

# Backup and Recovery in Laserfiche 8

*White Paper*

July 2008

**Laserfiche®**

*The information contained in this document represents the current view of Compulink Management Center, Inc on the issues discussed as of the date of publication. Because Compulink must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Compulink, and Compulink cannot guarantee the accuracy of any information presented after the date of publication.*

*This chapter is for informational purposes only. COMPULINK MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION IN THIS DOCUMENT.*

## Table of Contents

Backing Up Your Laserfiche Repositories .....	3
Basic Principles of Data Backup.....	3
Test Your Backups.....	3
Back Up Often .....	3
Back Up All Necessary Files .....	4
Update Your Backup Procedure When Your Installation Changes .....	4
Destruction and Retention .....	4
Components .....	4
SQL Databases .....	5
Index files.....	6
Volumes .....	6
Audit Data .....	7
Backup Order .....	7
Running Backups .....	8
Retention Policy and Document Destruction.....	9
Are Your Records Really Deleted?.....	9
Volumes and Backup Media .....	9

## **Backing Up Your Laserfiche Repositories**

The importance of regular backups is difficult to overstate. Hardware failure, viruses, user error, natural disasters, or any of a host of other problems can wipe out gigabytes of data in an instant which will take weeks or months to re-scan and re-create.

Because of the dire risk posed by data loss, most companies want to have detailed backup plans for their important files. This guide describes the basic principles of backing up your Laserfiche system. Because each organization's configuration is unique, no single backup plan will be appropriate for all. The way you implement your backup plan will depend on your preferred backup software, your database management system, your hardware, and a variety of other factors.

This guide is principally intended for Laserfiche installations which use an edition of Microsoft SQL Server (such as Microsoft SQL Server (MSSQL), Microsoft SQL Desktop Edition (MSDE), Microsoft SQL 2005 Express, et cetera) as their database management system. Oracle users will need to consult their Oracle documentation for additional, more specific information on how to back up their database files.

### **Basic Principles of Data Backup**

Data backup is a crucial task, and there are a few principles you should commit to heart.

#### **Test Your Backups**

Many organizations have implemented solid, thorough backup plans, executed them faithfully on a regular basis, and then discovered too late that their backups had been failing silently for weeks or months. Always test your backups on a regular basis to make sure that you can restore your system in full.

#### **Back Up Often**

Make sure you back up your data as often as possible. We recommend a differential backup at least once a day, with a full backup of the entire system once or more per week.

## **Back Up All Necessary Files**

The Laserfiche server includes several different components, each of which must be included in your backup plan, in particular the volume files and the SQL databases. This guide will help you determine what you need to back up.

## **Update Your Backup Procedure When Your Installation Changes**

Certain changes to your Laserfiche installation may require you to update your backup procedure. For example, migrating from a previous version of Laserfiche to a new version will usually involve creating a new database, and you should make sure that, once the migration is complete, you are backing up the new database rather than the old and now-unused database. Similarly, if you move some of your volumes to another hard drive to distribute the load, or create new volumes, you need to make sure that your backup procedure backs up the volumes in their new location, or the files will not be backed up.

## **Destruction and Retention**

Most organizations are most concerned about preventing loss of data. Some file plans and retention schedules, however, require documents to be destroyed after a certain period of time has passed. In these cases, ensuring that such data has been properly and completely removed from the system can be just as important.

Backup plans, by definition, result in several redundant copies of documents, sometimes in several different locations. This can make it difficult to ensure that a document is really destroyed, since even after it is removed from the main system other copies of it may still exist on backup media. If your organization is required to destroy documents, you must take particular care to design your backup plan to avoid this problem.

To help make sure no extra copies are left, create a formal procedure for destruction of backup data, and coordinate your backup schedule with your document retention periods. If you back up individual volumes, have a procedure either for erasing or destroying your volume backup media. CDs and DVDs, for instance, can be shredded when the information on them is no longer needed.

