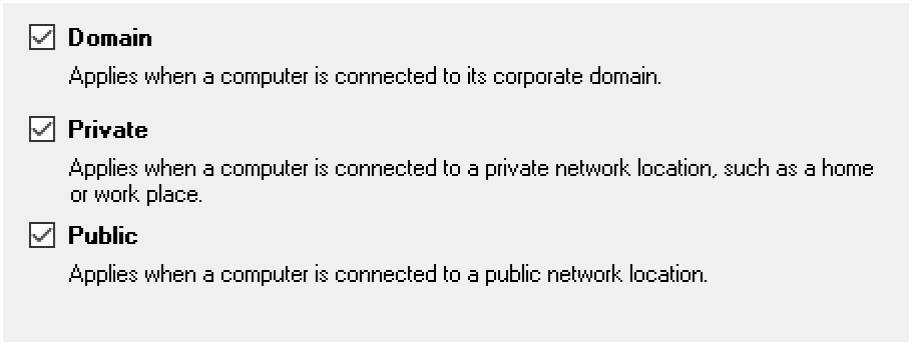

Finally, give the new **Firewall rule a name**, for example:

7.6.4 Setting up Port Forwarding in your Internet Router

To make sure that incoming requests from Syncovery on the other PC can reach the SFTP Server, we need to configure your Internet Router to forward the connection attempts to the PC where the SFTP Server is running. The relevant configuration pages in the router might be called

“Internet -> Permit Access”, “Port Sharing”, or “Port Forwarding”. You may need to select the PC from a list of connected devices, or type its IP address in the LAN, and then specify the port you want to use. Only do this if you specified a **sufficiently complex password** for all SFTP users, as documented above!

An example **port forwarding screenshot** is shown below:

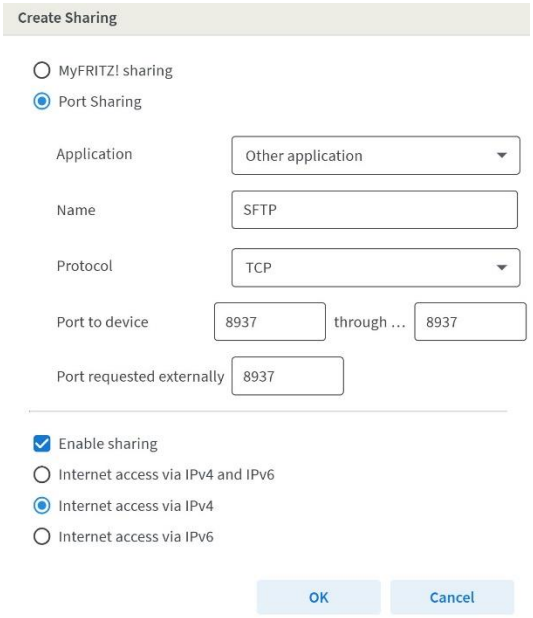

7.6.5 Creating the Syncovery Profile and Verifying the SFTP Server Fingerprint

It’s now time to take one final look at the **Auxiliary Services Control Panel**. If you click on the **SFTP Server** tab sheet once more, you will now see the **SFTP Server Fingerprints**. By comparing these fingerprints when Syncovery makes the connection, you can ensure that you are connecting to the correct SFTP Server and that there is no man in the middle attack. The dialog box with fingerprints looks as follows:

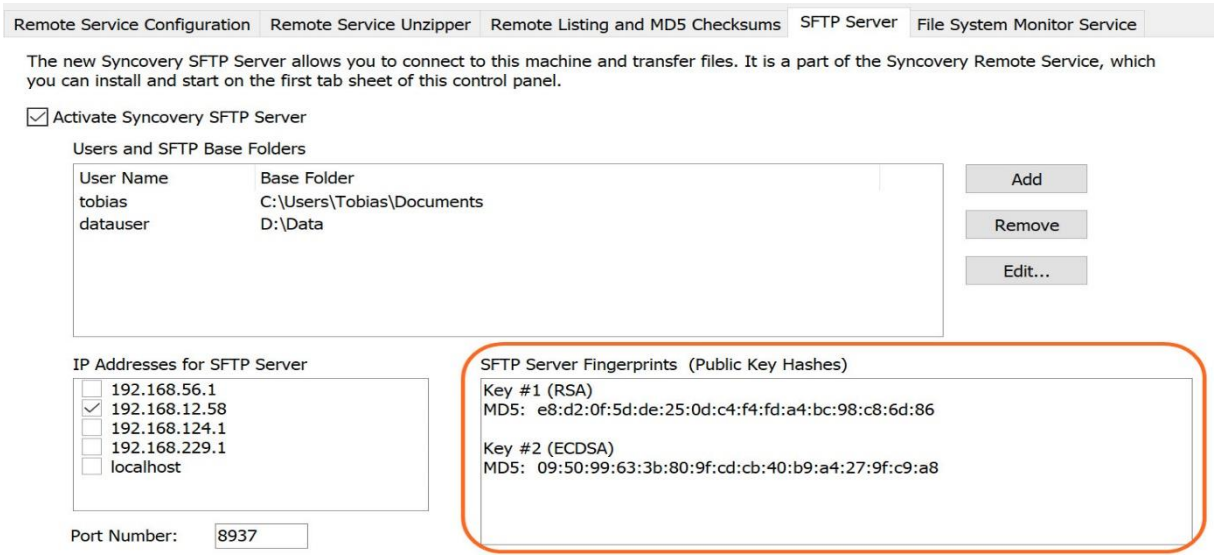

Finding out the server's public IP address

Since we are going to connect to your SFTP server over the Internet, we need to know the IP address or domain name under which your router connects to the Internet. If the router has **Dynamic DNS** features, you can use a domain name that resolves to your IP address. Alternatively, you can visit a page such as www.syncovery.com/myip.php from the PC where the SFTP server runs to find out the IP address.

Creating the Syncovery profile

Finally, you can go to the other PC and create your Syncovery profile (or job). Choose a local folder on one side, and click the Internet button on the other. Choose the protocol **SSH/SFTP**. The Internet Protocol Settings dialog can be completed like the following. Remember to put the other router's public IP address or dynamic DNS name in the URL field:

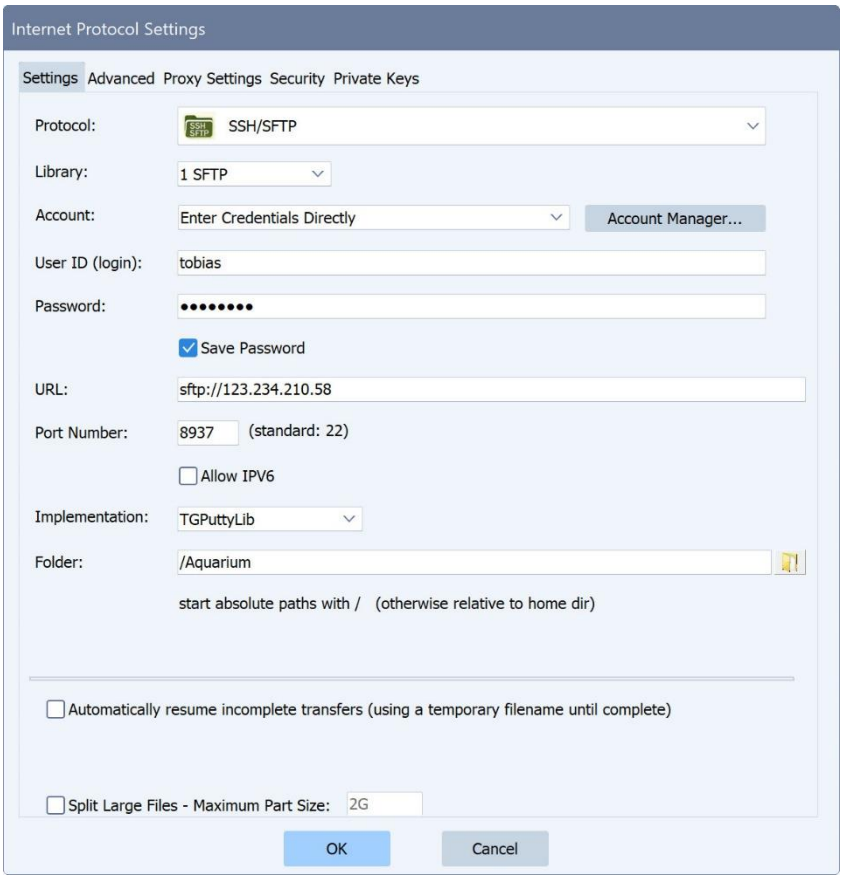

When the first connection is made, Syncovery will show a prompt asking you to confirm the **SFTP server's fingerprint**. Please verify it against the fingerprints shown by the server to ensure that your peer to peer connection will be safe.

When you click OK, Syncovery may ask you if you want to change the port to 22 (the SFTP default). Make sure you **don't let it change the port**! We really want to avoid using the default port number, as it is prone to continuous hacking attempts.

