

# Designing a Laserfiche 8 Backup and Recovery Plan

*White Paper*

April 2011

**Laserfiche®**

## Table of Contents

Summary .....	4
Core Repository Architecture .....	5
Repository Database .....	5
Expansion Databases .....	6
Volumes .....	6
Repository Directory .....	7
Search Catalog .....	8
Licensing Information .....	9
Laserfiche Server License Database .....	9
License Manager Database .....	9
Backup Best Practices .....	11
Backup Types .....	11
Full Backup .....	11
Differential Backup .....	11
Incremental Backup .....	12
Backup Type Summary .....	12
Backup Storage .....	12
Backup Schedule .....	14
Backup Medium .....	14
Syncing Backups .....	15
Database Backup .....	17
Search Catalog Backup .....	18
Volumes and Repository Directory Backup .....	19
Recovery Simulation .....	20
Schedule .....	20
Environment .....	20
Simulation Process .....	21
Other Laserfiche Products .....	24

Laserfiche Audit Trail.....	24
Audit Logs.....	24
Audit Database .....	26
Saved Reports.....	26
Configuration File.....	26
Recovery Tasks.....	27
Laserfiche Workflow .....	27
Workflow Databases.....	27
Configuration Files .....	28
Custom Activities .....	28
Recovery Tasks.....	28
Laserfiche Quick Fields.....	29
Session Files.....	29
Quick Fields Agent Data Folder .....	29
Recovery Tasks.....	30
Import Agent .....	30
Recovery Tasks.....	31
Agenda Manager .....	31
Agenda Manager Database .....	31
Volume .....	32
Recovery Tasks.....	32
Web Access .....	32
Recovery Tasks.....	32
WebLink.....	33
Configuration Folder .....	33
WebLink Customizations.....	33
WebLink 8 Reporting Information.....	33
Recovery Tasks.....	34

# Summary

Designing, implementing, and maintaining a dependable backup and recovery plan is a crucial administrative task for your Laserfiche system. This paper will provide important and detailed information about this effort. It is written for Laserfiche administrators who have been tasked with backup/recovery responsibilities and assumes some administrative Laserfiche knowledge on the reader's part.

# Core Repository Architecture

Before an administrator can think about *how* to back up Laserfiche data, he/she must have a firm understanding of *what* needs to be backed up and *why*. A leading cause of administrators failing to create reliable Laserfiche backups is not knowing, or forgetting, to back up a crucial component.

Keep the following points in mind when reading this section:

- The components described in this section of the paper comprise the core repository architecture. Each item contains data that should be backed up.
- The majority of this paper, including this section, will only cover backup/recovery concepts for core repository components, not components for other Laserfiche products (e.g., Workflow, Quick Fields, Agenda Manager, etc.). Designing a backup/recovery plan for the other components will be covered in the [Other Laserfiche Products](#) section.

---

**Important:** If you are auditing user actions, which is technically performed by the Laserfiche Server, ensure you read the [Laserfiche Audit Trail](#) section.

---

- While the information in this paper explicitly applies to Laserfiche 8.2, it may apply to other versions as well.
- If you have multiple repositories or Laserfiche product installations, ensure you back up all of the necessary components for *each* repository.

## Repository Database

Each repository has a Microsoft SQL Server or Oracle database that stores a significant amount of important information (e.g., metadata, security, repository structure, trustee settings, etc.). This database must be backed up.

---

**Note:** In many cases, the backup/recovery administrator is different than the administrator who originally installed and set up Laserfiche. As a result, this paper will explain how to determine the location of each component that must be backed up.

**Note:** This paper will only cover Microsoft SQL Server database concepts; Oracle will not be discussed. While some of the minor details will be specific to SQL Server, most of the concepts still apply to Oracle.

---

### To find your database:

1. Open the Laserfiche Administration Console.

2. Log in to a repository as a user with the **Manage Repository Configuration** privilege.
3. Under the repository's main node, expand **Repository Options**.
4. Right-click the **Settings** node and select **Properties**.
5. Select the **Attributes** tab to view the database's properties:
  - The type of database (e.g., SQL Server or Oracle) will be listed under **Driver**.
  - The name of the machine hosting the database will be listed under **Data source**.
  - The name of the database will be listed under **Database name**.

## Expansion Databases

If you used a personal, desktop, or express edition of SQL Server in a previous version of Laserfiche, and, after upgrading to Laserfiche 8, are still using one of these editions, your repository may be associated with *multiple* databases.

In this situation, the repository will still have only one *main* database (i.e., the database described above), but will also have one or more *expansion* databases (where only thumbnails and word location data are stored). Expansion databases are on the same database server as the main database.

After upgrading to Laserfiche 8, you must back up the expansion databases once. However, after that, it is not necessary to regularly back them up, as they will never be modified again (all future thumbnails and word location data will be stored in volumes). The main database, which is listed under **Database name** in the **Attributes** tab described above, should be backed up on a regular basis.

## Volumes



Volumes, which are stored in the Windows file structure, store a significant amount of important repository information (e.g., scanned images, OCR'd text, electronic files, etc.). A repository can have as many volumes as necessary, and they can be spread out across different machines. All volumes must be backed up.

---

**Note:** It is important that backup/recovery administrators understand Laserfiche volume basics, especially the differences between physical and logical volumes. For more information, watch [Overview of Volumes Part 1](#) and [Part 2](#) (videos) on the Laserfiche Support Site.

---

## To find your volumes:

1. Open the Administration Console.
2. Log in to a repository.
3. Under the repository's main node, expand **Volumes** to see all of the repository's volumes. Keep the following mind:
  - Logical volumes (which appear next to  icons) periodically roll over and create new subvolumes, known as physical volumes (which appear next to  icons).
  - Expand a logical volume to view its physical volumes. Ensure you back up all physical volumes that currently exist inside a logical volume, plus those that may be automatically created during rollover in the future.

