
CYA (Covering Your Assets): One Day Data Backup Will Save Your Law Practice!

by Ross L. Kodner

©2005 Ross L. Kodner, All Rights Reserved

(Appeared as a two part series in August/September and October/November 2005 issues of Law Office Computing magazine)

Let's face it: there is **NOTHING** more boring for a group of lawyers to talk about than data backup. Federal telecommunications regulations. The earliest derivation of the word "hopscotch." Quantum mechanics and its applicability to modern cuisine. **Anything** is more exciting than the dreaded six letter techno.word: "backup." In animated fashion, with more than occasional one-upmanship involved, the conversation in a room full of techno.enabled lawyers is typically peppered with comparisons of the speed of the latest Pentium M 2.1Ghz laptops they're toting, how long their laptop batteries last, how big their laptop screens are (um, lending credence, with an admittedly geeky twist to the old adage . . . well you know how it goes). Never, ever would you overhear such a group bragging about the alternating backup cycle they use or their "hot" new Ultrium (super high capacity, super high speed digital linear tape) backup system. Why? Because backup is EXCRUCIATINGLY DULL!

The periodically painful reality is that lawyers should spend more time thinking about the backup systems that should be protecting their increasingly electronic-based irreplaceable client and firm data. It seems inevitable that there will be a successful legal malpractice suit brought at some point in the future by a client angry about their lawyer's failure to protect work product against virus attacks or a catastrophic hardware failure. It is shocking how many lawyers and their entire firms, for that matter, do not have "adequate" data backup systems and procedures in place. My own informal and admittedly anecdotal polls of CLE audiences show as high as 75% of lawyers can't say, unequivocally, that if their offices burned down overnight, that they could restore their computer programs and data from a backup made the night before. That's frightening! The "it can't happen to us" mentality seems to rule--which of course tempts the "pixies" and defies the inevitability of Mr. (or Ms.?) Murphy's famous law.

At some point, the hard drive(s) in your office PCs **will** fail; it's guaranteed--as certain as death and taxes. The questions to ask are: (1) will your software survive? (2) will your critical irreplaceable client data survive? (3) will your equally critical firm systems (such as your billing/accounting and calendaring/docketing systems) survive? Let's take a short look at the latest twist on the time-honored legal tradition of "CYA": Covering your Assets.

First a quick definition: the process of backing up one's system simply means making a copy of the contents of one's hard drive(s) onto some other storage medium. In the event of a failure of any hard drive, the idea is that you can restore or rebuild your programs and your data--your documents, your calendar entries, your e-mail, your billing data, your conflicts information--by copying the information you previously copied and protected onto a new or repaired hard drive. Pretty simple, at least conceptually. So why is it that so many law firms find the process of implementation and execution of backup systems so torturous?

Normally backup involves some process that copies information stored on your PC(s) hard drives to some other type of media which can include other hard drives, tape drives, write CDs or DVDs, or even backing up information to specialized data backup websites.

There are some basic rules about data backup that every lawyer and law office staffer needs to memorize, to chant, to hear ringing over and over again in their subconscious thoughts (yes, just like that infernal "It's a Small World" that you can't get out of our head **for months** after treating the kiddies to Disneyworld). My mantras of using backup systems to help you Cover Your Assets are:

- ◆ **"We do not backup for the sake of backing up . . . we backup ONLY so we can restore when the techno.chips are down,"** and
- ◆ **"If you are doing anything in your backup process that could affect your ability to quickly restore your system . . . STOP IT NOW!"**

There are also several sub-corollaries, equally important to faithfully adopt:

- ◆ **Never Trust Your Backup System**--periodically perform a "mini test restore" to confirm the system is fully capable of bringing back your ENTIRE system's contents. If you believe your backup software when it tells you it backed up successfully the night before, 1 time out of 1000 it will lie - and that will be the night before the Mother of All Network Hard Drive crashes.

(To perform regular "mini-test restores" you can create some variation on a simple procedure involving locating two document files that had been backed up the previous evening. Then, copy those files to another folder for safekeeping. Delete them from the original location. Then use your backup software to tag and restore those two files. Then retrieve them using the application that created them to confirm they were properly restored. The point of this is that your ability to

randomly select and restore two files is probably indicative of the ability to restore all the files from the backup session, validating the process in the only way that makes sense - proving the restore capability. How often? There is no such thing as “too” often, but definitely such a thing as not often enough. Prudence and caution dictate daily or at least several times per week).