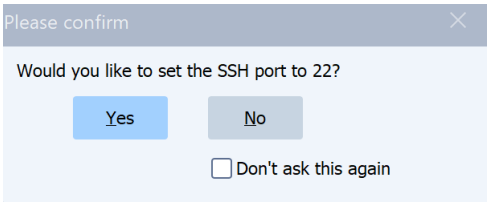

7.6.6 Major Syncovery Profile Settings

Under **Main Settings** in Syncovery 11, choose the **Copying Directions** and the **Sync Operation Mode**. Note that you don't actually have to let Syncovery copy any files – you can also just use it to compare the folders and show you the differences.

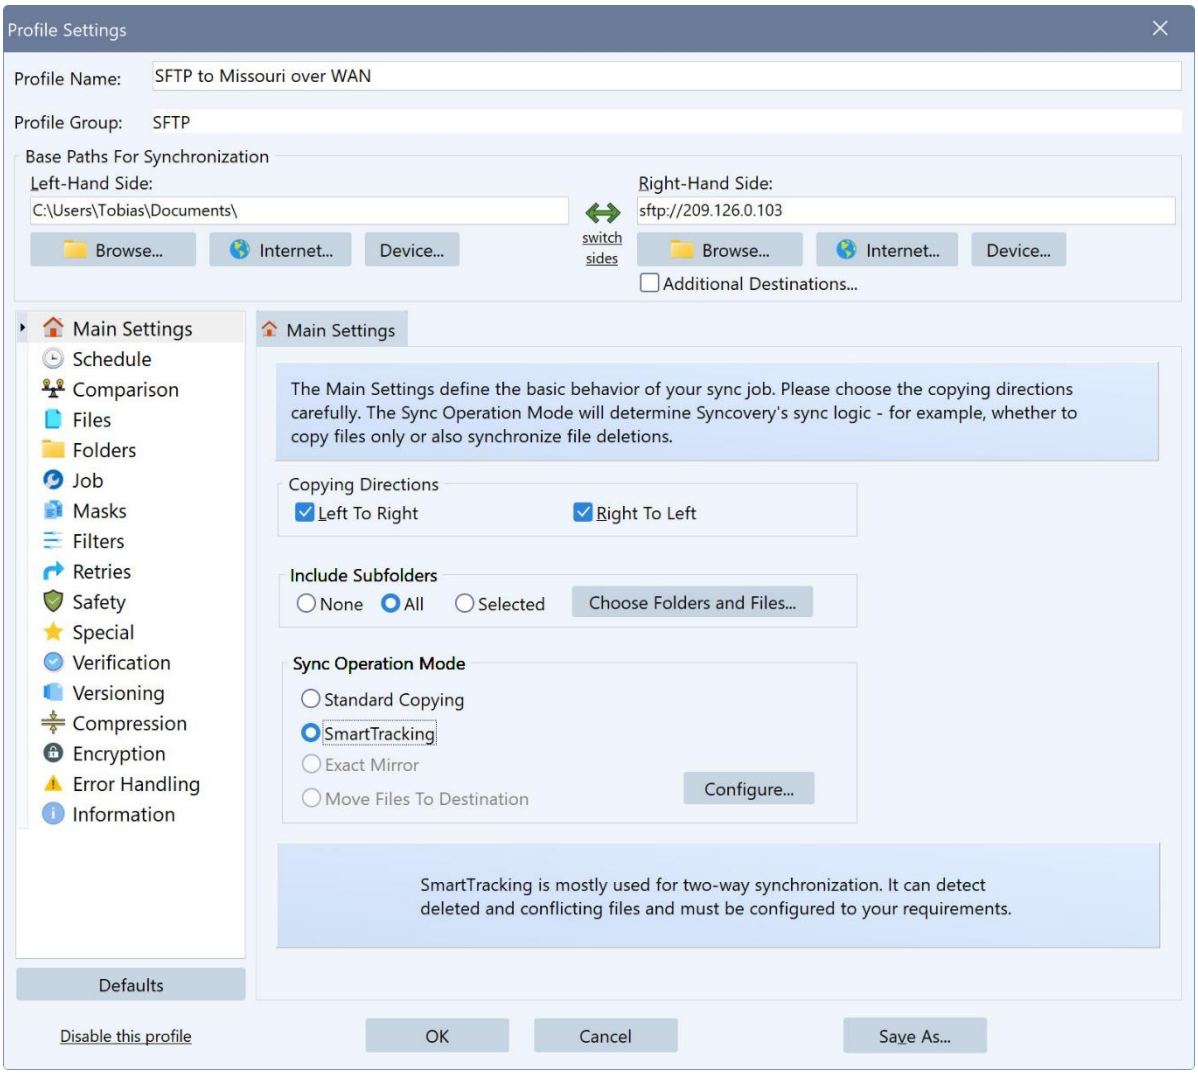

To ensure that no broken transfers are left as incomplete files, please choose the setting **Automatically resume (copy with temporary filenames)** in the **Files** category in the profile.

Finally, to greatly speed up the folder listing, specify the the Remote Service should be used for the listing. This works only if the profile's base path is specified under **Remote Listing** on the other PC, which we did as one of the first steps near the beginning of this guide.

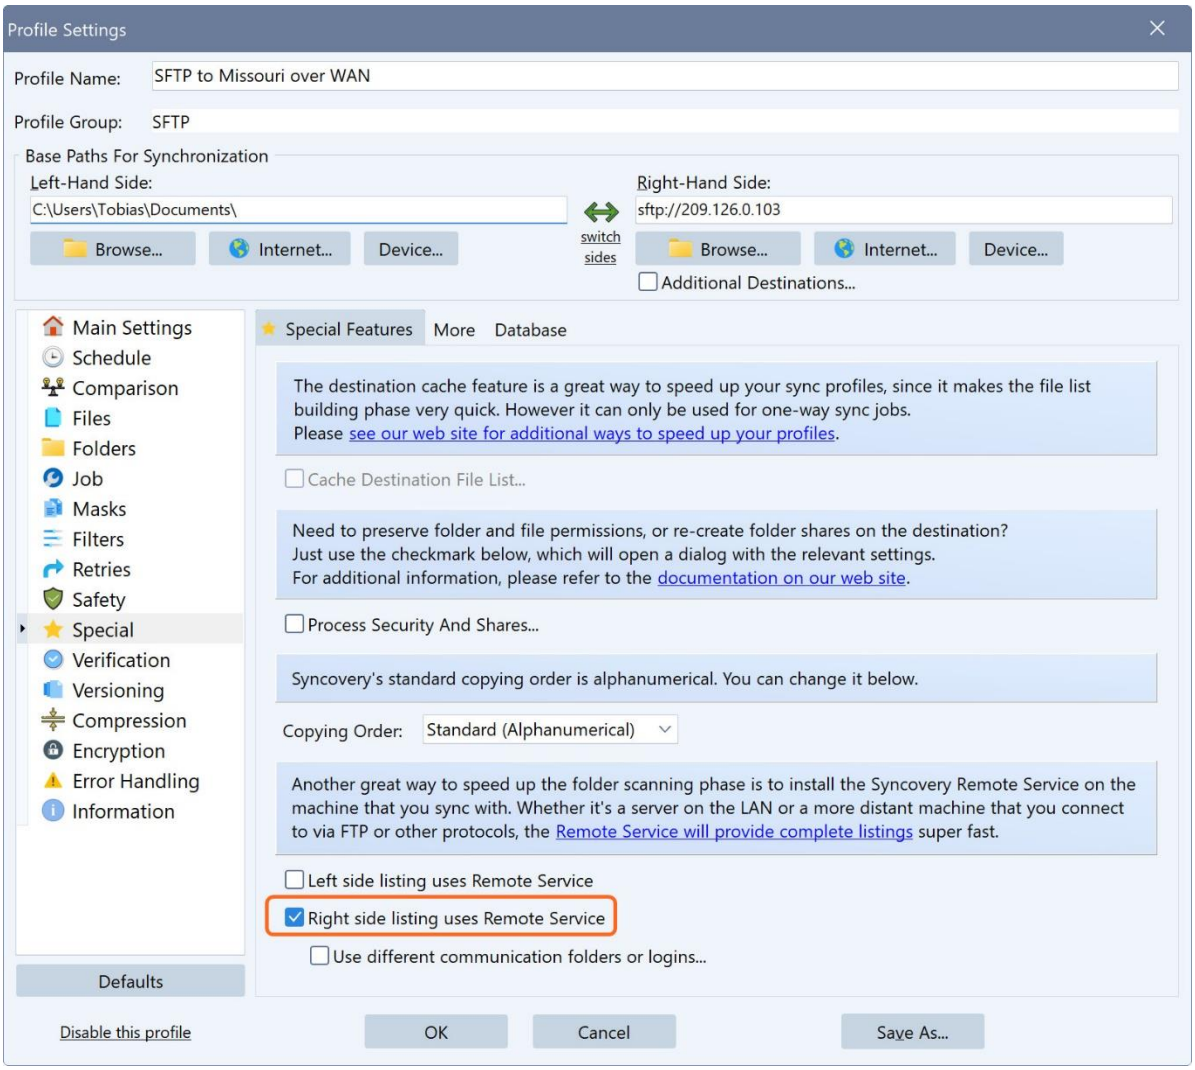

Starting the Syncovery Profile

You can now start the job. Remember that there are various ways to run a Syncovery profile. It can be started manually or by the scheduler. It can run in **Attended Mode**, or in **Unattended Mode**, or in the **Background**, or in the **Background With Preview**.

