Whitepaper

Migration to OpenCloud

Reliable and without downtime

https://opencloud.eu

OpenCloud

Table of Content

Migrate to OpenCloud without data loss or downtime

Many file-sharing solutions grow alongside an organisation's needs but eventually reach their limits. The reasons for switching are varied: rising licence and operating costs, increasing compliance requirements, performance issues as user numbers grow, or the desire for greater digital autonomy. Sometimes it is also because a manufacturer discontinues or sells its product.

However, the specific reason is secondary. What is crucial is ensuring that the transition goes smoothly. Many IT departments put off the change because the migration itself is seen as the real risk.

This whitepaper shows that the switch to OpenCloud is less complex than many people expect. The database-free architecture significantly simplifies the migration process from a technical perspective. Four clearly defined phases ensure that data, structures and permissions are transferred completely and securely. The aim is a transition that is as seamless as possible for end users – without data loss, without long downtimes and without unnecessary extra work in their day-to-day operations.

OpenCloud is based on a fork of the open source software 'ownCloud Infinite Scale' (OCIS), whose components were co-developed by developers from the science organisation CERN and other active contributors. OpenCloud is now being further developed by the Heinlein Group with new ideas and a clear focus on data protection, interoperability and sustainable digitalisation.

Why established infrastructures make the transition difficult

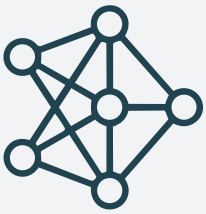

Many organisations maintain file-sharing platforms that were introduced years ago and have grown organically ever since. User accounts, group structures, file shares and permissions reflect years of administrative decisions. What was once manageable eventually becomes a complex web of interdependencies that is rarely fully documented.

Added to this are outdated technical systems. Traditional file-sharing platforms are based on technology stacks developed for a different era: interpreted code, relational databases for metadata, and complex caching layers. The larger the installation, the more time-consuming maintenance, updates and troubleshooting become.

Organisational challenges should not be underestimated either: when migrating to a new platform, administrators must reconfigure desktop clients on all end devices and switch existing connections to other systems to new endpoints.

All this does not make switching to a new file-sharing platform impossible, but it is complicated enough to put it off. As long as the existing system is somehow still working, concerns about the migration project take precedence.

The result: organisations remain stuck with platforms that hamper their operations and tie up a disproportionate amount of resources.

How OpenCloud simplifies migration

OpenCloud is structured completely differently from traditional file-sharing platforms, and this has a direct impact on the migration process. OpenCloud does not use a relational database at all. The platform stores all data (including metadata) directly in the file system. This offers significant advantages when migrating to OpenCloud: on the destination side, there is no database schema that needs to be prepared, populated or checked for compatibility. Files and metadata arrive via standardised interfaces; OpenCloud stores them correctly itself. This significantly reduces the technical complexity on the destination side.

Data is transferred via open, documented interfaces. WebDAV is used for files and directory structures; this is a standard that has been established for decades and is supported by virtually every file-sharing system. The LibreGraph API handles structures and permissions. The REST API, based on the OpenAPI 3 standard, is modelled on the Microsoft Graph API and is compatible with it. As OpenCloud supports these protocols, the file-sharing platform can be integrated directly.

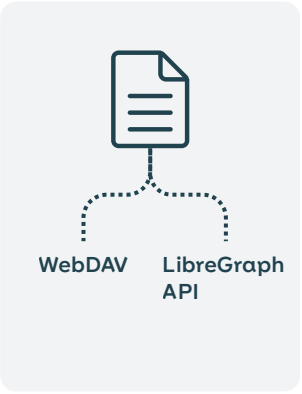

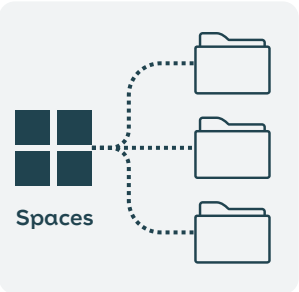

One key difference from other platforms remains: OpenCloud organises shared content into Spaces – self-contained workspaces with their own permissions, responsibilities and storage quotas. Anything that exists in the source system as a group folder or shared directory must be mapped to Spaces in a meaningful way during the planning phase. This step requires some planning, but is not a technical hurdle. The result is a cleaner, more manageable structure than before.