- ◆ **Alternate Your Backup Media** - Alternate between at **least** five backup tapes or drives (labeling them “Monday” through “Friday” is a very easy approach - more tapes are always better though). Alternate between them to spread your risk of relying on a single tape/cartridge over several days. Also it means you have snapshots of five or more days of work to restore from as needed. Many firms use 10 tapes or drives (two full work weeks) plus 12 “monthly” tapes/drives used the last day of each month and then archived permanently. More is better.
- ◆ **Store Backup Media OFF-Site, Not in Your Building** - Store your most recent backup tape/drive out of your office (actually out of your building--**not** in a purportedly fireproof, theftproof safe or file cabinet) to protect against fire, theft/vandalism or some other disaster. Diligently backing up your system every day and then having the tapes melt in the office fire doesn’t accomplish much.
- ◆ **Retire Old Backup Media** - Retire backup tapes after 1-2 years of use and replace them. Tape stretches over time and multiple use. It becomes less effective as a result of normal wear and tear. Don’t chance it--toss ‘em and replace ‘em after two years. Hard drives used for backup should last longer (3-4 years depending on how many are in the alternating cycle (the more you have, the less often each is used, the longer the usage period before replacement. Be sure to TRULY erase these discarded tapes--remember that they are storehouses of your confidential client and firm data!! Think about buying a “bulk eraser” from Radio Shack--about \$30--a couple of swipes of these vibrating, super-magnetic field-emitting gizmos and most traces of recoverable data will be wiped out. Incineration works also, if not particularly environmentally conscious. Of course, do this FAR away from your PCs, other hard drives, and probably even pacemakers). For more information on the safe and ethical disposal of backup media (as well as PCs themselves) *see* “Dumpster Disasters: Tips for Ethically Retiring Your Old Computers,” Kodner, Ross and Kennaday, Courtney, ABA GPSolo Magazine, December 2004.

-
- ◆ **Never Trust Your “Liveware”** - Train at least **two (or more) people** in your office on how to perform backups as well as conduct and regular and frequent test ‘Restore’ operations. Never trust the process to just one person, who could let you down. Have a procedure where each staffer with backup responsibility checks and confirms the work of the other. You don’t want to fall victim to the worst failure of all . . . defective “liveware.”

 - ◆ **Protect Against Downtime - Multiple Layered Protection** - Data backup processes seem particularly susceptible to Murphy’s Law and Worst Case Scenarios. Relying on a single backup method is tempting the fates. You should employ some type of secondary backup protection. Examples of secondary backup approaches might be:
 - * Using a real-time data backup system on your server and or individual workstations. Programs such as Second Copy 2000 (www.secondcopy.com) and Autosave (http://www.v-com.com/product/AutoSave_Home.html) work in the background, performing a real-time data backup of files to an external hard drive or mapped network drive letter. The files are stored in their native form meaning they are immediately accessible if the server fails. In this respect, this approach also serves to minimize downtime in the event of drive or server failure (until the primary backup can be restored to a system that has been returned to operation).

 - * Leveraging the built-in backup and/or synchronization functions of key software systems. More and more applications have “internal” backup functions. These usually will make a working copy of the core program and/or its data files on any drive letter specified, either per an automated schedule or manually initiated. Examples include programs like the Worldox document manager (www.worldox.com) which makes “mirror” copies of documents worked on to a designated drive/folder. Case managers like Time Matters (www.timematters.com) can synchronize to a local hard drive with a “clone” copy - a fully functional remote copy of the entire system. Both of these approaches can provide system functionality and minimize the frustration and cost of downtime with access to documents and practice management data in the event of a system failure.

-
- * Reduce downtime with hardware options - specifically, in your network server (if you have one), use a hard drive configuration that is “error-tolerant.” This means some version of hard drive mirroring or a RAID array of drives. Either method focuses on protecting against downtime caused by the failure of one or more hard drives in the server. Depending on the methodology used, this can either reduce downtime somewhat or entirely, while a replacement drive is deployed. Never mistake this approach, however, for a true backup process with critical off-site storage as a key element of the process.
 - * Reduce downtime by synchronizing a drive’s data to another drive - from the server to an external hard drive on a network station for example. Programs like Network Unplugged from Mobilit do this so easily that you will wonder why you hadn’t (a) known about it, and (b) been doing it already for years (<http://tinyurl.com/cg6rr>) (from \$80 to \$252 depending on version). A similar and newer product is called FolderShare (www.foldershare.com) (from \$49.95 to \$99.95 depending on version).