If you would like to see the list of proposed copying actions before letting Syncovery copy files, make sure you start the job in **Attended Mode** or in the **Background With Preview**. To see all choices for starting a profile, please right-click it from the Profile Overview in Advanced Mode, and you will see these menu items:

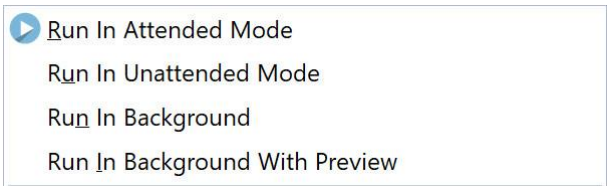

7.7 Connecting with Google Cloud Storage

Choice of Google Cloud Storage API

Syncovery supports two different protocols to access Google Cloud Storage: the older S3 compatible XML API, which is covered on a separate page. And the more **modern and powerful JSON API** , which is recommended to use and described on this page.

Since version 10.11.0, Syncovery is authorized for use with Google Cloud Storage using Service Accounts or the gcloud CLI. These methods allow flexibility in giving Syncovery either full Admin access to the Google Cloud Storage account, or allow it to work with individual buckets only. If you already have the gcloud CLI installed and connected to your GCS account, you're all set and you can start using Syncovery immediately.

Many steps described on this page are needed only if you want to authorize Syncovery directly, without the gcloud tools. If you have the **gcloud CLI**, it is much simpler. Just choose the Google Cloud Storage protocol in Syncovery and choose to authorize via **gcloud**.

Setting up a Service Account is not difficult, but it involves a few steps. This page will guide you through the steps of creating a Service Account for Syncovery and assigning the permissions it needs.

7.7.1: Creating a Service Account in the Google Cloud Console

This step is only needed if you don't have a fully authorized **gcloud command line tools** installation yet. A Service Account is easily created in the Google Cloud Console, on the Service Accounts page. Choose your project and click the button "**CREATE SERVICE ACCOUNT**":

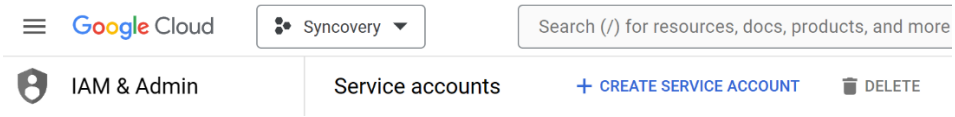

Next, you will see the following form, where you need to give the service a name and an optional description. The service account ID will be generated automatically. When done, click "CREATE AND CONTINUE":

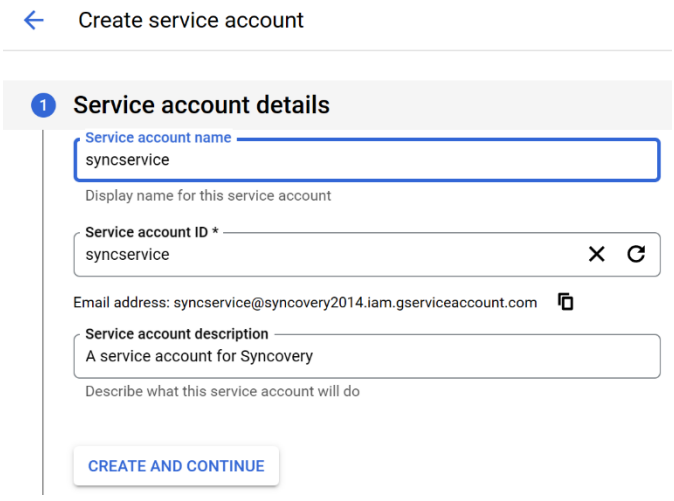

Now's the time to decide if Syncovery will have full control over the whole Google Cloud Storage account, or if you will later assign more granular permissions. If you don't want to bother with

individual bucket permissions, you can assign Syncovery the role “Storage Admin”. However this is not required.

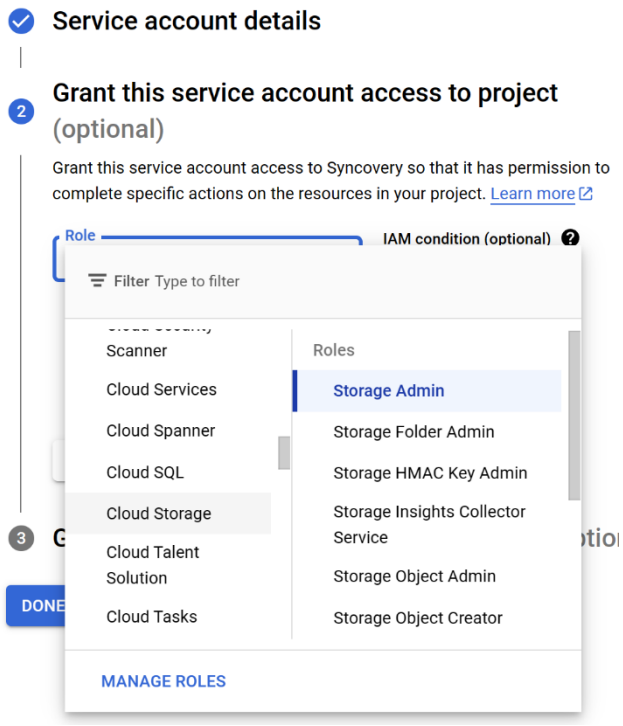

You can now click “DONE”, since we do not need step 3. The service account will now be created and you can see it in your list of service accounts. It will have an email address used for identification, such as “syncservice@syncovery2014.iam.gserviceaccount.com“. This is a good time to copy this address and save it somewhere for later use.

Syncovery can obtain access using the Service Account with a Private Key, or via the gloud CLI. If you have the gcloud CLI installed and activated, you do not need to give Syncovery a private key. Otherwise, you need to create a key for the account. Click on it and go to the “Permissions” tab as shown below, click on “ADD KEY”, and choose “Create new key”:

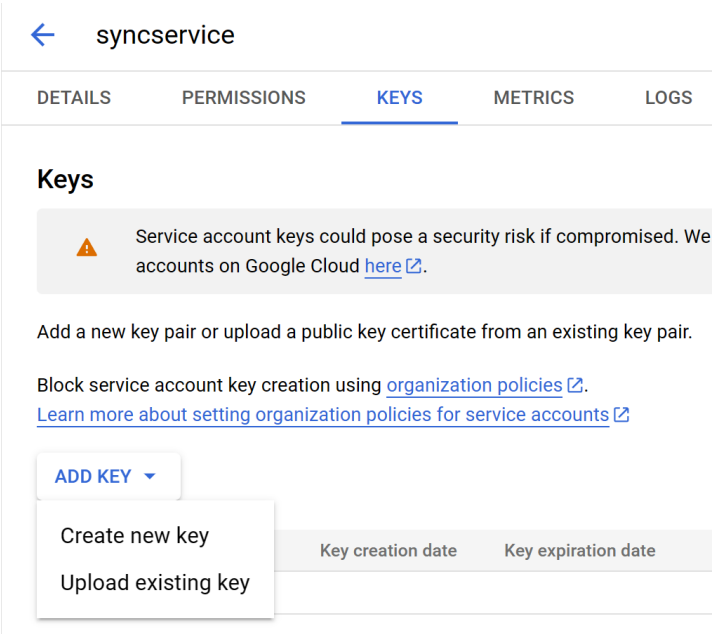

Choose the JSON format and click CREATE:

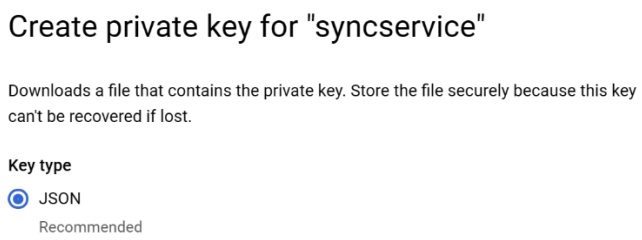

The Private Key is now saved to your computer in a json file, with a filename similar to “syncovery2014-9a29ca47fe28.json”. Save this file in a save place. The private key needs to be imported into Syncovery later on. It is extremely confidential and should be kept safe.

7.7.2: Assign Permissions for Individual Buckets

It is possible to assign permissions for each bucket separately. This step can be skipped if you assigned the Storage Admin role when creating the Service Account. To assign permissions on a bucket level, head over to the Google Cloud Storage Browser.