---

**Tip:** Unless otherwise configured, all physical volumes inside a logical volume share the same root directory. The easiest way to back up all of a logical volume's current and future physical volumes is to back up this directory. To find the directory, right-click the logical volume, select **Properties**, and choose the **Path** tab. Depending on the types of paths that are enabled, back up the fixed path, removable path, or both.

---

- While physical volumes can exist inside a logical volume, as explained above, they do not have to. Ensure all of a repository's physical volumes are backed up, whether they exist inside a logical volume or not.
- To determine where a physical volume is stored, right click it, select **Properties**, and choose the **Path** tab. Back up the fixed path, removable path, or both.
- Volumes paths listed in the Administration Console are relative to the Laserfiche Server machine. In addition, fully relative paths are relative to the repository directory.

## Repository Directory

A repository directory, which is stored in the Windows file structure, contains data the Laserfiche Server needs to communicate with a repository.

This directory contains:

- Temporary files
- Log files
- \*Volumes

- \*Audit logs
- \*Search catalog

The above \* items can be moved out of the repository directory if necessary, though they exist here by default. You should back up the repository directory even if all of the \* items are moved to somewhere else.

#### To find your repository directory:

1. Open the Administration Console.
2. Log in to a repository as a user with the **Manage Repository Configuration** privilege.
3. Under the repository's main node, expand **Repository Options**.
4. Right-click the **Settings** node and select **Properties**.
5. Select the **Attributes** tab. The location of the repository directory is listed under **Location**.

---

**Note:** The path listed is relative to the Laserfiche Server machine.

---

## Search Catalog

Each repository has a search catalog, which is a series of files used by the Laserfiche search/index engine to perform full-text searches (i.e., searching the repository for all documents whose text contains a specific word or phrase). Though we highly recommend backing up this catalog, it can be regenerated if necessary (this process can take hours or even days depending on the number of documents in the repository).

---

**Note:** If you do have to regenerate your search catalog, you may want to use the Laserfiche 8 Quick Reindex Utility, which can be significantly faster than choosing to reindex from the Laserfiche 8 Administration Console. For more information, [see the Administration Console help file topic on this feature](#).

---

#### To find your search catalog:

1. Open the Administration Console.
2. Log in to a repository as a user with the **Configure Search/Index** privilege.
3. Under the repository's main node, right-click the **Index** node and select **Properties**.
4. The location of the search catalog is listed next to **Path**.



## Licensing Information

Though it is always possible to create new license files, we recommend backing them up, as doing so will decrease recovery time. In almost all cases, a product's license is named **LF.lic** or **LF.licx** and is stored in its installation directory (e.g., the Laserfiche Server's license is usually stored in **C:\Program Files\Laserfiche\Server**).

---

**Note:** If a license file cannot be found in this location, the product may not require a license to run (e.g., the Laserfiche Client or Laserfiche Administration Console).

---

In addition, if your Laserfiche Server uses an Avante or Rio license, ensure you backup the relevant licensing databases.

### Laserfiche Server License Database

For Avante and Rio systems, backup each Laserfiche Server's SQLite database file that contains the Server's licensing information.

#### To find your Laserfiche Server license database:

- This file is always named **lfsnu.db** and is always stored in the Server's installation directory.

### License Manager Database

For Rio systems only, each License Manager Server has at least one License Manager Database that contains system-wide licensing information (not specific to any particular Laserfiche Server or repository). Each License Manager Database is a SQLite .db file and should be backed up.

---

**Note:** Rio users should back up the **lfsnu.db** file for each of their individual Laserfiche Servers, plus their License Manager Databases.

**Note:** In most cases, Rio administrators will have only one License Manager Database to back up. If you have more than one, back up each.

---

You can back up a License Manager Database using the License Manager Administration Console, even if your License Manager server is still running. This backup will preserve information about your installation, including your master license, your list of application instances, and your named user and device lists.

---

**Tip:** You can automate the process of backing up the License Manager Database using the License Manager API, known as License Manager Objects (LMO). More information on LMO can be found in the License Manager Administration Console help files.

---

**To manually back up your License Manager Database:**

1. Open the License Manager Administration Console.
2. Expand your License Server.
3. Right click your License Manager Database and select **Create Backup**.
4. Specify where to save the database file and what it should be named (e.g, C:\Backups\MyLicenseDatabase.db).
5. Click **OK**.

# Backup Best Practices

Once a Laserfiche administrator knows *what* needs to be backed up and *why*, he/she can consider how the backup should be performed. In this section, we will cover important backup/recovery best practices.

## Backup Types

There are three types of backups that are commonly used: full, differential, and incremental. These may be used to back up any Laserfiche component, and most organizations utilize multiple types simultaneously.

### Full Backup

This requires 100% of the data to be backed up each time, which requires a lot of storage space and total processing time. Restoring from a full backup alone, however, is fast and simple, as there is a single backup container to reinstate.

Most organizations perform full backups less regularly than any other backup type, and they most often use full backups in combination with differential or incremental backups.

### Differential Backup

This only backs up the data that has changed between now and the last *full* backup. For example:

- **Sunday:** A full backup of a volume is performed.
- **Monday:** Only the data in the volume that has changed between Sunday and Monday is backed up.
- **Tuesday:** Only the data in the volume that has changed between Sunday and Tuesday is backed up. And so on.

If differential backups are utilized, two separate backup containers are required to restore all data: the most recent full backup is restored, on top of which the most recent differential backup is restored.

---

**Example:** The Red Company performs a full backup once a week (on Sunday at 10 p.m.) and a daily differential backup (at 10 p.m. every day, except Sunday). If all production data is lost at 3 p.m. on Thursday, the organization would first restore the last Sunday's full backup, on top of which they would restore Wednesday's differential backup.