The migration process in four stages

To avoid any surprises during operation, the migration follows a tried-and-tested process: analysis, planning, testing and migration. Each of the four phases builds on the previous one.

1 Analysis: Fully capture the source platform's users, data and permissions

The first step is a detailed assessment. How many users are there, how many groups, and what is the amount of data? What usage patterns can be identified – are there any particularly large accounts, unusual file structures or external shares? Which desktop clients are in use, how many end devices are actively synchronising, and which other systems are connected via interfaces?

A structured questionnaire captures the key details and serves as the basis for initial planning conversations. It is not just about the figures. It is also important to understand whether there are any technical or contractual constraints that might affect the migration path, and what specific usage patterns exist – such as hidden tags, custom WebDAV properties or integrations with third-party systems.

It is also important to clarify at an early stage how user management works on the destination site: do the source and destination platforms access the same identity provider, or do user accounts and groups need to be created explicitly in OpenCloud?

Equally important is an understanding of the technical infrastructure on both sides: where are the source and target systems located, what is the network connection between them and the migration server, and what technical limitations – CPU, RAM, storage capacity – need to be taken into account?

2 Planning: Defining the process, scope and timeframe of the migration

Based on this analysis, a migration plan is drawn up. The key task is to map source to destination: which user accounts will be assigned to which locations in OpenCloud, and which group folders will be mapped to which spaces? At the same time, the parties involved determine exactly which data is to be migrated – whether deleted files are included and whether the full version history is transferred has a significant impact on the effort involved and the volume of data to be transferred.

Particular attention should be paid to how shared content is handled: should existing shares and group folders be mapped immediately as Spaces in OpenCloud, or should this be done gradually after the actual data migration? The latter is generally advisable. Anyone who opens OpenCloud after the migration and cannot find their files straight away because they have already been moved to Spaces will experience an unnecessarily bumpy start. A phased transition gives users time to familiarise themselves with the new concept.

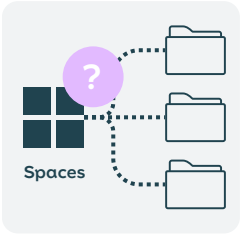

It is also important to clarify what cannot be migrated. Typical examples include comments, certain advanced user settings, or data that is not available via server APIs. The sooner this is established, the better expectations can be managed.

Before the actual transfer of data gets started, user accounts and groups must be set up in OpenCloud. This provisioning is an essential prerequisite – without the designated structure, data and permissions cannot be assigned correctly. Depending on the approach chosen, Spaces can either be set up in advance or created step by step after the migration.

The schedule should include sufficient buffer for unexpected situations. Key reference points are the availability of the OpenCloud target environment and the migration infrastructure.

Rather than migrating all users at once, it is advisable to migrate in groups, known as batches. These are organised by department or usage patterns and provide checkpoints. This helps identify problems at an early stage, before they affect the entire organisation.

If users are allowed to continue working on the source system during the migration, new and modified files can be synchronised in a second pass.

In general, there are three scenarios:

- Users should not make any changes to the source system during the migration.
- Users create new files or update existing ones – in which case a second run is required.
- Users move and delete files, which would require a full synchronisation and is not currently recommended.

In practice, scenario 2 is the norm for large-scale installations. Copying very large amounts of data can take several days – and operations cannot be interrupted for that long.

Recommended approach: First, all existing data is copied whilst the system is running. This is followed by a second run during a short maintenance window to process new or modified files. This keeps downtime to a minimum.
Until the official switchover, users should not have write access to the destination account.

Planning also involves communicating with users at an early stage. By informing users in good time about what is changing and what is expected of them, you can reduce the workload on support staff and avoid mistakes.

3 Test: Verify functionality and data integrity prior to going live

Prior to the actual migration, a multi-stage testing process takes place. Optionally, a preliminary test can be carried out even before the target platform is complete – for example, to migrate test users to a temporary environment. This step provides valuable insights, but cannot replicate all the complexities of the final migration.

During the functional test, the migration team transfers a representative selection of accounts that cover the full range of the source environment: large accounts, complex permission structures, and special cases such as file names containing special characters or particularly deep directory structures. Among other things, the functional test checks: Has all data been transferred correctly? Has the folder structure been preserved? Are corrupted files in the source system correctly flagged as errors? Has no data been altered in the source system? Is data ending up in the correct destination account?