Click on the bucket you want to work with, and go to the PERMISSIONS tab:

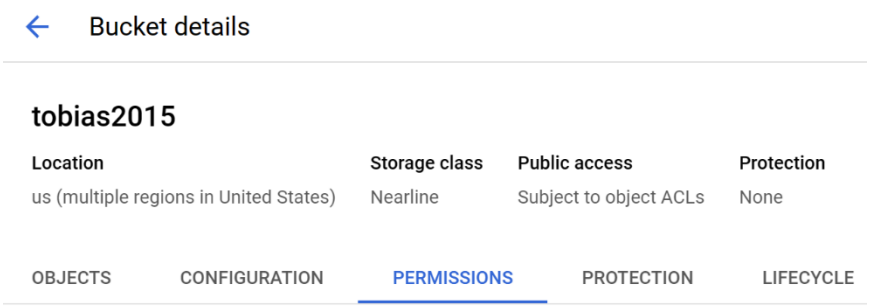

Further down, click on GRANT ACCESS:

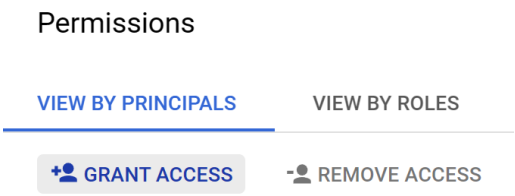

Next, you need to specify the service’s e-mail address, which is used as an ID. In this example, the e-mail address is syncservice@syncovery2014.iam.gserviceaccount.com.

For the best operation, please choose the role Storage Admin. This will give Syncovery full control of this bucket only, not of the whole Cloud Storage account.

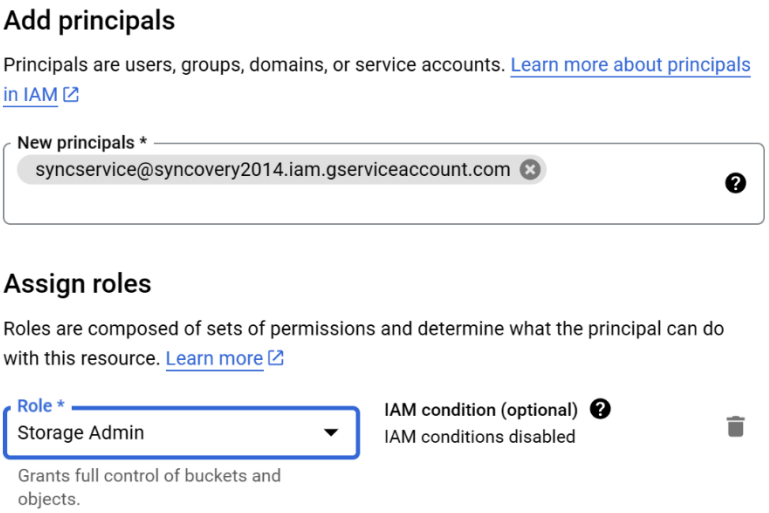

Click Save and you're done. The Service Account can now be used in Syncovery.

7.7.3: Using the Service Account in Syncovery

To access a Google Cloud Storage bucket in Syncovery, please click on the Internet button on one side in the profile, and change the protocol from FTP to Google Cloud Storage. If you will work with multiple different service accounts, you should type an optional internal ID into the field "Account (opt.)". This optional account name is only an identifier within Syncovery and does not have any real meaning. In particular, it does not have to be the storage account ID or e-mail address.

Click on the Settings... button to specify the Service Account Details:

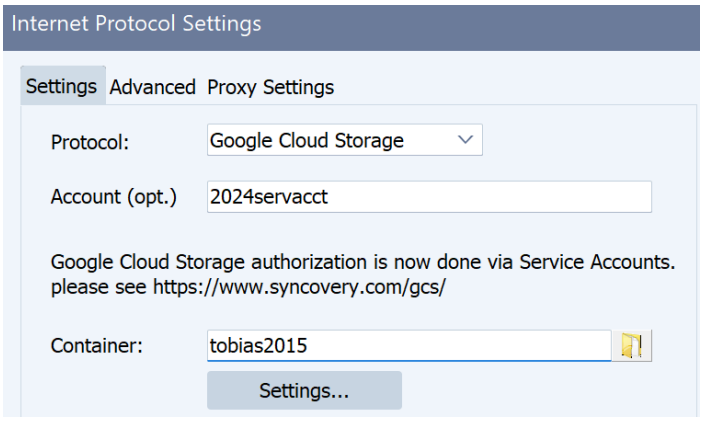

Syncovery will now ask you for the service account's email address:

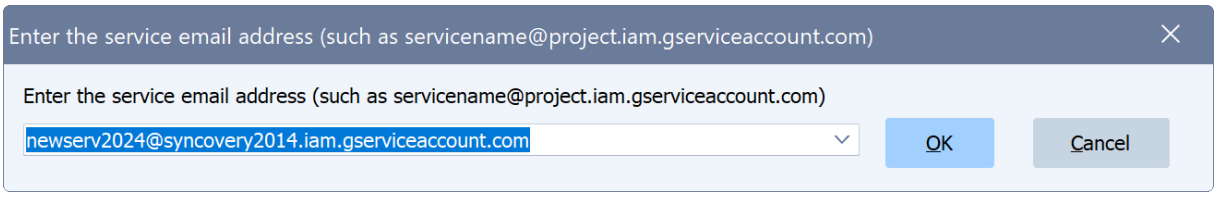

And for the private key. Open the previously saved json file in a text editor, and copy and paste its contents into the dialog box, which will look like this:

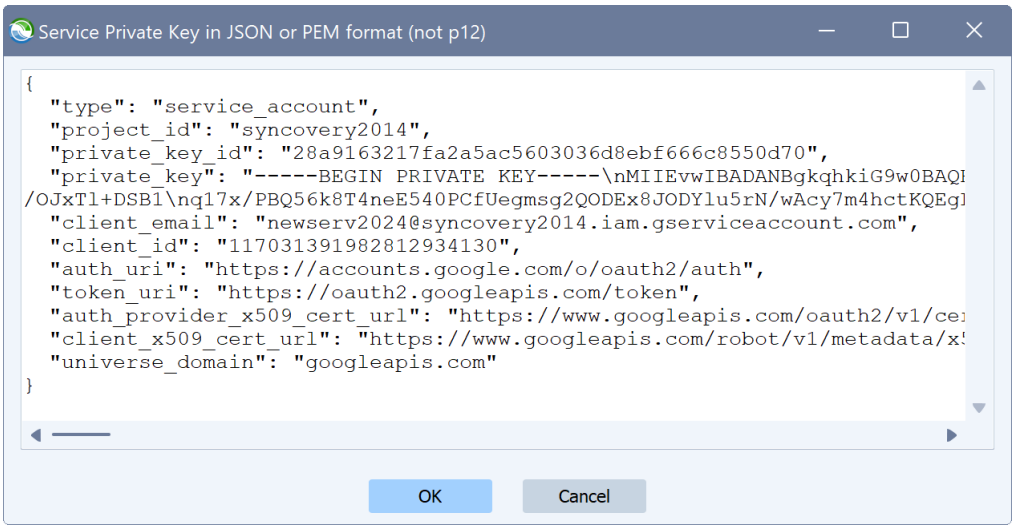

Now you're all set and the rest should be really easy.

If you know the bucket name that Syncovery should work with, you can type it into the "Container" field.

On the other hand, if you gave Syncovery full control over the storage account, you can click the Browse button next to the Container field to get a list of buckets to choose from.

Finally, you can click the second Browse button to choose a folder within the bucket.

7.8 Syncovery Monitoring Console – Manage Machines Remotely

Managing machines remotely with Syncovery's Montoring Tool

Syncovery includes a monitoring tool that allows you to view and manage Syncovery installations across multiple machines on a single dialog. You will find the Monitoring Tool in the File menu of Syncovery for Windows.

There are two ways how Syncovery can connect with remote machines:

7.8.1: Connect via Windows Networking (SMB/CIFS)

Connecting via Window networking is easy to configure, but it can only show the status of machines on the same LAN or VPN, and only if the machines are on and the Syncovery scheduler is running.

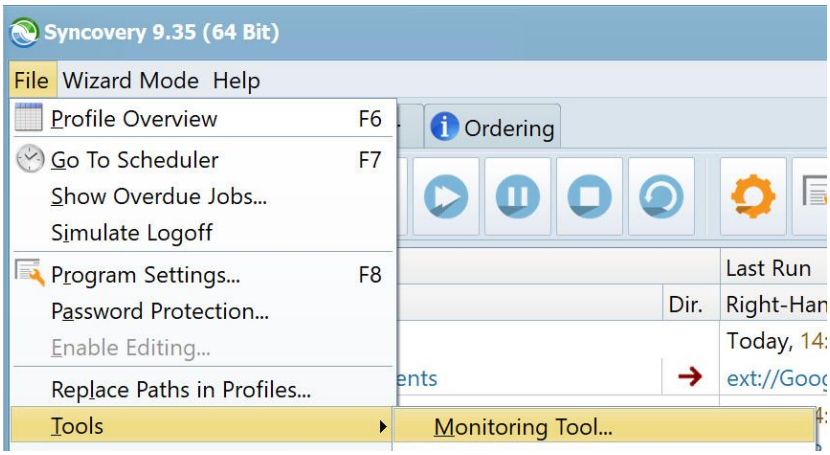

7.8.2: Exchange Information via Central Status Storage

Using a server storage location to save status files, you can manage any machine worldwide, and you can view the status of machines which are currently not connected, or even turned off. You can also monitor Syncovery running on Mac or Linux systems. Central status storage is a bit harder to set up, and the information shown can be slightly delayed. But it is a good choice if you can't connect directly via LAN or VPN, or if you don't have Windows networking credentials for the remote machine. The central status storage can be a folder on an (S)FTP server or use any of the other Internet/cloud storages supported by Syncovery.