**Note:** Each differential backup is a separate backup container that is dependent upon a full backup. No differential backup is dependent on any other differential backup. Therefore, it is up to

administrators to determine how many differential backups to keep on hand at any given time (e.g., an administrator could choose for today's differential backup to overwrite yesterday's, or he/she could choose to retain both).

---

## Incremental Backup

This only backs up the data that has changed between now and the last backup *of any type* (full, incremental, or differential). For example:

- **Sunday:** A full backup of the repository database is performed.
- **Monday:** Only the database content that has changed between Sunday and Monday is backed up.
- **Tuesday:** Only the database content that has changed between Monday and Tuesday is backed up. And so on.

If incremental backups are utilized, the restore process is more complex, as it involves multiple backup containers: the most recent full backup is restored, on top of which all subsequent incremental backups are restored (in the order in which they were created).

---

**Example:** The Blue Company performs a full backup once a week (on Sunday at 10 p.m.) and an incremental backup every six hours (at 12 a.m., 6 a.m., 12 p.m., and 6 p.m.). If all production data is lost on Friday at 1 p.m., the organization would restore the following backup containers (in this order): last Sunday's full backup, then all incremental backups performed between last Sunday at 10 p.m. and Friday at 1 p.m. (19 in all).

---

## Backup Type Summary

When designing your backup/recovery plan, utilize the backup types that best fit your situation. Below is a summary of the benefits of each:

- **Full backup (alone):** Fastest and easiest to restore from, but requires the most storage space and time to backup.
- **Differential (with full):** Middle ground between Full and Incremental backups.
- **Incremental (with full):** Fastest to backup and requires the least amount of storage space, but also takes the longest amount of time to restore from.

## Backup Storage

At the least, each backup you create (whether it's full, differential, or incremental) should be kept on three different media simultaneously. In other words, you should maintain redundant backup media, and not rely on just

one. A minimum of three is suggested because you should always assume one backup is corrupt and that another will become corrupt during the backup process, leaving you with one to recover from. If possible, we also suggest keeping the backups in different physical locations, to decrease the chances of all three being destroyed due to a disaster.

---

**Example:** If you create a full backup of a volume, store it on **Backup Disk 1, Backup Disk 2, and Backup Disk 3**. Later, if you also create a differential or incremental backup of the volume, also store it on the same three backup disks.

**Tip:** Ensure you document the location of each backup container and that all relevant administrators have access to this information. During a disaster, when a recovery plan is put into effect, often times organizations spend a significant amount of time determining where backups are located.

---

You should also determine how long it is necessary to keep each backup before deleting it or overwriting it with a new backup. The more backups you keep, the more history you have on hand, but the higher your storage disk requirements are.

---

**Tip:** In addition to regularly performing full backups and eventually deleting and overwriting them (e.g., once a week), some organizations perform special full backups intermittently that are permanently archived (e.g., once every six months). This gives you a permanent snapshot of all your data at a given time and can be invaluable if, for example, you realize many months after the fact that a user accidentally deleted a large folder of important data.

---

An offsite backup plan—which usually includes scheduled tape or DVD backups, after which the backup medium is stored offsite—are highly recommended for all Laserfiche deployments. Without offsite backup, all of your data can be destroyed by a campus-wide disaster, such as a fire or flood. Third-party data backup services, such as Iron Mountain Incorporated, offer offsite backup solutions, which in many cases include backing up to different geographic locations, ensuring your data is safe from regional disasters.

---

**Tip:** For mission-critical deployments of Laserfiche, protection of data alone is not sufficient, as a speedy return to service availability is also required. After a disaster occurs, waiting to stand up a new system, then restore from backup, may take longer than the organization can afford. It may be necessary to invest in a remote failover location to bring online almost immediately after a disaster.

---

## Backup Schedule

When deciding how often to backup, the key questions to ask are:

- How busy is the system (i.e., how often is data added, deleted, or modified)?

---

**Tip:** When answering this question, take into account both human interaction with the repository and automated interaction from software. For example, if you have Laserfiche Workflow, Quick Fields, or Import Agent, your repository may be automatically modified very regularly.

---

- How much work can you reasonably afford to lose?

---

**Example:** Assume an organization scans documents into a repository from 8 a.m. to 5 p.m. every day, and that it performs incremental backups once a day at 7 p.m. If water accidentally leaks onto the production disk and destroys this data at 6 p.m., an entire days worth of scanning is lost, since today's backup hasn't occurred yet. Can the organization afford to either lose all of today's data or spend another day rescanning everything? If yes, then this backup plan is sufficient. If not, then the organization should back up more regularly than once a day (e.g., every 6 hours).

---

At the least, we recommend performing a daily differential or incremental backup and a weekly full backup. While there may be some exceptions to this general guideline (e.g., Laserfiche systems that see very little traffic), in most cases it is the bare minimum in terms of a sufficient backup schedule.

---

**Example:** At Laserfiche, our internal Laserfiche system receives a differential backup every six hours and a full backup once per day.

---

## Backup Medium

Most organizations back up to hard disk or magnetic tape. Each medium offers its own unique advantages and disadvantages, which are, for the most part, beyond the scope of this paper. We recommend researching the two media to determine which is appropriate for you.

That said, below is a very simplified summary of the traditionally accepted advantages and disadvantages of each medium:

- In most cases, disk is faster, to both back up to and recover from. It has traditionally been more expensive, though the cost-difference between the two is quickly shrinking.
- Tape is more secure and longer lasting.

---

**Tip:** In January of 2011, the Spectra Logic Corporation, which sells data protection and backup solutions that use both disk and tape, released a white paper titled **Tape or Disk: Why Not Both?** The paper provides a detailed overview of the advantages and disadvantages of both media. While Laserfiche makes no express warranties about the paper or Spectra products, it may be a good place to start your research. The paper can be found on Spectra Logic's Web site or via an Internet search engine (e.g., perform a Google search for "Tape or Disk: Why Not Both? spectra logic").

