



The Leader In True Image Storage Solutions

PO Box 163088 - Altamonte Springs, FL 32716

OTM™ TECHNOLOGY MAKES BACKUP WINDOW A MEMORY

Many of today's most critical applications, such as email, web servers and transaction servers, can no longer be shut down for backup. They are constantly up and running...24 x 7. And yet, no data center manager would eliminate backup as a critical management task.

How, then, to address open file backup? The conventional approach is open file agents and managers that insured that all files were backed up intact and not skipped because they were open. But the complexity of applications and operating systems has spiraled, and data loss is once again a major problem facing the data center.

Strategic Research Corporation (Santa Barbara, CA) reports that users could not recover a usable intact dataset from their backup in over 35% of the time when they tried a restore function. Many times the files were successfully restored intact, but would not operate correctly since they do not match each other's timestamps. To address these problems, another enabling technology has been developed

Open Transaction Manager™(OTM)

This is an enabling technology that presents a stable, coherent, point in time snapshot alternative view of one or more volumes or physical drives to any backup application without impacting system performance. OTM permits a user to back up a server or workstation with all files and databases open and active. Popular databases like Oracle and Lotus Notes, SQL servers, and Email servers can be in use and running.

The technology creates an alternate "virtual drive" or static copy of the drive to be backed up. [Figure 1] Software looks for a short period of inactivity, five seconds or so, where no writes are occurring to any of the volumes or drives that have been selected for backup. Once this "quiescent period" is obtained, OTM maps in a virtual drive letter for each volume selected for backup. The backup application can now access this static virtual volume instead of the original volume, which is changing during the backup.

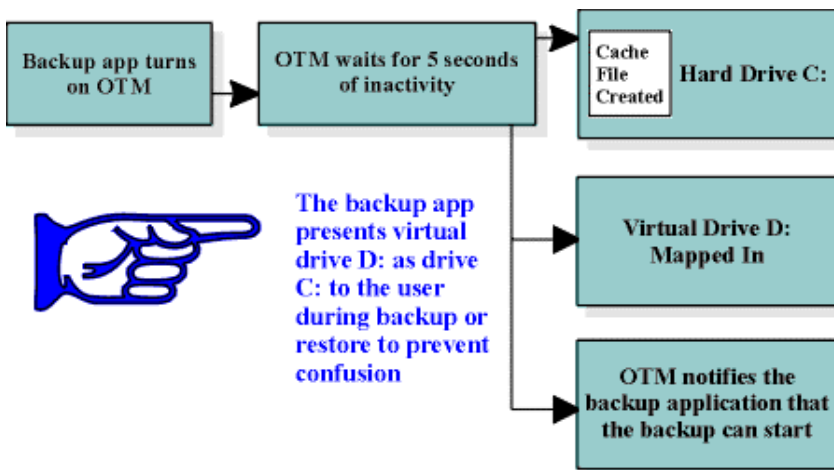


Figure 1

BACKUP PROBLEMS IN THE DATA CENTER

Conventional open-file agents or managers face significant problems in the data center. Here are a few:

The Timestamp Problem (Relational Integrity)

Most of today's applications have multiple associated files that are all updated together. When a backup is done while running an application, the backup copies file 'A' to tape. Now an application subsequently updates file 'A' and another associated file 'B' that must match each other for the dataset to load. The backup now backs up file 'B' to tape, which no longer matches file 'A' on the tape. Everything looks good to the administrator, until restore is necessary. Convinced that the file restore must have gone awry, he tries again and again, but it just won't run the data after 'restoring'.

Disappearing Backup Window Problem

Most of today's most critical applications such as email, web-servers and transaction-servers, can no longer be shut down for backup and are running 24x7. This makes backup integrity a virtual impossibility when utilizing open file agents or managers.

When a write occurs on the original volume, OTM pauses it and copies the old corresponding data to its cache file and immediately sends the original write data to the hard drive. This keeps the hard drive real time current and safe at all times during the backup. Read requests from all applications except the backup are passed directly to the hard drive with no intervention. Read requests from the backup package are passed to the OTM filter driver that determines if the data is already in cache.

If that data is in cache, OTM passes the cached data to the backup package. If not, the data is passed directly from the hard drive. Since OTM only needs to preserve the original data, additional writes to the same sector are not cached and passed directly to the hard drive.