7.8.3 Setting up Mode 1 (Windows Networking)

To monitor other machines on the LAN or VPN, you need to enter their computer names or IP addresses on the tab sheet "**Configuration via Direct Connection (CIFS)**", as shown in the following screenshot. The machines need to be visible in Windows Explorer, so for example to monitor **WIN10HOME**, you need to be able to see the network shares of **\\WIN10HOME** in Explorer.

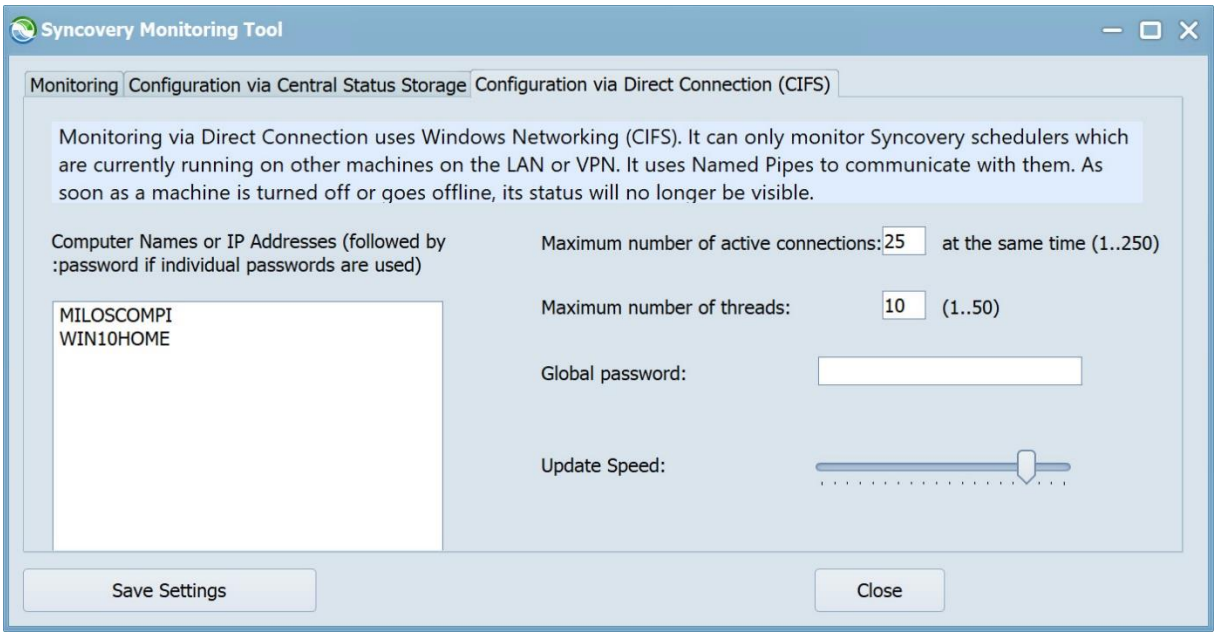

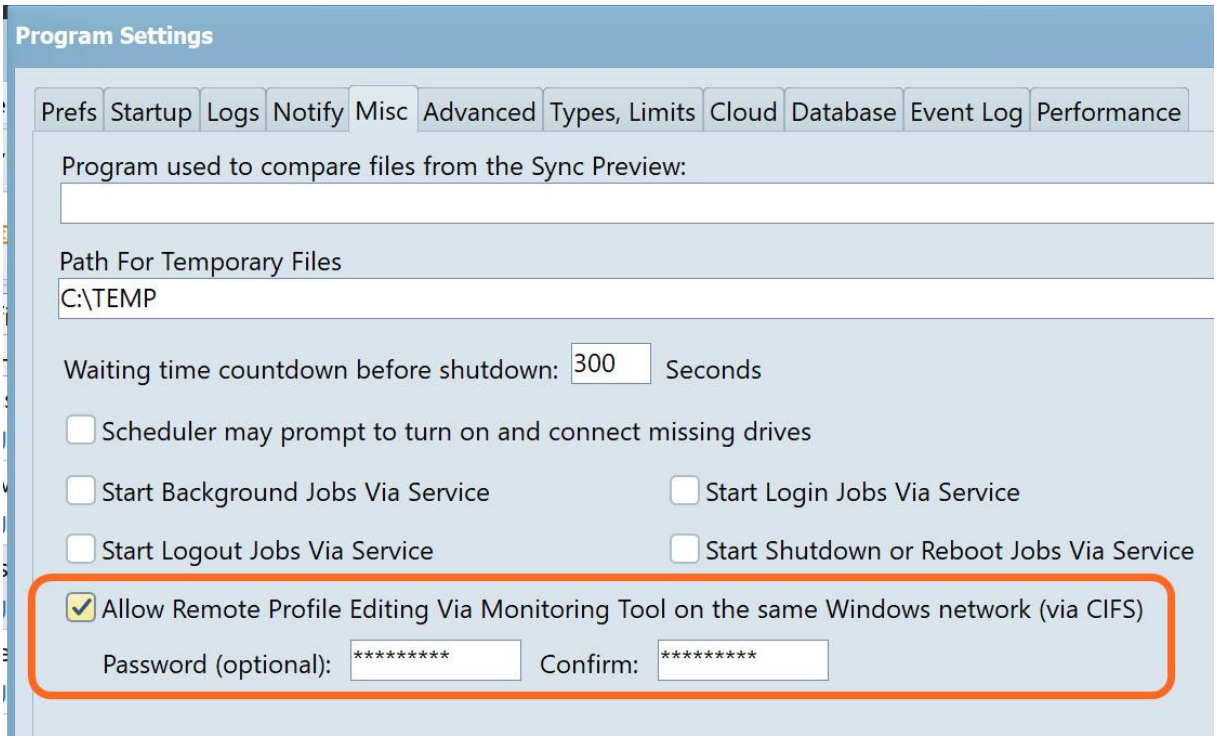

A password can be specified globally or for each machine. The password however is only needed to modify profiles, not to view the machine status. Permission to edit profiles must be given on the Program Settings dialog of each Syncovery installation that you want to manage. Please refer to the following screenshot.

The following screenshot shows how the machine status is displayed.

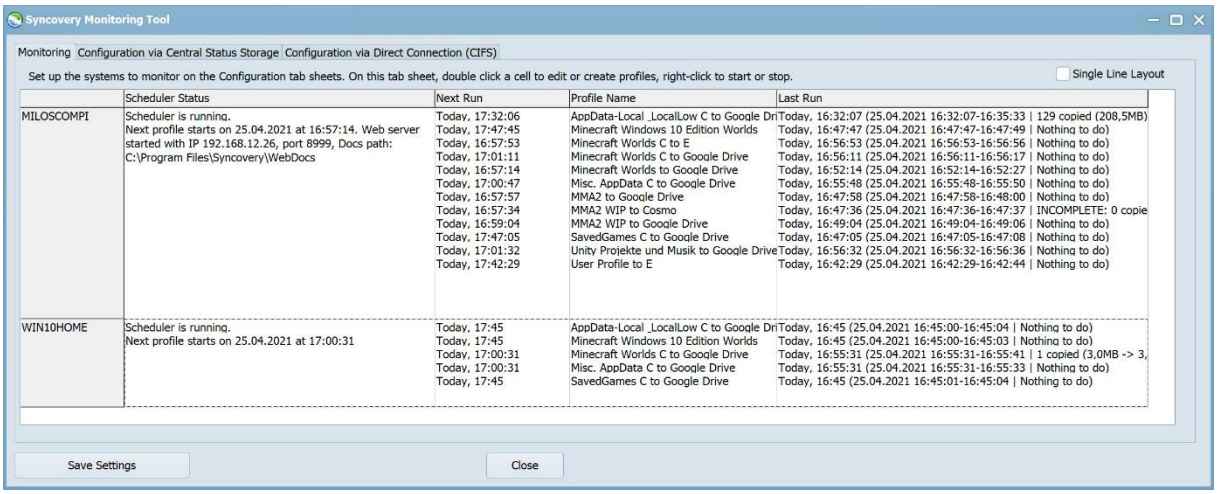

A context menu is available, allowing you to perform various administrative tasks, such as starting or stopping the schedule, and editing profiles. If editing profiles does not work, please check that "Remote Profile Editing" is allowed on the other machine, as shown in a previous screenshot, and that the password matches your configuration.