---

Keep in mind that many organizations use both disk and tape, employing a method known as **disk-to-disk-to-tape (D2D2T)**, which attempts to capitalize on each medium's strengths and minimize their weaknesses. This approach initially backs up data on the production disk to a backup disk, and then periodically copies data on the backup disk to tape. Since disk is initially used as the backup medium, the backup and recovery process (for recent data) is fast. However, since data on the backup disk is eventually copied to tape, the data is archived long-term on a more secure and longer lasting format.

## Syncing Backups

Backups can usually be scheduled outside of regular business hours, when no one is using the repository. If this is the case, and you can be sure that no users will ever modify the repository during the back up of any Laserfiche component, this section of the paper does not apply to you.

However, if this section does apply to you (e.g., if you need to back up in the middle of the day or your repository is used 24 hours a day), ensure that you prevent changes from being made to one Laserfiche component while another is being backed up, which would cause the two components to be out of sync. If this occurs, one component may recognize data that another does not, which can lead to problems.

---

**Example:** An organization schedules a full backup of the repository's only volume, immediately after which they schedule a backup of the repository's database. While the volume is being backed up, a user logs into the repository and adds 50 documents. Due to the timing of the changes, these documents are not included in the volume backup. When the database backup is performed, however, these documents are included. Therefore, the volume backup does not know about the documents, but the database backup does. If the organization restores from these backups, the documents will be listed in the Laserfiche Client (because the database thinks they exist), but, when a user attempts to open them, an error will occur, because they don't actually exist in the volume.

---

To prevent this issue, freeze your repository before starting any Laserfiche backups, then unfreeze it after all backups are completed. There are number of ways to do this:

- If you can temporarily provide zero-access to the repository while the backup is taking place, stop the Laserfiche Server, perform the backups, and then restart the Server.

---

**Tip:** This process can be easily automated. In many cases, third-party backup solutions can either directly start/stop Windows services or run a batch file before/after a backup. Alternatively, you can create a Windows Task Scheduler job that runs a batch file.

**Tip:** To create a batch file that stops the Laserfiche Server, create a text file, change its extension to .bat, and enter the following text into it: **sc**  
**\\LaserficheServerMachineName.domain.com stop LFS**  
(alternatively, change **stop** to **start** to start the Server). For more information on batch files, or on Scheduled Tasks, search [Microsoft.com](http://Microsoft.com).

---

- If you must at least provide read-only repository access during backups, do one of the following:
  - Set the read-only flag on all volumes, perform the backups, and then remove the flags.

---

**Tip:** This can be automated by creating a Laserfiche SDK program that is run before/after a backup (the program could be triggered directly by a third-party backup solution or via a Windows Task Scheduler job). If you create such a program, build in logic that ensures it doesn't remove the read-only flag on volumes that should *always* be read-only, regardless of backups. For more information, see the Laserfiche SDK documentation.

---

- Create (and maintain) a Laserfiche group that contains all repository trustees. Before starting any Laserfiche backups, set this group to be read-only, perform the backups, and then remove the read-only flag.

---

**Tip:** The process of setting and clearing the group's read-only flag could also be automated via a Laserfiche SDK program. Creating and maintaining the Laserfiche group (e.g., adding new users to the group) should be done manually.

**Note:** Recent versions of Laserfiche do not allow you to set the Everyone group to read-only. This is why it may



be necessary to create a custom group that contains all trustees.

---

## Database Backup

Historically, the most likely point of failure for a Laserfiche backup is the database. In these cases, the backup that is created is invalid or corrupt due improper methods used to create it.

There are two safe ways of backing up a SQL Server database:

- **Hot backup (database stays online):** Execute the **BACKUP** Transact-SQL command. During the backup, the database can still be accessed and modified, but it may become slow. Many third-party database backup solutions use this backup method.
- **Cold backup (database goes offline):** Detach the database from the SQL Server and back up the primary data file (.mdf), the transaction log file (.ldf), and, if it exists, the secondary data file (.ldf). The database will not be available to read-from or write-to during the backup and must be re-attached to the SQL Server once the backup is finished.

---

**Note:** There are two safe ways to detach a SQL Server database: using the **sp\_detach\_db** Transact-SQL command and using SQL Server Management Studio.

**Important:** We highly recommend detaching the database you want to back up instead of stopping the SQL Server that hosts it. If you stop the Server in an improper manner, the database files you back up may be invalid or corrupt. This is a common reason for failed Laserfiche backups.

**Important:** Before using a third-party backup solution, ensure it both supports database backups and that it uses one of the methods explained above. If the backup solution does not know it is backing up a database (e.g., if you point it directly at a database's data and log files without taking the database server into consideration), it will not take the necessary steps to perform a *safe* backup. This is another common reason for failed Laserfiche backups.

---

For both ease of use and dependability, we recommend performing hot backups using any of following methods:

- **SQL Server Maintenance Plan:** A feature in SQL Server Management Studio that enables users to create a workflow of tasks that ensures a database is regularly optimized and backed up. Historically, “express” versions of SQL Server do not include this feature. For more information, search Microsoft.com for “maintenance plan.”

- **Symantec’s Backup Exec:** A popular backup solution that includes a database option. For more information, query an Internet search engine for “Symantec Backup Exec.”
- **Stored Procedure + Batch File + Windows Task Scheduler (free):** This method involves creating a stored procedure in SQL Server’s master database, creating a batch file to trigger it, and then tying a Windows Task Scheduler job to the batch file. For more information, see Microsoft KB 2019698, or search Microsoft.com for “how to schedule and automate backups of SQL Server databases in SQL Server Express editions.”

## Search Catalog Backup

Due to the architecture of the Laserfiche search/index engine (i.e., Laserfiche Full-Text Indexing and Search Service), the search catalog backup process is more involved than simply copying the catalog from the production disk to a backup disk. If one of the methods below is not used, the backup version of the catalog may be invalid or corrupt.