The point is simple: don’t put all your backup eggs in one basket. Mr. Murphy, the Boogeyman and Gremlins will sniff you out and wreak havoc on the unwary practice relying on one backup methodology.

Backup What?

The next question to consider is exactly **what** do you back up? The **ONLY** way to truly protect your system’s information is to backup **EVERYTHING** on your key hard drives every single day. This means your software with the myriad of settings, configuration changes, tweaks and customization, special macros, etc. It also means your data: your billing and accounting information, your calendars, your case management system’s info, your documents . . . **EVERYTHING!**

Full backups of your system require several things, practically speaking. You need a backup system with enough capacity to store ALL the information that is located on the hard drives you are backing up. In other words, if you have a 120 gigabyte hard drive, your backup system should be capable of storing at least 120 gigabytes onto a single piece of “backup media” (without using the “compression” capabilities touted by data backup hardware and software vendors). You also need data backup software that can automatically perform a backup session, unattended, in the

middle of the night (or when your people aren't busy working on your system with many files held "open").

An undesirable, but all too commonly used backup approach is referred to as an "incremental" or backing up "changed files only". The idea here is that a full backup of a hard drive might be run weekly. In between, each workday a backup of only those files (programs and data alike) that have changed since the time of the last backup are in fact, backed up. This tends to make for pretty short daily backups and on its face, that would seem to be a good thing. Under the general theory that we only backup so we can restore at some point, restoring from a collection of disks or tapes containing incremental backup sessions can be nightmarish--assembling bits and pieces of one's system scattered across multiple media.

Can't I Backup to CDs or DVDs with My New "Burner"?

Inevitably, the same question arises periodically on the legal listserves - seemingly in four month cycles: "Can I backup to CDs or DVDs since I have a PC that can now burn both?" Or sometimes it is stated this way, as appeared on the Technolawyer list (www.technolawyer.com):

A community member wrote: "I back up over 10 years of data on one CD. I never back up programs. If disaster strikes, I will either restore from my original program CDs, which are off-site, or buy new copies of the programs. I keep my data on a separate partition from the programs. The CD-R simply backs up my data drive."

This kind of comment always amazes me: "I don't backup my programs because I have them all on CD."

Huh? What about all the settings and configuration tweaks in every program - none of which is being backed up? Do you remember all the settings you've made in every application? How about years of invaluable and irreplaceable macros in Word and WordPerfect? How about Bookmarks and Favorites? How about all the configuration settings and changes in your case manager, your billing and accounting systems? Is that information in the data folder or not? And what about the layers and layers of patches and updates that would then have to be reinstalled after you install the CDs, presuming you can actually find all your CDs. How about rebuilding Windows itself and patching and updating it and getting all the drivers you need? And go and BUY all your programs if you can't find the CDs? Do you have an infinite budget and endless time on your hands?

How long would all this take? Would you ever get the system back the way it was before the crash? Within a week? Doubt it.

When it is possible to buy a reliable, working, automated system that can do full backups every day for as little as \$500, it continually amazes me to see the tortured ways that people try and save a little money (and typically more than make up for it in scads of lost otherwise billable time). Curious aspect of human behavior. Just amazes me how people want to regularly choose the hard way when there is an easier way to do it.

What this person is really describing could be best classified as a "spot backup" approach, backing up data files only. In my CLE materials on the subject of data backup (entitled "Safety, Safety, Safety: Data Backup from A - Z" at www.microlaw.com/malpractice_prevention.html) I discuss the drawback of this approach as a primary backup method stating: "The ONLY way to truly protect your system's information is to backup EVERYTHING on your key hard drives every single day. This means: Your software with the myriad of settings, configuration changes, tweaks and customization, special macros, etc. Your data: your billing and accounting information, your calendars, your case management system's info, your documents . . . EVERYTHING!"