Banish The Backup Window

The mere fact that there is a "backup window" (that small slice in time when a backup can be performed without affecting system performance) is a major problem for most system administrators. All file-by-file backup applications have relied on secondary and tertiary agents and managers just to allow the backup of the open files while the system is live and in use. But, these file-by-file applications combined with open file agents and managers are extremely CPU and I/O intensive. This can bring server performance down to its knees during your backup.

OTM makes all of the file, database, Ethernet, CPU, and I/O subsystem resources available to users on a priority basis during the backup by throttling the I/O. This increased availability makes the time to complete the backup a secondary issue.

A faster backup would also minimize that time. OTM "sees" a True Image of the file server. All I/Os pass through OTM, even all of the OS's hidden system I/Os. Because of this, only OTM can safely give users priority over secondary processes like backup.

Since servers can task switch many times a second, once True Throttling is implemented no one should ever have to wait for their data. As backup applications have become more efficient at consuming all available resources, this background application can effectively shut down the more costly applications--users. OTM realigns these priorities according to common-sense management guidelines. The backup also takes less time, since the timestamps no longer need to be reset after backup on the virtual OTM drive.

I/O Performance Supports The Backup Function

Since OTM sees all I/Os, OTM pauses the backup I/Os if there are outstanding user I/Os. Since the system will timeslice, no one will ever have to wait for their data. When the system is at 100% I/O utilization, then the backup will wait, not the user. Because the I/Os from the backup are being paused in a 100% I/O saturation situation, the backup application itself will not be issuing more file requests to the OS until the current requests are fulfilled. If the backup file requests were to continue, the OS itself could quickly overburden the CPU, the network cable, and the storage subsystem. By preventing more backup file I/O requests, the OS, CPU and the Ethernet are free to service all user requests and the backup resumes only after the user requests are satisfied.

OTM works at the block level, and not the file level, all of its work is done within OTM and not the OS and the file system. This improves performance in read/write operations. Since write operations are typically less than 10% of reads, OTM's I/Os are generally invisible to the user or the backup.

Configuration Induced Loss

When using open file agents or managers, the only way to overcome this limitation is for the administrator to learn each applications' file structures and directories, then configure groups of all the associated files together in the open file agent or manager, and finally test to make sure that no files were missed in the grouping. Since many applications write these files in multiple directories, the chance for error is high.

Improved System/Backup Performance

Since open file agents and managers must operate at the file level, they can seriously degrade the server's performance by consuming huge amounts of CPU and I/O bandwidth. OTM can be operational without seriously impacting system performance as it acts at the disk sector level, bypassing the operating system, and thus CPU overhead. Open file agents and managers consume vast system resources like most backup software--users and vital processes are left out and can't run at the same time as the backup in this virtual 'musical chair' performance game.

no delay or system impact. For reads from the backup application's alternate "virtual" volume, these blocks reads are passed through to the hard drive if the data has not changed since backup commencement. For block reads from the virtual volume that have changed, OTM will read only those blocks from the hard drive's cache file, resulting in slightly slower file access to the backup.

For writes, either from the user or the backup application, the old data is first copied to the cache file before the write is passed to the hard drive. OTM speeds up this operation by utilizing the "lazy write" feature of the OS to write to the cache file. Additional performance can be achieved by placing OTM's cache file on a separate system. All subsequent writes to the same sectors will not be cached, but would be immediately written to the hard drive. Once volumes have completed their backups, OTM can be told to release them, thus allowing all subsequent reads and writes to be immediately passed to the hard drives without OTM.

The 'Daily' Upgrade

Dedicated open file managers and open-database agents must have intimate knowledge of the underlying database to work correctly. So anytime the database vendor makes a change; the open file agents vendor must make a corresponding change. So the customer must always be playing catch-up, installing new software and then going through extensive testing to insure that it is working and configured properly. Some open file agent vendors are now grouping all files together, which improves the odds of a good backup. Unfortunately, if server activity is high during the first 60 seconds, all the grouping goes away, and so does your dataset relational integrity.

Hidden Data Loss

As more and more hidden structures in the operating system are suddenly 'discovered', file-based open file agents and managers must constantly change to accommodate these new objects--and they may not be allowed by the underlying OS to properly protect them.

Copyright 1999 Columbia Data Products, Inc. All rights reserved. Open Transaction Manager, OTM, True Image and True Throttling are Trademarks of Columbia Data Products, Inc. Open File Manager is a Trademark of St Bernard Software. All other trademarks are the property of their respective companies. Reprinted by permission of Columbia Data Products, Inc.

Integrated backup solution partners . . .