- **Automatic:** Use a third-party solution (e.g., Symantec’s Backup Exec) that supports Microsoft’s Volume Snapshot Service (VSS) technology, which is required to back up a *live* search catalog. Ensure the Laserfiche VSS Writer service is running on the machine hosting the search/index engine (it is stopped by default) and point the backup program at the search catalog with VSS-enabled. The Laserfiche VSS Writer service will coordinate the process of stopping all catalog actions (full-text searching and indexing) while the backup process takes places, as well as resuming catalog actions after backup is complete.

---

**Note:** There is no way to have full-text searching and indexing remain functional while a search catalog backup is being performed.

---

- **Manual:** Stop the search/index engine, copy the search catalog to a backup disk, and restart the service.

---

**Tip:** To create a batch file that stops the search/index engine, create a text file, change its extension to .bat, and enter the following text into it: **sc \v-dev-xp86-10.laserfiche.com stop LfFTSrv** (alternatively, change **stop** to **start** to start the Server).

---

## Volumes and Repository Directory Backup

Volume and repository directory files can be simply copied from the production disk to a backup disk. We recommend automating this process using a third-party backup solution (e.g., Symantec's Backup Exec).

---

**Important:** By default, the repository's search catalog is stored inside the repository directory (though it can be moved out). If this is the case, in order to perform a safe backup of the catalog, you may need to exclude it from the repository directory backup and perform a separate search catalog backup. This is also true for any database files that might be stored in the repository directory.

---

# Recovery Simulation

Once a backup/recovery plan is designed and implemented, it should be tested on a regular basis. By simulating a situation where all production data is lost, and then attempting to recover solely from backup data, you can determine if your backup/recovery plan works.

## Schedule

While recovery simulations should be regularly scheduled, it is especially important to test your backup/recovery plan after significant changes are made to your production or backup environments. For example:

- Hardware or software is upgraded.
- Hardware is physically moved from one location to another.
- New applications are added.
- Major changes are made to existing software (e.g., a new repository is added to a Laserfiche Server).

When events like these occur, revisit and test your backup/recovery plan to ensure all important data is being backed up.

## Environment

You will need a test environment in which to perform a recovery simulation. The process of creating one depends on your Laserfiche licensing model.

- **Rio:** You have the ability to install and use an unlimited number of Laserfiche Servers, which allows you to easily create a test environment.
- **Avante/Team/United:** You are limited to the number of Laserfiche Servers you have purchased. Choose from one of the following options:
  - **If downtime is permitted:** Perform recovery simulations on the production environment.

---

**Important:** Be careful not to overwrite or modify any production data, and ensure you schedule the test at a time when the production environment is not needed by everyday users.

---

- **If downtime is not permitted:** Request temporary licenses from your Laserfiche Value Added Reseller (VAR) or your Laserfiche sales manager to create a temporary test environment.

## Simulation Process

The goal of the simulation is to attempt to recover all repository data from backup disks alone.

---

**Note:** The following simulation guide is specific to a particular Laserfiche repository. If you have multiple repositories, you should perform recovery simulations for each.

**Note:** If you have a Rio system, you should also perform recovery simulations for the entire system that takes into account back up versions of License Manager Databases. This process is beyond the scope of this paper.

---

### To perform a recovery simulation:

1. Copy the backup versions of each component to the relevant locations:

- **Repository database:** Copy to the machine hosting SQL Server.

---

**Note:** If the repository has any expansion databases, copy these to the same location.

---

- **Volumes:** After registering the repository (see Step 4 below), the Laserfiche Server will look for the repository's volumes in the same location it did on the production environment. In order to test the backup versions of the volumes, you will need to do one of the following:

- Copy the backup volumes to a location the Laserfiche Server can access them from, and then, after completing Step 2 below, modify the repository to look in this location (i.e., change each volume's path in the Administration Console).

---

**Note:** Volumes paths listed in the Administration Console are relative to the Laserfiche Server machine. In addition, fully relative paths are relative to the repository directory.

---

- Move the production volumes to an alternate temporary location, and then move the backup volumes to the location the Laserfiche Server expects them to be.

---

**Warning:** Be very careful not to accidentally overwrite the production volumes with the backup volumes. Also, ensure you move the production volumes back to the appropriate location after the simulation is over.

---

- **Repository directory:** Copy to the machine that will host the simulation environment's Laserfiche Server.
  - **Search catalog:** Copy to the machine that will host the Laserfiche search/index engine.
2. Install the Laserfiche Server, Client, and Administration Console on the simulation environment.
  3. Attach the backup version of the repository's database to a SQL Server.

---

**Note:** If the repository has any expansion databases, attach these as well.

**Note:** If you are using the same SQL Server for both production and for the backup simulation, attach the database(s) using different name(s), so as not to overwrite the production database(s).

---

4. Register the repository to the simulation environment's Laserfiche Server.
  - When prompted to define the database, point the wizard at the database attached to SQL Server in the previous step.
  - When prompted to define the repository path, point the wizard at the backup version of the repository directory.

---

**Note:** After registering the repository, you may receive a warning indicating the repository's search catalog does not exist. This message can be ignored.

**Tip:** Step-by-step instructions on registering a repository are provided in the Administration Console's help files.

---

5. Attach the backup version of the search catalog to the repository.

---

**Tip:** Step-by-step instructions on attaching a search catalog are provided in the Administration Console's help files. Search for "attaching or detaching a search/index catalog."

---

6. If the Laserfiche Server uses an Avante or Rio license, copy the backup version of the **lfsnu.db** file (which is the Laserfiche Server license database) into the Laserfiche Server's installation directory (in most cases, this will be **C:\Program Files\Laserfiche\Server**).
7. Restart the Laserfiche Server to ensure all components are available.
8. Log in to the repository using the Laserfiche Client and, at the least, test the following:
  - Open at least one document on each volume. If a document's content cannot be displayed, or you receive an error, the volume may not have been backed up or its path may be incorrect.