The point here is that people who say "I don't need to backup my programs. I've got 'em all on CD and I can just reinstall them" truly need to get their techno.heads examined. How many hours or days would it take to reinstall every program, every patch for every program, run through the myriad (morass might be a better word) of configuration settings and tweaks, rebuild Word or WordPerfect macros (or both), not to mention totaling reinstalling Windows, all its service packs and critical updates, replacing all your lost bookmarks/favorites, printer drivers and the updates for them, any other peripheral drivers or Palm software, recovering all your e-mail, e-mail archives, etc. The likelihood of EVER returning your system to the way it was pre-disaster is effectively nil.

And that doesn't address the question of what "data" you would actually protect. Documents certainly, but do you know where all your billing and accounting data files are? How about your calendar or case management system's info? How about your bookmarks/favorites? What other data is stored by your system in locations OTHER than your documents folders? So exactly what WOULD you backup?

I just plain don't understand why anyone wouldn't take the safe (and frankly easier) approach and do full system backups, in unattended automatic mode to tape or hard drive media, every single night? What's the point of trying to shortcut the process? To save time? That's absurd -- there's no time saved and certainly not if there is a system disaster. Why people fight this is just beyond me ... simple and safe are always better.

The “How” Factor: Different Backup Approaches

There are multiple approaches to backup. Examples include:

- * **Traditional:** tape backup systems
- * **Modern:** hard drive, disk-disk-tape systems, online backup, removable optical media (CD, DVD, Magneto-Optical disks)

Tape Backup: The Story in 2005

Backing up to tape is an approach that has been with us since the early days of legal computing. Today, backup to tape is still a proven, tested approach that can be very cost effective. There are a number of classes of tape backup devices available, generally categorized as follows, along with relative advantages and disadvantages:

Travan - lower capacity drives from a number of major manufacturers. The highest capacity tapes are referred to as “TR-5” and have a maximum capacity of only 20 GB per tape, with compression activated (which is never desirable, so effectively, these are 10 GB cartridges).

Advantages: Inexpensive.

Disadvantages: None, capacity is too small to be useful today. Further, Travan class tape units are notoriously unreliable, frequently either failing to backup information or worse, to restore it. Tapes also prone to physical failure.

Recommendation: Advice? Run the other way and never, ever rely on a Travan-class backup unit - the category might be better spelled as “Travesty.”

DAT - used to stand for “Digital Audio Tape” when initially released, now just referred to as DAT. The tapes used are called DDS (Digital Data Storage). Capacity is up to 36 gb per tape, uncompressed with the latest DAT72 standard. Characteristics are reliability, relatively rapid data transfer times and inexpensive tape cost.

Advantages: Inexpensive drive pricing (starting at about \$500) for the DAT72 drives from major vendors such as Hewlett-Packard (<http://tinyurl.com/auz9l>). Inexpensive tape cartridges (starting at about \$15 per tape). Proven reliability. Relatively quick transfer rates between 3 and 6 MB/s (megabytes per second).

Disadvantages: Relatively low capacity per tape. For larger scale backup, might require manual intervention to change tapes mid-process. Transfer rate too slow for larger capacities.

Recommendation: Perfectly acceptable as a lower-cost primary backup method for small firms with less than 36 Gb of server information to backup.

VXA-2 - An intermediate tape system that bridges the gap between DAT and LTO Ultrium drives, VXA is now the exclusive province of widely-known drive maker, Exabyte. VXA-2 offers a capacity per tape up to 80 Gb uncompressed. The emerging VXA-3 standard offers capacities up to 160 Gb uncompressed per tape. A combination of ruggedness and relatively low price have created a significant market for the VXA-class devices.

Advantages: Moderate drive pricing (starting at about \$650) for the VXA-2 drives. Inexpensive tape cartridges (starting at about \$32 per tape or VXA-3 at about \$75 per tape). Proven reliability. Relatively quick transfer rate of about 12 MB/s (megabytes per second).

Disadvantages: Still relatively slow transfer rate compared to the faster DLT-class LTO Ultrium drives.

Recommendation: Perfectly acceptable as a lower-cost primary backup method for small firms with less than 80 to 160 Gb of server information to backup.

LTO or LTO Ultrium - (Linear Tape Open) higher capacity drives in the DLT (Digital Linear Tape) family, descended from the “big iron” UNIX “minicomputer” world, these are rugged, robust, fast and not inexpensive. The LTO Ultrium-class tape drives provide data storage capacity of up to 400 Gb (uncompressed) and a maximum transfer rate of 80 to 160 megabytes per second (MB/s). Prices start at about \$850 for the 100/200 Gb models and can surpass \$2500 for the highest capacity 400/800 Gb models.

