



LiveVault® from Backup Direct™ Protecting Open Databases and Files

Overview

Active databases and open files are typically your most important data; therefore it is critical that you employ a data protection solution that natively supports backing up these mission-critical files. Backup solutions that require complex plug-ins or costly third party tools to protect open databases and files, such as Open File Manager (OFM) or Open Transaction Manager (OTM), are vulnerable to serious compatibility and operational problems between the backup solution, the database, and the database plug-ins. Such backup solutions are not always compatible with latest release of OFM/OTM. In addition, upgrading a database can cause serious compatibility operational problems between the backup solution, the database, and the database plug-ins. Unfortunately most of today's backup solutions require these third-party tools to back up open databases and files.

Open Database Protection with LiveVault

LiveVault solutions from Backup Direct™ remove risk and complexity through built-in support for backing up open files and databases. The LiveVault backup and recovery service continuously protect open and changing databases such as Microsoft® Exchange Server, Microsoft SQL Server, Oracle® and Lotus Notes® without the need for complex plug-ins or costly third party tools. This is a unique and powerful advantage of LiveVault.

This document describes two technologies employed by LiveVault that work together to safely and effectively protect open files and databases: **1) Snapshots** to ensure access to open files and to ensure transactional consistency of databases; **2) a software agent** utilizing Microsoft's NT Filter Driver APIs to interact directly with file system-level functions to efficiently record changes to files as they occur.

Job Phases

Individual backup processes are referred to as jobs which transition between two phases: **initial backup** and **continuous or scheduled backup**.

Initial Backup

The first time a server is configured for backup, an initial backup occurs. A copy of each file and directory selected to be backed up is efficiently and securely sent to the vault, creating a baseline image of the files and file structure. The initial backup runs until completion. If the initial backup is interrupted (due to server or network outages) the backup will resume automatically from where it left off. When the first copy of all data selected for backup has been replicated to the vault, restore jobs can be successfully run to recover data.

Continuous or Scheduled Backups

After the initial backup of a file, the LiveVault solution only needs to transfer the changed deltas within the file. After the initial synchronization, LiveVault always performs delta

backups, and never needs to perform another synchronization. This enables the LiveVault solution to efficiently back up data 24x7 (typically sending changes to data on 15 minute intervals), while minimizing network bandwidth usage.

Replication Technologies

Two powerful technologies are used in parallel during the phases described above.

Point-in-Time Versions via Snapshots

During backup jobs (typically at 15 minute intervals), LiveVault instantiates snapshots. LiveVault uses snapshot technology for two primary reasons:

- 1) Snapshots give the LiveVault agent full access to a server's file system without the need for open file managers and without risk of impacting files currently in use.
- 2) Snapshots lock the entire file system at a point in time. This allows for transactional integrity of restored databases.

Snapshot definition

Snapshots are a file system function. They provide an isolated, virtual file system for specific applications (such as backup) to access files as they exist at a moment in time and without having to compete with other applications using those files.

In practice, when the LiveVault agent requests a snapshot an empty snapshot file is created. After this, if an application modifies a file, a "pre-image" of the blocks of the file being modified are moved to the snapshot file. When the LiveVault solution then asks for this file, the snapshotter sends LiveVault's request for this data to the snapshot file. Once the backup job is complete, the LiveVault solution deletes the snapshot file.

Efficient Delta Backups via File System Filter

The Microsoft Filter Driver framework is a well-defined set of programming interfaces developed by Microsoft to allow third-party services to interact with actions occurring within the file system.

LiveVault solutions use a file system filter to monitor the changes that occur between snapshots. This way, LiveVault is able to efficiently capture and send just the changed deltas after a snapshot is instantiated.

The filter removes the need for the LiveVault solution to scan the file system for each backup job to determine the changes to the file system. This significantly improves the performance of backup jobs (which typically finish in seconds to minutes).