---

**Tip:** The **Within Volume** search allows you to search for documents by volume.

---

- Perform a full-text search for a word you know exists in at least one document. If no results are returned, or if an error is returned, the search catalog may not have been backed up (if so, the repository may be in the process of re-indexing), or the catalog may be invalid or corrupt.

# Other Laserfiche Products

In addition to backing up core repository components, you should also back up important data associated with other Laserfiche products. This section will list the data that should be backed up for each, as well as provide a summary of each product's recovery tasks.

Keep the following points in mind when reading this section:

- When discussing recovery tasks, we will assume the product in question must be completely reinstalled and that all of its production data and configuration settings were lost. We will also assume that, before attempting to recover the product in question, you have already recovered the relevant Laserfiche Servers and repositories.
- This list covers all major Laserfiche products and their most important data. That said, it is not exhaustive, as some minor settings and/or components may not be covered.
- Keep in mind that some of the concepts covered in previous sections of this paper may apply to backing up the products discussed below (e.g., how often to back up, backup media, recovery simulation, etc.).
- Though we recommend backing up all of the items listed for each product, not all are necessarily mandatory. Some items represent unique data (which cannot be recovered without a backup), while others represent configuration settings that determine how the product runs (which can be redefined, though doing so may be time-intensive and it may be difficult to exactly match the original settings).
- For more information on any of the steps below, search the product in question's help files.

## Laserfiche Audit Trail

The following Audit Trail data should be backed up: audit logs, an audit database, saved reports, and a configuration file.

---

**Note:** While the information below explicitly applies to Audit Trail 8.2, it may apply to other versions as well.

---

## Audit Logs

The Laserfiche Server saves auditing information to audit logs, which are the most important Audit Trail components to back up, as they contain all of your auditing information.

There are two types of logs: those managed by the Laserfiche Server and those that are not. Both should be backed up.



### To find audit logs managed by your Laserfiche Server:

1. Open the Administration Console.
2. Log in to a repository as a user with the **Manage Audit Settings** privilege.
3. Under the repository's main node, expand **Audit**.
4. Right-click the **Settings** node and select **Properties** to open the **Auditing Properties** dialog box.
5. In most cases, the Laserfiche Server will stop writing to a log file when it reaches a certain size or age, at which point it will roll over to a new file. In this situation, the easiest way to back up all log files (current and future) is to back up the folder listed under **Rollover directory**.


---

**Note:** If neither the maximum size nor the maximum age option is selected, rollover is not enabled. In this situation, it is only necessary to back up the log file listed under **Audit log path**.


**Note:** If rollover is enabled, but the **Audit log path** location does not exist within the **Rollover directory**, back up both.

---

### To find audit logs NOT managed by your Laserfiche Server:

1. Open the Audit Trail Web Reporter's configuration page by browsing to **http://<MachineName>/AuditTrail8Config/Configuration.aspx**, where **<MachineName>** should be replaced by the name of machine hosting the Audit Trail reporting service.
2. Select the **Repositories** tab.
3. Data sources listed next to  icons represent folders that contain audit logs not managed by a Laserfiche Server, all of which should be backed up (a machine name and a path is listed for each).

---

**Note:** Data sources listed next to  icons represent audit logs managed by a Laserfiche Server.

**Note:** The method to find unmanaged audit logs described above assumes an administrator has provided the Web Reporter with the correct location for each log. If the location is incorrect, this method will not work. In addition, note that the **Repositories** tab only lists audit log folders that an administrator has explicitly configured the Web Reporter to work with. If your organization has additional logs that are not currently in use, back up these as well. Bottom-line: While the above method may work in many cases, it is ultimately the administrator's job to know where each log file is and to ensure it is backed up.

---

## Audit Database

Using the Web Reporter and date ranges, administrators identify a subset of the data stored in audit logs, which represents the data that will be used to generate reports. The subset is copied from the logs into an audit database, which the Web Reporter queries to build reports.

Since the audit database contains the same information stored in the audit logs, it is not crucial to back it up. Nonetheless, we recommend that you do, as it will reduce your recovery time (i.e., you won't have to rebuild and repopulate the database).

### To find your audit database:

1. Open the Web Reporter's configuration page.
2. Select the **Database** tab. The name of the machine hosting the database server will be listed next to **Server Name** and the database name will be listed under **Audit Database Name**.

---

**Note:** If you are using the Web Reporter with multiple repositories, each may have a different audit database, all of which should be backed up.

---

## Saved Reports

After creating an Audit Trail report, you can save it. We recommend backing up all saved reports.

### To find your saved reports:

1. On the machine hosting the Audit Trail reporting service, browse to the Web Reporter's installation directory. In most cases, this will be **C:\Program Files\Laserfiche\Audit Trail 8**.
2. Browse into **WebAuditReporter** and then **App\_Data**.
3. Ensure the **report\_templates** folder is backed up.

## Configuration File

The settings defined in the Web Reporter's configuration page are saved to an .xml file, which should be backed up. Doing so will save you from having to set up the Web Reporter again during a recovery.

### To find your configuration file:

1. On the machine hosting the Audit Trail reporting service, browse to the Web Reporter's installation directory.
2. Configuration settings are saved to **Config.xml**, which should be backed up.

## Recovery Tasks

Follow the steps below to recover Audit Trail using backup data:

1. Attach the audit database to the database server.
2. Copy all audit logs to their original location.
3. Reinstall the Web Reporter.
4. Replace the Web Reporter's configuration file with the backed up version of this file.
5. Replace the Web Reporter's saved reports folder with the backed up version of this folder.
6. Restart the Audit Trail reporting service.

## Laserfiche Workflow

The following Workflow data should be backed up: Workflow databases, configuration files, and custom activities.