Advantages: Fastest tape backup method. Proven reliability. Quick transfer rates between 80 and 160 MB/s (megabytes per second) on the LTO Ultrium series make backup of large hard drives very practical in a “middle of the night” automated scenario. Very rugged backup tapes.

Disadvantages: Most expensive backup media in terms of both drive cost and tape cost (the latter starting at about \$45 each).

Recommendation: Rock-solid approach - bar none, the best approach in a tape-based backup system, albeit if a bit costly as “data insurance.”

There is a hybrid approach in the tape backup arena worth mentioning. Certance produces a series of hybrid devices in their CP3100 series (www.certance.com). These devices combine a tape backup unit and a hard drive in a single chassis. The system performs a backup initially onto the faster internal hard drive (160 or 320 Gb capacity - and sustained data transfer rates of up to 79 Gb/hour). Then the system backs up the information to an internal DAT72 tape system for off-site storage. The idea is that the performance drain on a network is lessened v. traditional tape because the initial process of copying first to a hard drive is significantly faster. Interesting concept. Prices start at about \$1200.

Tape Backup Tips:

- * Always buy a cleaning tape and run it monthly - replace the cleaning tape after 12 uses. Maximizes backup and restore read/write reliability.
- * Consider buying a second, spare tape drive. Why? What if you need to restore three years down the road and your tape drive, attached to the network server, melts in the fire that burns down your office? If your drive is no longer manufactured, it could pose problems restoring your data. If you buy a spare drive and KEEP IT OUT OF YOUR BUILDING, you have a spare instantly available to you.

Hard Drive Backup:

An approach that is gaining in popularity in part because of the striking drop in hard drive costs and in part because of the inherent speed advantage offered. Hard drives in capacities exceeding 1 Tb (Terabyte) today are more and more affordable. With fast SATA (Serial ATA) interface drives of 250 Gb in capacity regularly below \$150 per drive, and with fast Firewire and USB2 interfaces for external models, this is a growing methodology.

Most data backup software systems support backup to hard drives, in addition to tape and other media types. As such, it has become entirely practical for the small firm to acquire a batch of five

or more external USB2 or Firewire-connected portable hard drives and treat them the same as if they were tape cartridges, for an alternating full nightly system backup. Portable drives such as those packaged in smaller chassis like those produced by Iomega, are most definitely light enough for off-site transport. Be sure to use a padded carrying case to protect the drives - cases designed to hold portable DVD players are ideal (for example, the Targus DVD001 (<http://tinyurl.com/9xj6m>)) or the ruggedized cases from Pelican such as the 1150 model (www.pelican.com).

There are four approaches in hard drive-based backup:

- 1) **External portable hard drives** - as previously indicated, from 80 Gb to 1.6 Tb in capacity, with either USB2 or Firewire 400/800 connections. The sweet spot in early 2005 would be those in the 250 Gb range, starting at about \$160 each. Look for lighter and smaller chassis units for easier off-site transportability. Most drive mechanisms are made by either Maxtor, Seagate or Western Digital although some may be Hitachi, Fujitsu or Samsung drives - all perfectly capable.

Same rules for multiple media apply - 3-5 drives at a minimum, more always being better for an alternating backup approach like the time-honored "Grandfather-Father-Son" rotation. This scheme uses daily (Son), weekly (Father), and monthly (Grandfather) backup sets. This approach begins with the daily backups. According to the online backup resource "The Three R's of Data Protection" (www.dlftape.com/ThreeRs) and equally applicable to both tape and hard drive-based backup:

"Typically, four backup media are labeled for the day of the week each backs up; for example, Monday through Thursday. Each tape is recalled for use on its labeled day. If only a one-week version history of files is maintained, then each tape is overwritten each week. In order to maintain a 3-week version history of files (recommended), more tapes are required. For example this week's Monday tape will not be overwritten for 3 weeks. (See When is it safe to overwrite for more information).

Weekly backups follow a similar scenario. A set of up to five weekly backup media is labeled "Week1," "Week 2," and so on. Full backups are recorded weekly, on the day that a "Son" media is

not used. Following the example above these would be "Friday" tapes. This "Father" media is re-used monthly. Five weekly tapes are required in order to maintain a one-month history of files, as some months have 5 weeks.