Not all Microsoft OS versions include the filter driver framework. When the framework is not present, the LiveVault agent software scans the file system to identify files and blocks that have changed. This will take longer than when use of the filter driver is possible, depending on the number of files and the amount of data in the file system. For very efficient delta backups using the LiveVault filter you should be running:

- Windows 2000 with SP4 Update Rollup 1
- Windows 2003 SP1

Putting It All Together

By using snapshots and filter drivers, LiveVault is able to natively protect open files while ensuring the transactional integrity of databases. This is accomplished with a minimum of resource consumption. Through the use of these technologies, LiveVault's solutions address two very important obstacles in backing up open and changing files: 1) files open for exclusive read and 2) actively changing files.

Backing Up Open and Changing Files

Snapshots, by definition, provide complete read access (with no ability to modify) to all files on a system to applications using the snapshots.

With the filter driver monitoring file system activities, the LiveVault agent is able to 1) view files as they are opened and 2) log the location of changes occurring in the files as commands are executed on open files. Changes are tracked and their locations tracked (no data is cached during this process, only records of the changes and their locations).

Database Transactional Integrity

In a 24x7 production environment, many database or application files are always open and actively changing. If the server suffers a power failure and is restarted, one of the functions of the database software is to bring the database into a transactionally consistent state as part of its startup activity by "rolling out" any incomplete transactions. To do this many databases keep transaction log files as well as the main database files. All modern databases such as SQL, Exchange, Oracle, Sybase and Lotus Notes maintain transactional consistency in their databases.

Backup products typically require either that the database be brought into a quiet state or shut down before backup to insure that the collection of database files (log files, database files, etc.) are in a consistent state between each other.

In contrast, LiveVault backup does not require any database plug-ins or that the database be shut down or that any pre- or post-processing occur. The reason is that through the snapshot technology LiveVault can guarantee that all the database files are "captured" at a specific instant in time which reflects the on-disk state of all those files at that instant. If the collection of database files is later restored, when the database starts up it will "see" the files in a state that could have resulted from a power failure and, if necessary, will make the database transactionally consistent as part of its startup.

Handling Log Files

With traditional backup, the normal scheme is to retain a database's transaction log files until after a known good backup has occurred. Often the backup software plug-in must reset the log files or command the database software to do so. In one way or another the size of the log files must be contained so that they don't eventually fill up the disk.

LiveVault does not need to interact with the database software so a strategy is needed to control log file growth. The normal approach depends on the database. For example:

- For SQL, use the database consistency checker (DBCC) on a regular basis.
- For Exchange, use circular logging
- Oracle always uses circular logging. However, if ARCHIVELOG mode is on, you will need a simple script or procedure to periodically clean the archived redo log files.

Using Native Database Backup Features with LiveVault

Many database packages provide features for handling or assisting with backup. These features can be used in parallel with LiveVault if so desired. A few typical examples are mentioned below:

- SQL provides a SQL Backup capability to make dumps of the database and/or transaction logs into a separate backup directory. LiveVault's continuous offsite protection often makes the use of SQL Backups unnecessary since recovery to a 15 minute point-in-time in the recent past is normally possible with LiveVault; however some customers choose to continue to make SQL Backups a few times a day as extra protection.
- ORACLE offers a logical backup facility as well as the RMAN utility. Oracle also provides ARCHIVELOG mode to make possible the recovery from online backups. As discussed above, LiveVault performs online backups in a unique way. With LiveVault, you backup the main database files (as they are changing) as well as the collection of online redo log files (as they are changing). You do not need to backup the archived log files, although you can certainly do so if you wish.

Conclusion

Open databases and files are the most difficult to protect. Yet they are typically your company's most valuable data assets. LiveVault's backup and recovery solutions help you overcome this challenge by providing reliable, automated protection of your company's valuable data, including open databases and files. Through the use of sophisticated technology, the LiveVault online backup and recovery service and the LiveVault InControl remote office backup solution both provide secure, continuous protection for open files and databases without the need for third-party plug-ins. With LiveVault, critical open database applications such as Microsoft SQL and Exchange, Lotus Notes, Oracle and other core business applications can be continuously protected 24x7 without the need for complex backup software, third-party add-ons or burden associated with traditional backup.

The two most common applications that are protected by LiveVault's customers are SQL and Exchange. This is hardly a surprise since these applications normally contain the mission critical data and therefore warrant the best possible protection.

For more information on LiveVault's backup and recovery solutions, please visit www.backupdirect.net or call 08000 789 437.