---

**Note:** While the information below explicitly applies to Workflow 8.0.1, it may apply to other versions as well.

---

## Workflow Databases

Both the Workflow Server and Subscriber have databases that contain important information (e.g., workflows, starting rules, workflow history). Both components can share the same database or they can each use separate databases.

### To find your Workflow databases:

1. Locate each component's database using the appropriate utility:
  - **To find the Workflow Server's database:** On the machine hosting the Workflow Server, click the Microsoft Windows **Start** button, **All Programs, Laserfiche, Workflow 8.0, and Workflow Server Configuration Utility**.
  - **To find the Workflow Subscriber's database:** On the machine hosting the Workflow Subscriber, click the Microsoft Windows **Start** button, **All Programs, Laserfiche, Workflow 8.0, and Workflow Subscriber Configuration Utility**.

---

**Note:** The steps below apply to both the Workflow Server and Subscriber. Complete these steps once for each component to determine where its database is stored.

---

2. Click **Next** until you get to the **Workflow Database** step.

3. The name of the machine hosting the database server will be listed next to **Choose SQL Server**. The name of database will be listed next to **Choose database**.

## Configuration Files

The Workflow Server and Subscriber each have a configuration file that should both be backed up. Doing so will save you from having to configure certain Workflow settings again during a recovery.

### To find your configuration files:

1. On the Workflow Server machine, browse to Workflow's installation directory. Ensure that **Laserfiche.Workflow.Service.exe.Config** is backed up.
2. On the Workflow Subscriber machine, browse to Workflow's installation directory. Ensure that **Laserfiche.Workflow.Subscriber.exe.Config** is backed up.

---

**Note:** The Workflow Server and Subscriber may be on the same machine, in which case both files are in the same folder.

**Tip:** In most cases, Workflow's installation directory will be **C:\Program Files\Laserfiche\Laserfiche Workflow**.

---

## Custom Activities

If you have added custom activities to Workflow, they should be backed up.

### To find custom activity files:

- On the Workflow Server machine, browse to **<Workflow Installation Drive>:\Program Files\Common Files\Laserfiche\Workflow\Activities**. All files in this folder should be backed up.

---

**Note:** This is the recommend location to store custom activity files. If you have opted to store them elsewhere, back up this location instead.

---

## Recovery Tasks

Follow the steps below to recover Workflow using backup data:

1. Attach the Workflow Server and Subscriber database to the database server.
2. Reinstall the Workflow Server, Subscriber, and all Designers.
3. Stop the Workflow Server and Subscriber.
4. Replace both the Workflow Server and Subscriber's configuration file with the backed up version of these files.

5. Deploy all custom activity files to the Workflow Server and all Workflow Designer machines. Skip this step if you don't have any custom activities.

---

**Note:** For more information on custom activity deployment, download [Creating Custom Laserfiche Workflow 8 Activities](#) from the Laserfiche Support Site (log in required), open **Building LFSO Workflow Activities (VB.NET)**, and read the deployment section.

---

6. Restart both the Workflow Server and Subscriber.

## Laserfiche Quick Fields

Quick Fields session files and the Quick Fields Agent data folder should be backed up.

---

**Note:** While the information below explicitly applies to Quick Fields 8.0.2, it may apply to other versions as well.

---

## Session Files

All sessions should be backed up by exporting them as .qex files, as this will save following data for each session:

- Configuration information that specifies how Quick Fields should handle documents.
- Personal settings, such as your toolbar layout, pane layout, and shortcut settings.
- Sample images.

---

**Note:** If you save a session, instead of exporting it, only its configuration information will be saved. Personal settings and sample images will not be retained.

**Note:** Ensure you create and back up a new .qex file each time the session is significantly modified.

---

### To export a session:

1. With a session open, select **Export Session** under **File** in the menu bar.
2. Browse to the location where you want to export the session.
3. Click **Save**.

## Quick Fields Agent Data Folder

If you are using Quick Fields Agent to schedule sessions to run unattended, we recommend backing up the Agent's data folder (which contains the schedules you've created and the Agent's history, among other data).

### To find Quick Fields Agent data folder:

1. In Windows, ensure hidden files and folders are being displayed.
2. On the machine hosting Quick Fields Agent, browse to **<System Drive>:\ProgramData\Laserfiche**.

---

**Note:** If this location does not exist, instead browse to **<Quick Fields Agent Installation Drive>:\Documents and Settings\All Users\Application Data\Laserfiche**.

---

3. Ensure the **Quick Fields Agent** folder is backed up.

### Recovery Tasks

Follow the steps below to recover Quick Fields using backup data (if you don't have Quick Fields Agent, skip the steps involving this component):

1. Reinstall Quick Fields and Quick Fields Agent.
2. Copy the backup version of all session files to the Quick Fields and Quick Fields Agent machine.
3. Import each .qex file into Quick Fields and save them as .qfx files.

---

**Note:** Ensure you save the sessions in the location where Quick Fields Agent expects them to be.

---

4. Replace Quick Field Agent's data folder with the backed version of this folder.
5. Restart the Quick Fields Agent service.

### Import Agent

All Import Agent profiles should be backed up.

---

**Note:** While the information below explicitly applies to Import Agent 8.1, it may apply to other versions as well.

---

### To find your profiles:

1. On the machine hosting Import Agent, open Windows' Registry Editor.

---

**Tip:** On most operating systems, this can be done by typing "regedit.exe" into Windows search or Run bar and pressing ENTER.

---

2. Ensure one of the following registry keys is backed up, depending on your operating system:
  - **32-bit machines:**  
HKEY\_LOCAL\_MACHINE\SOFTWARE\Laserfiche\Import Agent 8\Profiles

- **64-bit machines:**  
HKEY\_LOCAL\_MACHINE\SOFTWARE\Wow6432Node\  
Laserfiche\Import Agent 8\Profiles

---

**Tip:** A manual way of backing up a registry key is to select it the Registry Editor and choose **Export** under **File** in the menu bar.