## **Components**

In backing up your Laserfiche repository, you will need to back up two different components:

- The SQL databases that store the folder structure, metadata and other administrative information about the repository. For an MSDE, Microsoft SQL Server: Personal Edition, or Microsoft SQL 2005 Express database, this will include expansion databases as well as the main database. For an Oracle server, this information will be organized within the Laserfiche schema.
- The volumes that contain the images and text associated with the documents. These are stored as normal files within the Windows file system and must remain synchronized with the SQL database.

Backing up your DBMS may require different software than backing up your volumes and index files, although some third-party backup software can back up both SQL databases and Windows files.

This guide will explain some ways of choosing backup methods and methods for ensuring that your backups are current and functional.

## SQL Databases

SQL databases may be backed up in a variety of ways, depending on what software you have and what type of SQL DBMS you have installed.

**Note:** If you are using a migrated repository that was previously on the Team version of Laserfiche 6 or 7, you will need to be sure to back up your expansion databases as well.

- **Native SQL Server backup.** Oracle and Microsoft SQL Server include a variety of powerful and flexible backup tools. With MSSQL, you can use Enterprise Manager or SQL Management Studio to back up your repository; the Administration Console help files include step-by-step instructions for these tools. Oracle users should consult their Oracle documentation for detailed information about using the backup tools included with the Oracle server.
- **Third-party software with SQL option.** Third-party companies make software that can backup both your volume files and your SQL database files. Some of these products are compatible with both Oracle and MSSQL servers.
- **Detach and copy SQL database.** This method may be used with all editions of Microsoft SQL server. Stop the Laserfiche service, without stopping the SQL service, and perform the detach operation on your

SQL database through the Microsoft SQL server. Both the detached .MDF and .LDF files need to be copied. To restore the database, simply re-attach it to your server. Note that if you apply this method to a migrated repository used by Laserfiche Team (MSDE, Microsoft SQL: Personal Edition, or SQL 2005 Express) you will also need to detach and copy the associated expansion databases. Note that you should not simply turn off SQL and copy the databases; you must detach them before copying then and then reattach them.

- **Manual oSQL backup.** Microsoft SQL Server is installed with a command-line utility called oSQL, which can be used to back up your MSSQL database. While this utility will work for all versions of MSSQL, it can be complicated and difficult to use. We recommend using Enterprise Manager or SQL Management Studio instead. Microsoft SQL Management Studio Express can be downloaded from the Microsoft website.

**Warning:** You should not attempt to back up your database by manually copying .mdf and .ldf files. This type of backup requires you to stop and start the Laserfiche service to ensure that the files are not in use when you copy them. This type of backup is not recommended. It is inefficient and can result in incompatibilities if files are opened with different versions of SQL server.

## Index files

It is not, strictly speaking, necessary to back up index files; the server can recreate them if they are accidentally lost. However, for large repositories, the re-indexing process can be extremely lengthy, so we recommend that you back up index files whenever possible. The index files are stored in the SEARCH8 folder of the repository directory on the same machine as the Laserfiche Server.

## Volumes

Volumes are standard Windows directories, and so they may be backed up using whatever method you prefer to use for Windows files. The main consideration when backing up volumes is keeping them synchronized with your SQL database backup files. You can do this by making your volumes temporarily read-only.

If your industry is regulated, you may be required to archive your data. In this case, you can use volume rollover to keep your volumes at a manageable size,

and export them periodically onto archival media. Exported and archived volumes can be an excellent disaster recovery resource, because they include the database structure and metadata of the documents stored within them in XML format.

## **Audit Data**

If you are using Laserfiche Audit Trail and wish to ensure that you do not lose audit data, you will need to back up the Audit Trail binary log files. By default, these files are stored in the Audit folder of your repository directory, though you can configure another storage location for them. As Audit data is stored only in these logs, and is no longer present in the Laserfiche database, losing these logs may result in loss of your audit data; it is therefore imperative to include them in your backup plan. For more information, please see the Backing Up Your Audit Data section of the Administration: Reference.

## **Backup Order**

When performing a database backup, it is critical to synchronize backup data for Laserfiche volumes with the information stored by your DBMS. By keeping volume and database backups consistent, you can ensure a quick and easy recovery. If these backups are not synchronized, then the volume may be missing files that are in the database or it may contain files that are no longer in the database.