The final set of three media is labeled "Month1," "Month2," and so on, according to which month of the quarter they will be used. This "Grandfather" media records full backups on the last business day of each month. If your backup plan follows a corporate fiscal calendar, then your monthly tape will take the place of the week 4 or week 5 weekly/Father tape, depending on the month. If your backup schedule follows calendar months, then your monthly backup will vary throughout the year, replacing a daily or weekly tape. Typically, monthly tapes are overwritten quarterly or yearly (recommended), depending on version history requirements."

- 2) **Internal removable hard drives** - in this approach, there is actually a removable hard drive bay mounted in the network fileserver - internally. A good example would be the removable Data Express series of backup drive systems from StorCase (www.storcase.com/dataexpress_bu/backup_overview.asp). Individual hard drives (3.5") are mounted in removable "carriers." These slide into the mounting bay in the fileserver system. The connection system is either PATA (Parallel ATA) or the faster SATA (Serial ATA). These are available under the Kingston Technology brand as well. A set of a single "receiver" (starting at about \$110) and five "carriers" (starting at about \$70 each), ready to have hard drives mounted in each carrier, would start at about \$460, plus the price of SATA hard drives up to 400 Gb in capacity.

This approach offers higher potential data transfer speeds than external USB2 or Firewire 400/800 drives through the SATA interface method. Up to a theoretical 150 megabytes per second for SATA and up to 300 Mb/s for the newer, faster SATA II - class drives.

It goes without saying (but is worth emphasizing) that the portable drives should be stored off-site nightly and carried in a well-padded carrying case - the same cases as for the portable external drives should suffice.

3) **Other Removable Media - CDs and DVDs** - the quick answer is “no,” writable CDs and DVDs are not suitable for system backup. Why? (Yes, it’s a quiz!) Simple - CDs only hold about 680 Mb of data and DVDs hold up to about 4.7 Gb (or about 3.5 times that with dual-layer DVDs). The problems are two fold:

- * Simply not enough capacity on the media to perform automated full system backups.
- * Even if one conducted a manual backup, changing the media to a new CD or DVD when prompted to do so by the data backup software, when would you do this? After 5 PM? For a couple of hours? Or paying a staffer to do this? It doesn’t make sense, from a purely functional perspective, if not logistically.

CD and DVD can make be a very capable archival medium - perfect for smaller scale “spot” backups or “secondary” backup approaches. But they just do not work for full system backup, no matter how many ways proponents might try to argue their way around this one.

4) **Network Attached Storage (NAS)** - this class of network hard drive systems has been around for years. These devices can be attached to a Windows or Novell network and semi-miraculously configure and map themselves as an available network drive letter. The original product in this space was called the Snap Server - deriving its name from the idea that it could be connected and configured literally in a “snap” (actually true). Today there are a broad range of such devices. Most portable drive manufacturers offer some sort of networked version of their hard drives. Companies like Seagate, Buffalo, Western Digital, Maxtor and Ximeta all offer simplistic NAS devices.

But for network-smart systems that recognize and connect to Windows Server and Novell Netware systems, a step up the ladder is warranted. Snap Servers have become Adaptec Snap Appliances (www.snapappliance.com). With models ranging in capacity from 80 Gb to a mind-boggling 30 Tb, and just about everything in between, these devices can make an excellent secondary backup system. Why not primary? Probably too bulky to take off-site on a regular basis.

With Guardian OS and BakBone Netvault software, these systems are ready to serve as capable and powerful inhouse backup systems (<http://tinyurl.com/cg6rr>

for more info). For smaller firms, pricing ranges from about \$475 (street pricing) for the SnapServer 1100 model with 160 Gb of capacity to about \$1200 for a 500 Gb SnapServer 2200 model.

In short, great secondary internal backup system.

##AMANDA - CONTINUE BELOW FOR PART TWO:

Online / Internet-Based Backup - the proverbial 64 Gb question today is “why should I take responsibility for backup at all when I can outsource it?” After all, it is an indisputable fact that data backup is not only the single most boring technology topic to discuss, but it’s even more excruciating to be responsible for actually getting it done. So why not outsource and automate data backup to Internet-based data protection service providers? They store the data out of your office - already a promising capability; they handle backups automatically, per your scheduled approach which can take place in the middle of the night. Sounds great, right?