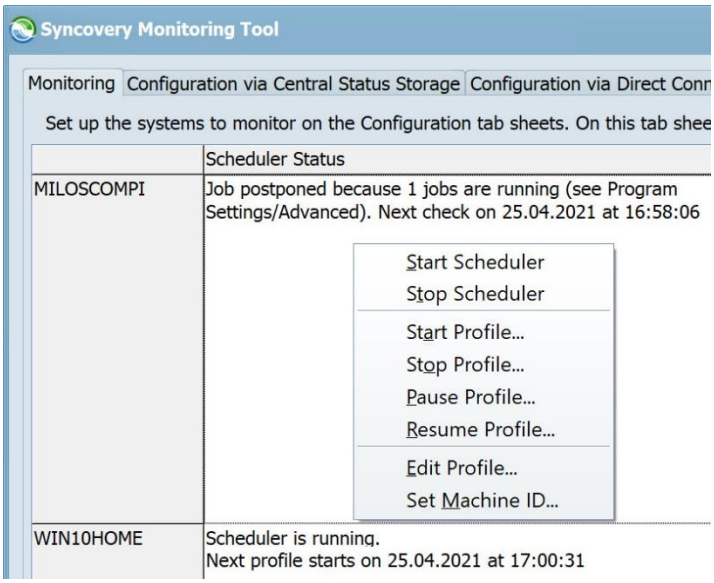

7.8.4 Setting Up Mode 2 (Exchange Information via Central Status Storage)

Setting up the monitoring via central status storage is more complex. It is done mainly on the tab sheet shown in the screenshot below, but also involves editing Syncovery.ini directly on the clients to be monitored. You will find detailed steps and explanations below the screenshots.

The main requirement is that you need a folder on a server or some cloud storage that can be used. In this example, we will use **SFTP** to access **monitoring.syncovery.com** with user name **admin** and password **xyz** *(these credentials are just an example and do not work in real life)*.

Two separate SFTP folders are actually used:
a) one for **status files** that the clients send to the monitoring tool, and
b) another one for **command files** that the monitoring tool sends to the clients. These commands can include starting or stopping the scheduler, editing profiles etc.

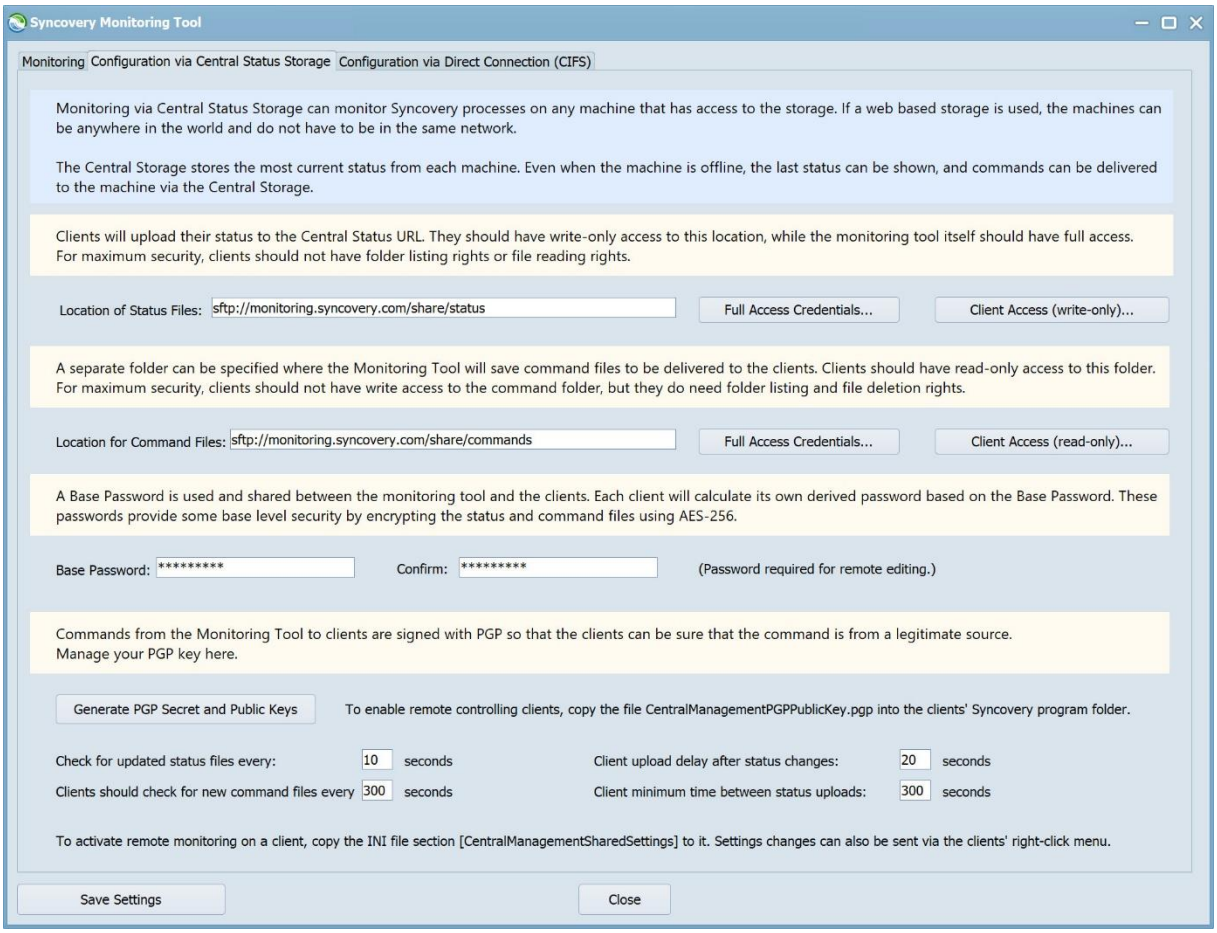

Step 1: Specify Location for Status Files

In the first step, we will specify the SFTP location for **status files**. Clients will upload their status and progress information as small files to this folder. Each client uses a unique file name, so that the files of many clients can all exist in the same folder. You need to click the "Full Access Credentials..." button to specify the Internet Protocol, and user name and password. You can optionally specify a separate user for the clients with write-only access, but this is only necessary if you want to implement the highest possible security standard.

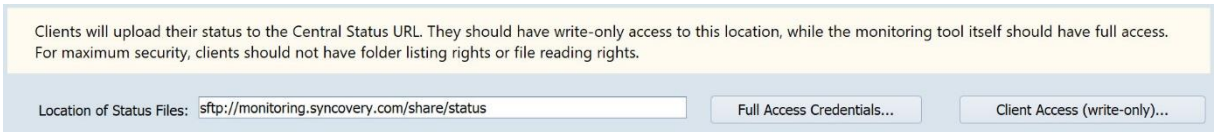

Step 2: Specify Location for Command Files

To be able to send commands to clients, you need to specify a separate folder for **command files**. The Monitoring Tool will upload commands to this folder when you want to start or stop the scheduler, edit or manually run profiles, update Syncovery etc. Again you need to click the "Full Access Credentials..." button to specify credentials. You can optionally specify a separate user for the clients, who need read-only and delete access (they have to be able to delete command files after executing them).

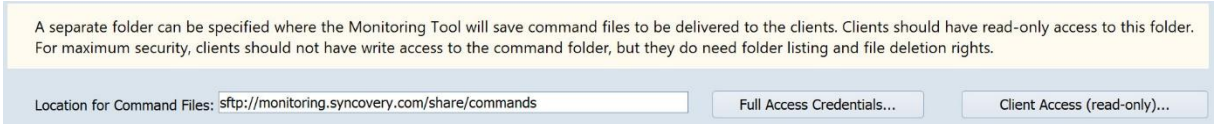

Step 3: Specify a Base Password

A base password is needed to encrypt the communication. Each client will actually derive its own password from the base password, so that clients cannot see each others' statuses. But all can communicate with the monitoring console.

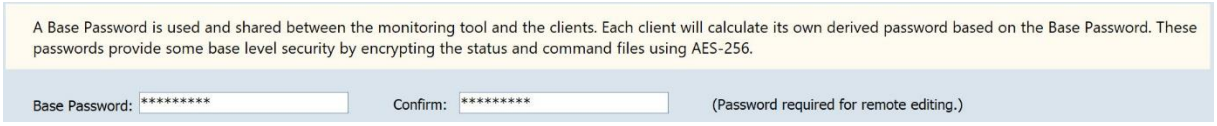

Step 4: Set up PGP to make command files trustworthy

To ensure authenticity of the commands that you send to clients, they are encrypted and signed with a PGP secret key. Please click the button Generate PGP Secret and Public Keys and follow the Wizard steps to generate your keys.

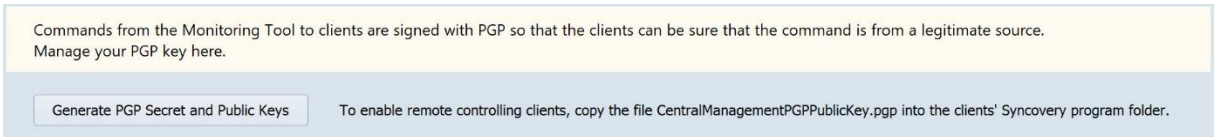

At the end of the key generation, a message box will inform you about the location of the generated files. They are usually similar to these:
C:\Program Files\Syncovery\CentralManagementPGPPublicKey.pgp
C:\Users\Tobias\AppData\Roaming\CentralManagementPGPPrivateKey.secret

The secret key, along with its password, will give you the power to send commands to clients. Because there can be only one, you may need to share it with other administrators who need this capability. The secret key is not shared with client installations.