For example, if your volume data was backed up at 8 P.M. while your database was not backed up until 11 P.M., and if changes continued to take place to the repository in between those times, then recovery may generate a Laserfiche repository with missing data or data that is not associated with the repository. Documents that were added between the volume backup and the database backup will be reflected in the folder structure but will be missing all of their pages and text. It may be impossible to retrieve the contents of these documents. Therefore, it is vital to ensure that backups are synchronized.

The only way that you can ensure that volume and database backups are synchronized is to perform them when the repository is not in use. You can ensure this in several ways. You can temporarily disable the volumes, you can temporarily set the Everyone group to read-only, or you can take your repository offline.



## Running Backups

Once you have settled on a backup method, you will need to set up a schedule for running your backups, determine the locations where the backed-up data will be stored, and develop a procedure for testing backups to ensure that they can be restored.

**Differential backups.** Differential backups are much quicker and less processor-intensive than full backups, providing a way to get many of the benefits of daily backup without the costs in disk space and computer time. Most backup software allows you the option of running a differential backup, which will back up only the files that have changed since your last full backup.

**Frequency.** You should run a full backup at least once a week, and differential backups at least once per day. If your installation is very large, with several thousand pages added every day, you should strongly consider running differential backups two or more times per day. They may make the difference between losing an hour's worth of work, or losing an entire day. When possible, you should run your backups (especially full backups) at times when your system has a relatively light load, such as at night or on weekends.

**Automation and testing.** Most backup solutions will allow for some type of automation, whether through scripting or through built-in automation methods. You should automate the backup process as much as possible. However, automation cannot replace human supervision. Whomever is responsible for backing up your Laserfiche system should check the system frequently to ensure that backups are successful, and test to ensure that the data can be fully restored.

**Synchronization.** Running backups takes time, and it's possible that your SQL database may be changed while you are backing up your volumes, or vice versa. There are a few different ways to address this problem.

- Temporarily disable your volumes while you perform the backup; this will ensure they are not modified during the process. If you wish to fully automate your backup, you can write a Toolkit script to perform this step.
- Temporarily set your Everyone group read-only. This will allow users to continue to view documents, but not to modify them. (Users with the Manage Trustees privilege will still be able to modify the repository during this time, and so should be informed that the repository should not be modified until the backup is complete.)

- Take your repository temporarily offline while running the backup. This will guarantee that no changes are made during the backup process, but makes the repository unavailable for the duration of the backup. It is therefore a good idea to take this step if there is a time (for instance, very early morning) when few or no users will be attempting to access the repository. You can use a WMI script to unmount your repository to automate this step; see Introduction to Windows Management Instrumentation in the Administration: Reference for more information.

## **Retention Policy and Document Destruction**

Deleting documents can be as important to some organizations as preventing them from being deleted. This may be due to privacy concerns, at organizations dealing with sensitive patient or customer data, or security concerns at organizations where proprietary information is commonly used. In some cases ensuring proper destruction is simply a form of housekeeping—freeing up disk space for other needs and ensuring that repositories don't get too large. Your document destruction policy and your backup policy will interact with each other, so it's important to make sure that they're appropriately synchronized.

### **Are Your Records Really Deleted?**

The existence of backups can stymie full document destruction schemes, because many documents will exist in two or more locations. Destroying documents in the main repository could leave a document manager with the false impression that a document has been safely disposed of, when in fact it is still present on backup tapes.

To avoid this problem, coordinate your backup schedule with your retention schedule. Your backup schedule should be no longer than your retention schedule—if you destroy documents after six months, you should keep no more than six month's worth of backups.

### **Volumes and Backup Media**

Many organizations back up and archive Laserfiche files by detaching older volumes and storing them on backup media such as a DVD or CD. If you do this, you will need to physically destroy the media when it is time to destroy the files on it.



Backup and Recovery in Laserfiche 8

July 2008

Author: Constance Anderson

Technical Editor: Justin Pava

Compulink Management Center, Inc.

Global Headquarters

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.

Copyright © 2008 Compulink Management Center, Inc.

All rights reserved