**Note:** Ensure you back up this registry key whenever significant changes are made to profiles.

**Tip:** You can automate this process by creating a batch file. To do so, create a text file, change its extension to .bat, and enter the following: **regedit /E**

**C:\IA8Profiles.reg**

**"HKEY\_LOCAL\_MACHINE\SOFTWARE\Laserfiche\Import Agent 8\Profiles"** (this batch file, which only works with 32-bit versions of Import Agent, saves a .reg file to the root of the machine's C drive; change the key name (to the 64-bit key), output location, and file name, if necessary). This batch file could be called by a third-party backup solution or tied to a Windows Task Scheduler job.

---

## Recovery Tasks

Follow the steps below to recover Import Agent using backup data:

1. Reinstall Import Agent.
2. Double-click the backed up .reg file. When asked if you want to add the information in the file to your registry, click **Yes**.
3. Restart the Import Agent service.

## Agenda Manager

The Agenda Manager database and volume should be backed up.

---

**Note:** While the information below explicitly applies to Agenda Manager 8.0.1, it may apply to other versions as well.

---

## Agenda Manager Database

The database contains important information about your meeting types and users, among other things.

**To find your database:**

1. On the machine hosting the Agenda Manager service, click the Windows **Start** button, **All Programs**, **Laserfiche**, **Agenda Manager Server**, and **Agenda Manager Server Configuration**.

2. On the wizard's first screen, look for **SQL Server Instance** and **SQL Server database**.

## Volume

The volume contains your Word templates, attachments, published agendas, and published agenda packets.

1. Open the Agenda Manager Server Configuration wizard.
2. On the wizard's first screen, look for **Volume location**.

## Recovery Tasks

Follow the steps below to recover Agenda Manager using backup data:

1. Attach the database to the database server.
2. Copy the backup version of the volume to the machine hosting the Agenda Manager service.
3. Reinstall all necessary Agenda Manager components.
4. Run the Agenda Manager Server Configuration wizard.
  - When prompted to identify the database, point the wizard at the database you attached in Step 1.
  - When prompted to identify the volume, point the wizard at the volume you copied in Step 2.

## Web Access

Settings defined on the Web Access configurations page are saved to a configuration folder, which should be backed up. Doing so will save you from having to set up Web Access again during a recovery.

---

**Note:** While the information below explicitly applies to Web Access 8.2, it may apply to other versions as well.

---

### To find your configuration folder:

1. On the machine hosting Web Access, browse to the Web Access installation directory. In most cases, this will be **C:\Program Files\Laserfiche\Web Access 8**.
2. Browse to **Web Files**.
3. Ensure the **Config** folder is backed up.

---

**Note:** Ensure you back up the **Config** folder and not the **Configuration** folder.

---

## Recovery Tasks

Follow the steps below to recover Web Access using backup data:



1. Reinstall Web Access.
2. Replace the Web Access configuration folder with the backed up version of this folder.

## WebLink

The following WebLink data should be backed up: a configuration folder, customizations, and WebLink 8 Reporting information.

---

**Note:** While the information below explicitly applies to WebLink 8.0.2, it may apply to other versions as well.

---

### Configuration Folder

Settings defined in WebLink's Administrator's Utility are saved to a configuration folder, which should be backed up. Doing so will save you from having to set up WebLink again during a recovery.

#### To find your configuration folder:

1. On the machine hosting WebLink, browse to WebLink's installation directory. In most cases, this will be **C:\Program Files\Laserfiche\WebLink 8**.
2. Ensure the **Config** folder is backed up.

### WebLink Customizations

All customizations made using the WebLink 8 Designer should be backed up.

#### To back up WebLink customizations:

1. On the machine hosting WebLink, browse to WebLink's installation directory.
2. Browse into **Utilities** and open **WebLinkSettingsBundler.exe**.
3. In the WebLink Settings Bundler, define an export location and click **Export**.

---

**Note:** Perform these steps whenever you make significant customization changes to WebLink.

---

### WebLink 8 Reporting Information

WebLink 8 Reporting is a Web-analytics package that helps administrators see specific actions users are performing in a repository via WebLink. If this feature is enabled, you should back up the analytics data that is collected.

#### To find analytics data:

1. On the machine hosting WebLink, browse to WebLink's installation directory.

2. Ensure the **AnalyticData** folder is backed up.

## Recovery Tasks

Follow the steps below to recover WebLink using backup data:

1. Reinstall WebLink.
2. Replace WebLink's configuration folder with the backed up version.
3. Using the WebLink Settings Bundler, import the folder containing your backed up customizations.
4. Replace the WebLink 8 Reporting folder with the backed up version of this folder.



Designing a Laserfiche Backup and Recovery Plan  
April 2011

Author: Jonathan Powers  
Editor: Tammy Kaehler  
Technical Editor: Justin Pava

Description:

Designing, implementing, and maintaining a dependable backup and recovery plan is a crucial administrative task for your Laserfiche system. This paper will provide important and detailed information about this effort. It is written for Laserfiche administrators who have been tasked with backup/recovery responsibilities, and assumes some administrative Laserfiche knowledge on the reader's part.

Laserfiche  
3545 Long Beach Blvd.  
Long Beach, CA 90807  
U.S.A

Phone: +1.562.988.1688  
[www.laserfiche.com](http://www.laserfiche.com)

Laserfiche is a trademark of Compulink Management Center, Inc. Various product and service names references herein may be trademarks of Compulink Management Center, Inc. All other products and service names mentioned may be trademarks of their respective owners.

Laserfiche makes every effort to ensure the accuracy of these contents at the time of publication. They are for information purposes only and Laserfiche makes no warranties, express or implied, as to the information herein.

Copyright © 2011 Laserfiche  
All rights reserved