The public key however needs to be copied to clients, so they can verify the authenticity of incoming commands. So you need to copy the file **CentralManagementPGPPublicKey.pgp** to **C:\Program Files\Syncovery** on all client computers that you wish to control.

Step 5: Additional Settings

The additional settings allow you to fine-tune how frequently the clients (and the monitoring tool) will connect with your status storage server. Depending on the number of clients, you can allow smaller or larger intervals. If the intervals are larger, it will take longer for status information to update, or for commands to be processed.

Step 6: Copy Configuration to Clients

Before copying the configuration to clients, please click on all the **four credentials buttons** once more and verify the settings. Ensure that all folders are absolute paths that start with a slash on the four dialogs.

Your Syncovery.ini file (usually in C:\ProgramData\Syncovery) will now contain two new sections related to monitoring. These lines contain all of the above settings.

```
[CentralManagementSharedSettings]
ClientCommandCheckSeconds=300
ClientStatusUploadDelay=20
ClientMinimumStatusUploadPause=300
CentralManagementURL=sftp://monitoring.syncovery.com/share/status
CentralManagementClientAccessData=+ $a 1Œ 2ž}¤‘\ #Fd„yTº /b± ?\°C;cøÍr„ÀAºË2Ýbd 1V #WÍSq
CentralManagementCommandsClientAccessData=+ $a 1Œ 2ž}¤‘\ #Fd„yTº /b± ?\°C;cøÍr„ÀAºË2Ýbd
CentralManagementCommandURL=sftp://monitoring.syncovery.com/share/commands
CentralManagementBasePasswordUtf8=
[CentralManagementSecretSettings_DO_NOT_SHARE]
CentralManagementFullAccessData=+ $a 1Œ 2ž}¤‘\ #Fd„yTº /b± ?\°C;cøÍr„ÀAºË2Ýbd 1V #WÍSqÍ
CentralManagementCommandsFullAccessData=+ $a 1Œ 2ž}¤‘\ #Fd„yTº /b± ?\°C;cøÍr„ÀAºË2Ýbd 1V
```

Please save only the **[CentralManagementSharedSettings]** to a text file and add this section to **every client's Syncovery.ini** if you want its status to appear in your Monitoring Tool. In addition, remember to copy the file **CentralManagementPGPPublicKey.pgp** to **C:\Program Files\Syncovery** on all client computers that you wish to control, as mentioned in Step 4. The Syncovery scheduler or service should be completely stopped when making these changes, and then restarted.

Step 7: Check Functionality and Troubleshooting

After setting this up, and restarting the scheduler on the client machine, you should begin to see client information on the Monitoring tab sheet relatively quickly.

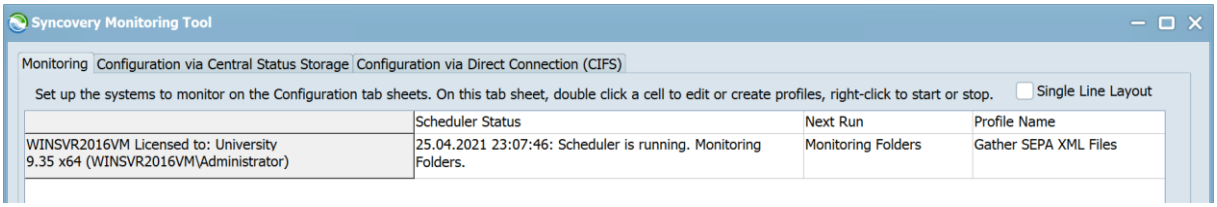

If you don't see such information, it's usually because of the Internet/FTP settings. The client might not have been given the correct credentials or the correct URL. Check if files like this appear on the status storage:
WINSVR2016VM.{1B6D0D72-3285-4D5D-8AD7-B38041BF0F77}.syncoverystatus

Sometimes you need to be patient for a status to arrive. You can probably change the additional settings and especially reduce the "client minimum time between status uploads".

If there are still issues, you can have Syncovery create log files for the status file transfers by adding this line to the [CentralManagementSharedSettings] section of Syncovery.ini:
LogMonitoringTransfers=1

Step 8: Try Sending Commands to Clients

You can use the context menu to send various commands to the client. The menu is similar to the screenshot near the top of this page. You can start/stop the scheduler, edit or create profiles, as well as update Syncovery. Because the process of transmitting the commands can take a few minutes, you will see the processing state in the grid, such as:
*** Command Uploaded ***
*** Command Downloaded ***
*** Command Processed ***

Step 9: Activate Monitoring for Linux and Mac Clients

On Linux machines, you can import the CentralManagementSharedSettings by saving the relevant Syncovery.ini section to a separate, small ini file and importing it via the Program Settings dialog, tab sheet “Startup” by clicking the button “Import Config Lines (INI Style)...“.

On macOS (since version 9.35a), there’s a button “Import Config Lines” on the Prefs tab sheet of the Program Settings dialog to import the configuration.

7.9 Linux Documentation

7.9.1 Linux and FreeBSD Platforms that Syncovery Runs On

Syncovery runs on virtually all recent Linux and FreeBSD distributions as well as many NAS brands. It is available for various different CPU types, including Intel/AMD and ARM/AArch64 in both 32-bit and 64-bit.

Syncovery supports these **Linux** distributions and NAS systems:

- **Debian** and derived Linux types, including **Ubuntu**, Linux MINT, MX Linux, Kali Linux, **Raspberry Pi OS**, and many others
- **Red Hat** and distributions using the RPM package manager, including **Fedora, CentOS, openSUSE**, Rocky Linux and AlmaLinux
- Additional distributions where dpkg can be used, including **Arch Linux / Manjaro**, **UGREEN NASync** and **Netgear ReadyNAS** devices
- Native installation packages are available for NAS models from **QNAP, Synology, ASUStor**, and Western Digital (WD)
- NAS models suported via our Unified Installer for Windows: **Seagate**, Thecus, and Zyxel.
- **FreeBSD** and derived systems, including **FreeNAS/TrueNAS** and XigmaNAS
- Simple tar.gz distribution packages are available for any other type of Linux

7.9.2 Syncovery Installation Guides

Various methods are available to install Syncovery, your favorite Linux Backup and Sync software, on your system. Native installation packages exist in different formats, and generic tar.gz archives can also be used. If you are aware of the system and CPU type that you want to install Syncovery on, it is easy to pick the correct download. On the other hand, our Unified Syncovery Linux Installer is available to install Syncovery automatically and avoid having to pick a particular download package.

After the installation, you can reach Syncovery's **web GUI via port 8999**. For example, enter the following line in your web browser (if your Linux system has a desktop GUI with browser):
localhost:8999

If you are running your web browser on a different system, you need to enter the Linux system's IP address, for example:
192.168.1.58:8999

The initial **user name** is **default** and the password is **pass**.

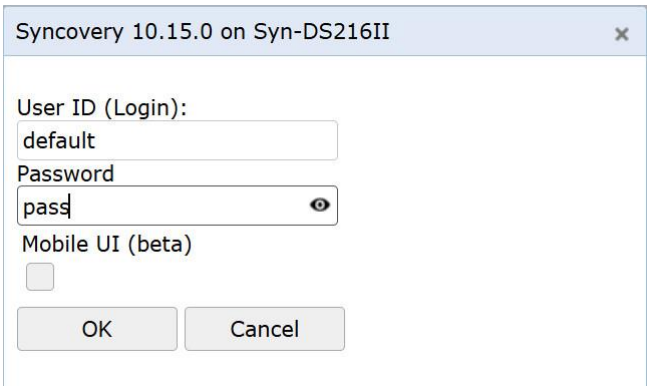

The Unified Syncovery Linux Installer

The Unified Syncovery Linux Installer is a tool that runs on Windows and installs Syncovery on virtually all supported platforms using an SSH shell connection. You can use it if you have an SSH login that has admin or root privileges (or is allowed to use sudo). This useful installer tool is available from our Syncovery for Linux Download Page.

If you prefer to install Syncovery without using a Windows tool, you will find instructions for the different platforms in the following paragraphs:

Installing Syncovery on a Synology NAS

You will find native installation packages for Synology on our Syncovery Download Page for Synology NAS. These can be installed on Synology DSM 6 or 7 using the manual installation in Package Center. It is important to choose the correct download, depending on your CPU type and DSM version. In addition, if you have DSM 7 or later, you need to grant Syncovery access to the folders that you want to work with.

Installing Syncovery on a QNAP NAS

Our QNAP Download Page contains packages for installation in the QNAP App Center, using the manual installation button. You will also find an older Syncovery version in the QNAP app catalog, but it is highly recommended to use the latest version directly from our web site. Please check which CPU your NAS has and download the correct package.

Installing Syncovery on ASUStor, WD and others

Syncovery can be found in the ASUStor App Central and we also offer manual downloads on our ASUStor Download Page. There's also a separate download page for Western Digital (WD) NAS devices. Additional NAS brands are supported by our Unified Syncovery Linux Installer.