My opinion on whether law practices should consider Internet-based backup approaches is an absolutely, unequivocal, unwavering . . . maybe.

The positive aspects on online backup are seductive - unattended automated process means it gets done - no liveware failures in this process. The offsite storage issue is covered. Capacity is virtually unlimited. So what’s the catch?

The downsides of off-site backup are:

- * Potential ethical issues - do your state’s rules of professional responsibility address your ability to store confidential information with a third party, outside your control. This is probably not different than offsite storage of closed paper files, but it should be checked for each state.
- * Downtime in the event of server failure - it is especially important to enable secondary backup approaches if using online backup systems. This is because of the need to restore to either a new server, repaired server or temporary server - all could take time to locate and prepare. This can mean more downtime.
- * Risk - all the online companies are dotcoms. What happens when even the most solvent-seeming provider - who has 100+ Gb of your confidential firm and client data - ends up in Chapter 7? While all the companies have service-level agreements

that protect confidentiality, what happens with a successor entity - a U.S. Bankruptcy Trustee for example. Are they bound by the agreements the debtor entered into? Maybe. Maybe not. The point is that this risk must be evaluated.

- * Cost - the cost to store an entire server's contents, or multiple servers can start to add up rapidly. To explore the financial impact of online backup, several providers were queried. Let's explore the Q&A with these vendors.

Q&A with Online Backup Companies:

To further explore the viability of online backup, I went to the source, posing a series of questions to several online backup service providers. The questions were as follows:

- * Can you please describe your pricing structure?
- * Can you provide specific information on how you restore data in a situation where there is a failed server that is replaced with a new unit?
- * Also, do you backup Windows Server's Active Directory/Novell Netware's NDS (Netware Directory Services - both the core user/network resource settings files essential to backup in order to effectively restore and rebuild a server rapidly)
- * Do you backup Windows Server/Novell Netware server user and print queue setting?
- * Do you address backing up open files and Windows Exchange Server files?

EZVault.com responded by pointing out it provides very flexible pricing with many options to meet the needs of a broad set of customers. According to Richard Heitman, Director of Product Management at Evault, "we have a several hosted offerings which come in a "standard" Protect Edition and a Small Business Edition. The hosted offering is for customers that want to backup their data to our secure data centers. Pricing is based on the amount of data being backed up, and ranges anywhere from as low as \$50 a month up to several thousand dollars per month.

A key component of our overall solution is our DeltaPro™ technology, which speeds up backup operations and drastically reduces storage costs. Only new and changed data blocks within a file are backed up. All backup data is compressed and encrypted before being transferred and remains

that way on the storage vault, increasing security and reducing storage consumption (and cost). In addition, open files are backed up seamlessly, enabling backups to occur without interruption.

Our application plug-ins integrate seamlessly with critical business applications, and we have a plug-in designed specifically for Exchange 5.5, 2000, and 2003. Admins [Network Administrators] have a lot of flexibility in choosing exactly what data in Exchange to backup and restore.”

The emphasis in the answers provided was on the backup process, rather than the more critical restore process. More on this point at the end of this section.

Jim McManus from Sparksource.com indicated that his company’s “LiveVault InSync starts at \$119 per month. An example: \$199 per month per server for up to 5 GB of storage, based on a 30-day retention plan with no set limit on historical levels.” McManus responded to the query about restoring data as well, saying that “to restore to a new server, simply install LiveVault agent and request a restore. LiveVault updates server to match restored images including System State and all Applications and data.” Further, regarding open file backup and backup of Microsoft Exchange Server and SQL Server databases, McManus noted that “LiveVault backs up all system data including System State and Active Directory. LiveVault continuously protects all chosen files on a server – this is true of user files on a file server as well as open and active databases like MS Exchange and SQL server.”

Larry Sanders, Tech Support Lead from Remote Data Backups, Inc. (www.remotedatabackups.com) said that to address a failed server that has been replaced by a new system, with his service one would “install Remote Data Backups software on your new server. (You will need the account number, encryption key, and the password if used. [Then] select the data you wish to restore, click on “restore now.”

Regarding the question of open file backup, he indicated “our software does not backup open files, but it can be used in conjunction with any backup software that does (for example Veritas).”