During the pilot phase, the team is expanding the test to include a full trial run using real data. The aim is to demonstrate that the environment is ready for production. Only then will the live test begin with selected users who are willing to be the first to use the new environment and provide feedback. If necessary, it is possible to switch back to the source system. Following this initial step, there will be a pause for evaluation before the migration continues.

4 Migration: Users and data will be monitored and gradually transferred to OpenCloud

Once the analysis, planning and testing have been completed, the actual migration to OpenCloud begins – the source system continues to operate as normal for all users who have not yet been migrated. The migration team transfers files, folder structures and permissions, whilst continuously monitoring throughput, error rates and system utilisation on both sides.

The first groups are being monitored particularly closely. The volume of data, the number of files, the performance of the source and target systems, and the connection between them determine how long each step takes. These findings are incorporated directly into further planning: How large can the next groups be? What timeframes are realistic? What failure scenarios need to be accounted for?

The user cutover – the point from which a user works exclusively in OpenCloud – only takes place once all data that was present in the source system at the start of the migration has been transferred completely and correctly. Errors reported by users are first reviewed by the internal support team: not every issue is related to the migration. Anything that is actually caused by the switch is then passed on specifically to the migration team.

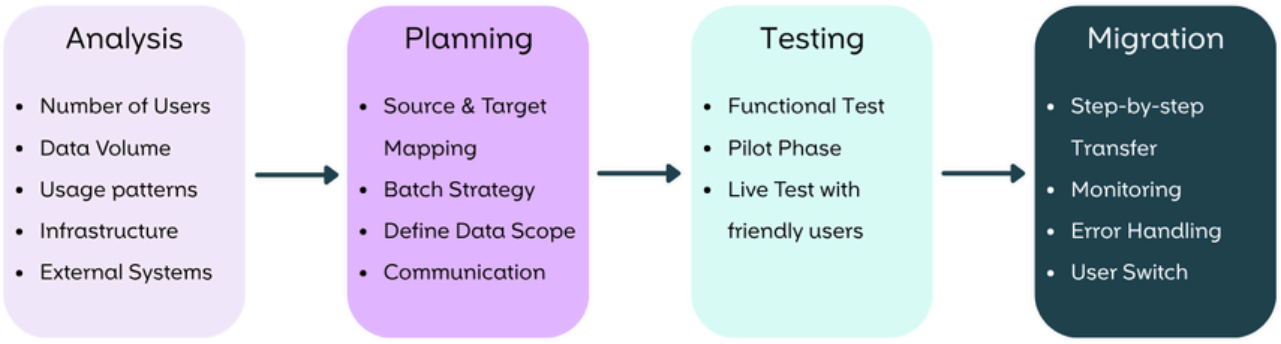

Challenges during migration: What to keep in mind

No migration is ever completely free of challenges. What matters is how well a platform handles them. OpenCloud is designed so that typical stumbling blocks either don't arise in the first place or can be resolved quickly.

- **Re-configuring desktop clients:** Users who synchronise files locally must switch their desktop client to OpenCloud following the migration. Before doing so, it must be clarified whether users should resynchronise all files after the migration or whether the new client should recognise existing local data – and whether there is sufficient local storage space available for this.

- **Adapting legacy configurations:** Clients and third-party systems configured to use old hostnames and paths must be updated to use the new endpoints following the migration. This applies not only to desktop clients, but also to all automated processes and integrations that access the source platform.

- **Active users during migration:** Anyone who continues to use the source system during migration will need a second run – this will transfer all changes made since the first run. The effort involved is usually manageable.

Professional support: OpenCloud and audriga

Whether a migration goes smoothly depends not least on how well it is supported. OpenCloud provides direct support – from the initial consultation right through to live operation.

For large-scale projects, audriga is a partner specialising in data migration. audriga analyses existing structures, oversees the data transfer and helps adapt work processes to the new platform. This includes specialised tools for analysing the source environment and transferring data.