Installing Syncovery on Debian, Ubuntu, Linux MINT, Raspberry Pi, UGREEN NASync, Netgear ReadyNAS and other Debian-based systems

Installation on these systems is done using the .deb packages. Please download the correct package from our Linux Download Page and install it, either using the installer that opens it on your Linux desktop, or using a Terminal or SSH shell window. Here's an example for the installation command line:
sudo dpkg -i Syncovery-10.15.8-amd64.deb

In some cases, dpkg may show an error, and if you are certain that you have the correct package, you can override any errors using this command line:
sudo dpkg -i --force-all Syncovery-10.15.8-amd64.deb

Installing Syncovery on Red Hat, Fedora, CentOS, openSUSE and other RPM-based systems

These systems use .rpm packages to install software, including Syncovery. Just download the correct RPM package from our Linux Download Page and install it with a command line such as:
sudo rpm -i Syncovery-10.15.8-amd64.rpm

Alternatively, you can use yum:
sudo yum install Syncovery-10.15.8-amd64.rpm

Installation on FreeBSD, FreeNAS/TrueNAS and XigmaNAS

You will find installation instructions for FreeBSD based systems on our FreeBSD Download Page.

Generic Linux Download Page

For an alternative overview of the installation types, including links to the various download pages, please see our Generic Linux Download Page.

7.9.3 Using the Syncovery Web GUI for your Linux Backup and Sync

When you install a NAS, Debian or RPM package, the Web GUI should become available automatically.

If you choose the .tar.gz download, you can activate the web GUI by running these SyncoveryCL commands in a Terminal window:
./SyncoveryCL SET /WEBSERVER=localhost *(to configure the web server)*
./SyncoveryCL start *(to start SyncoveryCL – the start command line parameter is recognized since v8.25)*

Open the web GUI in a Browser window by entering: localhost:8999. Unless you are accessing the Web GUI from localhost, it will ask for username and password. The defaults are: user name = default, password = pass.

This is the complete command line to configure the web server:
SyncoveryCL SET /WEBSERVER=localhost /WEBUSER=username /WEBPASS=password /WEBPORT=port /WEBDOCSPATH=path_to_web_docs_folder

To turn the web server OFF (requires a restart if SyncoveryCL is already running):
SyncoveryCL SET /WEBSERVER=OFF

The Web Docs Folder is called WebDocs and it is included in the tar.gz archive that you download from this page.

7.9.4 Additional Information

SyncoveryCL will create a .Syncovery folder for configuration, logs, and databases in the HOME folder. It depends on the HOME environment variable. You can also use SYNCOVERY_HOME, which takes precedence. **HOME should not be /.**

7.9.5 Syncovery Command Lines, not only for Linux Backup and Sync

Here are some example command lines to configure and run Syncovery. Additional command lines can be found here. You only need to use command lines if the web GUI is not used or not ready yet.

Run a job without saving it to the configuration file:
SyncoveryCL RUN /LEFT="/home/tobias/Documents/"
/RIGHT="/home/tobias/DocumentsCopy" /L2R

Add this same job to the configuration file:
SyncoveryCL ADD /NAME="LocalTest" /LEFT="/home/tobias/Documents/"
/RIGHT="/home/tobias/DocumentsCopy" /L2R

Run an existing job:
SyncoveryCL /RUN="LocalTest"

Configure a scheduled job (every 30 minutes):
SyncoveryCL ADD /NAME="LocalTest" /LEFT="/home/tobias/Documents/"
/RIGHT="/home/tobias/DocumentsCopy" /L2R /Sched /Rep /Days=0 /Mins=30

List Jobs in Config File:
SyncoveryCL /LIST

Disable a job:
SyncoveryCL CHANGE "SFTPTest" /Disabled

Enable a job:
SyncoveryCL CHANGE "SFTPTest" /Disabled=No

Show job configuration:
SyncoveryCL SHOW "SFTPTest"

The easiest way to get the command line with additional parameters is to create the job in the Windows or Mac version, and go to the Information tab sheet in the profile editor, where you will see the profile XML and command line.

Start the Scheduler (only necessary if you don't use the Debian or RPM packages):
SyncoveryCL *(runs in the foreground)*
or
SyncoveryCL start *(becomes a daemon – supported since v8.25)*
or
SyncoveryCL & *(runs in the background but not as a daemon)*

Stop the Scheduler:
SyncoveryCL /STOPTIMER

See Scheduler Status:
SyncoveryCL /STATUS

See Continuously Updated Scheduler Status:
SyncoveryCL /CONTSTATUS

Upload to FTP:
SyncoveryCL ADD /NAME=FTPTest /LEFT="/home/tobias/Documents/"
/RIGHT="ftp://username:password@yourdomain.com/FolderName" /L2R

Upload to SFTP:
SyncoveryCL ADD /NAME=FTPTest /LEFT="/home/tobias/Documents/"
/RIGHT="sftp://username:password@yourdomain.com/FolderName" /L2R

Upload to SFTP with a certificate:
SyncoveryCL SET /CERT=/home/tobias/tobias_rsa
SyncoveryCL ADD /Name=WithKey /Left=/home/tobias
/Right="sftp://tobiaskey@192.168.10.20/Tests" /L2R
/RightFTPSettings="SFTP:Port=2222,AbsolutePath=N,Flags=UTF8+NoCertPass+UTC,"Cert=tobias_rsa""
SyncoveryCL /RUNX=WithKey /ACCEPTSERVER

Upload to WebDAV (example: Strato HiDrive):
SyncoveryCL ADD /NAME=WebDAVTest /LEFT="/home/tobias/Documents/"
/RIGHT="https://user:pass@webdav.hidrive.strato.com/users/yourusername/folder" /L2R
/EXCL=.*

Upload to Amazon S3:
SyncoveryCL ADD /LEFT="/home/tobias/Documents/"
/RIGHT="S3://yourid:yoursecretkey{.:CRED:.}bucketname/Documents" /L2R
/CacheDestination=Yes /EXCL=.* /NAME=S3Test

Upload to Microsoft Azure:
SyncoveryCL ADD /NAME=AzureTest /LEFT="/home/tobias/Documents/"
/RIGHT="AZ://yourid:yoursecretkey{.:CRED:.}container/Documents" /L2R

Upload to Amazon Glacier:
SyncoveryCL ADD /NAME=GLTest /LEFT="/home/tobias/Documents/"
/RIGHT="GL://yourid:yoursecretkey{.:CRED:.}vaultname/Documents" /L2R /EXCL=.*

Upload to Rackspace:
SyncoveryCL ADD /LEFT="/home/tobias/Documents/"
/RIGHT="ext://yourid:yourkey{.:CRED:.}containername/Documents" /RProt=Rackspace /L2R
/EXCL=.* /NAME=RackTest

Upload to SugarSync:
SyncoveryCL ADD /LEFT="/home/tobias/Documents/"
/RIGHT="ext://your@email.com:yourpass{.:CRED:.}My SugarSync/Documents"
/RProt=SugarSync /L2R /EXCL=.* /NAME=SugarTest

Upload to Backblaze B2 Cloud Storage:
SyncoveryCL ADD /LEFT="/home/tobias/Documents/"
/RIGHT="ext://account_id:application_key@BucketName/Documents" /RProt=B2 /L2R
/NAME=B2Test
You can get your Account ID and Application Key from the “Buckets” page on Backblaze B2.

Cloud Services with OAuth – a prompt will appear to open the OAuth URL in a browser.

Upload to box.com:
SyncoveryCL ADD /LEFT="/home/tobias/Documents/"
/RIGHT="ext://OptionalAcctID@Box/Documents" /RProt=Box /L2R /EXCL=.* /NAME=BoxTest

Upload to Microsoft OneDrive:
SyncoveryCL ADD /LEFT="/home/tobias/Documents/"
/RIGHT="ext://OptionalAcctID@OneDrvNew/Documents" /RProt=OneDrvNew /L2R /EXCL=.*
/NAME=OneTest

Upload to DropBox:
SyncoveryCL ADD /LEFT="/home/tobias/Documents/" /RIGHT="ext://DropBoxV2/Documents"
/RProt=DropBoxV2 /L2R /EXCL=.* /NAME=DropBoxTest

SyncoveryCL /RUN=DropBoxTest

General Configuration

Set up email noticiations using Gmail:
SyncoveryCL SET /GMAIL=youremail@gmail.com /SMTPPASS=xxx

Specify mail recipients (if different from the Gmail address from previous command line):
SyncoveryCL SET /EmailRecipients=email1,email2,email3

Apply global settings similar to the [Main] section of the INI file on Windows:
SyncoveryCL SET /SettingName=Value

For example:
SyncoveryCL SET /S3PartSize=536870912

Export a profile to XML:
SyncoveryCL /EXPORTONEXML="Profile Name"

Import profile(s) from XML:
SyncoveryCL /IMPORT="/path/to/theprofiles.xml"

See also The Syncovery Command Line

Imprint

This manual was released by:

Super Flexible Software GmbH & Co. KG
Buddenstr. 29-31
48143 Münster
Germany

www.syncovery.com