What organisations gain after migration

Less maintenance, lower energy consumption, no proprietary dependencies: switching to OpenCloud pays off immediately. OpenCloud runs without a relational database, without interpreter overhead and without licensing risks. Backups are reduced to a simple file system snapshot, updates run without any downtime, and performance remains predictably stable even as user numbers grow.

What remains is an infrastructure that grows in line with requirements – open, efficient and entirely under your own control. Organisations decide for themselves where their data is stored, who has access to it and how the system is developed. That is digital sovereignty in practice.

Would you like to migrate your file-sharing platform to OpenCloud? We can help you with the planning and implementation.

Get in contact with us at sales@opencloud.eu.
We look forward to hearing from you.

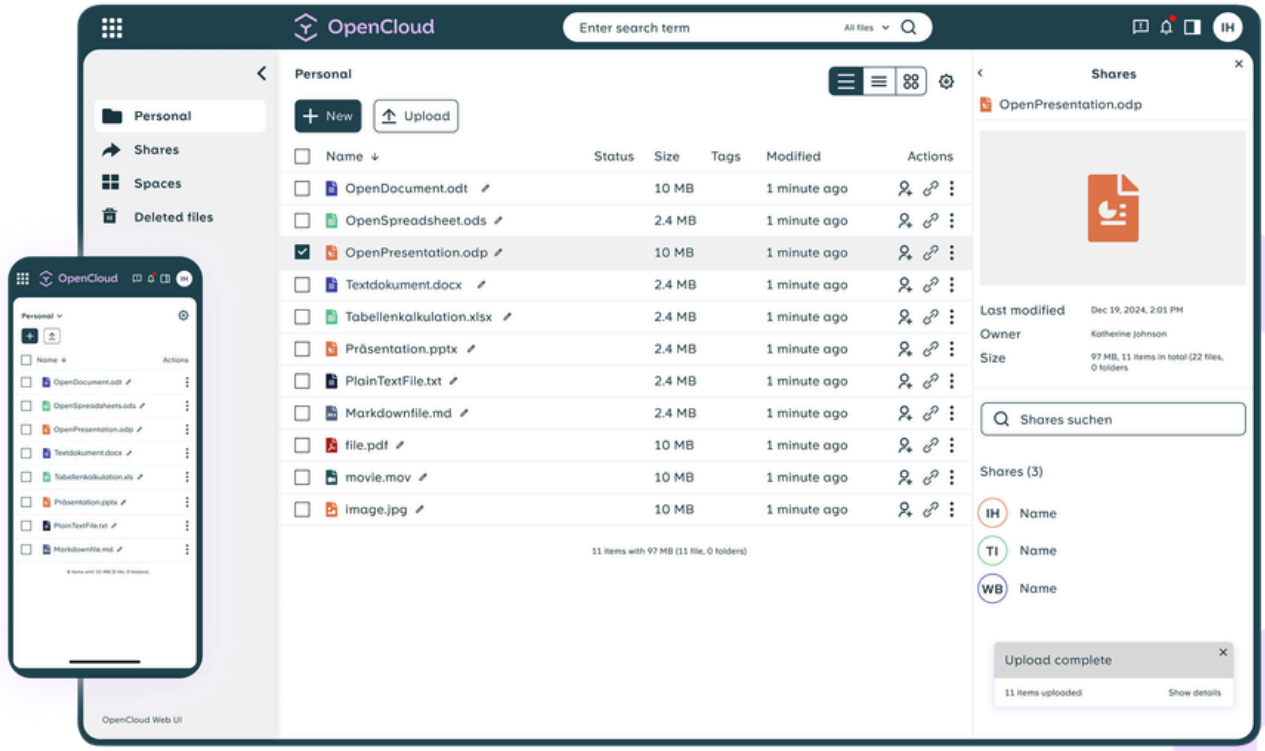

OpenCloud

OpenCloud is the powerful open source solution for file management, file sharing and content collaboration that meets the highest standards of user-friendliness, performance, stability and scalability. The software can be operated flexibly on-premise or securely procured via partners. With intuitive operation, seamless integration, person-independent file management ("Spaces") and secure versioning, OpenCloud is ideal for the public sector, educational institutions, research and large companies and providers.

Together with OpenTalk and mailbox, OpenCloud is one of the three strong pillars of the Heinlein Group, which provides secure and digitally sovereign IT infrastructures for Europe.

https://opencloud.eu