Susan Klees, Vice President of Data Protection Services, LLC (www.dataprotection.net) provided very detailed information. According to Klees, “for the initial backup, all data must be transmitted. For example, if we assume a 50% compression, the initial backup of 100 GB of data would be 50 GB. All subsequent backups only add the changed portions of changed files. The difference between DPS backups (which we call SmartFull backups) and incremental backups is that every DPS backup is a full backup (there is never a need to “build back” to a certain backup or date). So the client gets the benefit of every backup being a full backup, yet adds volume like

an incremental backup. Again, if the first backup is 50 GB and the client is doing a typical grandfather-father-son retention (keeping daily backups for a week, weekly backups for a month, and monthly backups for a year), then at the end of a year (with 21 backups stored in our vault - dailies, weeklies, and monthlies) we would expect a typical client to be storing 53-55 GB.”

Klees points out that “the DPS backup is a "scheduled" backup, not a continuous backup. The difference is that continuous backup services also require continuous processing of your computer and continuous drain on your bandwidth. With a scheduled backup, backups occur only when you schedule them - whether that is daily, hourly or every few minutes. In my opinion, scheduled backups are far superior.”

As to the costs, Klees indicated “there is a one-time license fee of \$450 . . . and the license is good for as long as he remains a client. The monthly cost under your 100 GB scenario would depend on compression. If I assume a 50% compression, the cost would range from \$640/month to \$800/month. This is for a fully redundant solution, with data stored in vaults in two separate remote Tier 1 data centers.”

The data restore process - what happens with an online backup. According to Klees, “when there is a fully dead server, the client would . . . obtain another server with a minimal operating system so that Internet access could be acquired. The client would then download and re-install the DPS software from our website, go through a process we call “synchronization” (which is mostly an authentication that the new server is a valid receiver of the data), then select what data should be restored and click "restore." It's as simple as that. We also have the ability to restore to other media (CD's or DVD's) and, for large volumes of data, we can ship to client on a Mobile Vault. Our Rapid Response Team can prepare and ship (via any method the client chooses - including same-day air) within hours of receiving the request.

DPS is clearly a premium service with broad flexibility for different server systems and in its technical capabilities. DPS backs up Windows Server Active Directory systems, as well as Novell NDS network information. A “plug-in” allows backup of Exchange Server and SQL database files, which tend to remain open and active.

Two other companies, Connected.com and Netmass.com were queried but did not respond.

Summarizing online backup, presuming the local ethical hurdles are cleared, online backup may be an effective data backup option. However, for a small firm for full system backups, the costs may be prohibitive. In the case of Data Protection Services, a small firm with 100 Gb of information, and presuming 50% compression of the files, would incur costs of approximately \$7680 to \$9600

per year. Add to this staff time to coordinate period test restores to confirm proper backup operation - about 15 minutes per week or 13 hours per year. At \$20/hour, that adds \$260 annually.

In a comparison to local backup done in one's own office, the costs of staff time do have to be considered to be fair. The small firm might spend \$1500 on backup components including software and media. If a staff person spends 5 minutes each workday removing and replacing media, and updating a backup log sheet, that translates to about 1250 minutes per year or about 21 hours. Add weekly "mini test restore" sessions to confirm the backup is working adequately at approximately 15 minutes per week - 780 minutes, or another 13 hours per year. The total cost of 34 hours of staff time based on a \$20/hour average wage is \$680.

That yields a marked cost differential in favor of local backup. However, other factors do come into play that favor online backup. These include the practical consideration of reliability - the online backup will occur as long as there is an active high-speed internet connection. It is not dependent on a person remembering to insert the backup media each evening - a potential point of failure in the local backup process. Another element is the sophistication and reliability of the equipment - the online backup provider's data center undoubtedly incorporates far more sophisticated and extensive system technology than the small law firm could have in place.

But an area of significant concern is one that should be considered paramount - the most important aspect of the entire data backup process: restoring a failed system. After all, the entire point of backing up is not just backing up - the point is to be able to restore when the chips are down. This subject tends to either not be addressed directly or glossed over by most online backup advocates.

An effective restore process, particularly for smaller firms who may not have internal IT resources at the ready in a tight circumstance, should take place as quickly and easily as possible. The goal is to NOT have to rebuild and reconfigure a network fileserver from scratch prior to being able to run the data backup software to restore the system. Such an effort takes up valuable time - downtime that can significantly disrupt a small practice with substantial economic costs. Instead, the goals of any competent local backup/restore process should be twofold:

- * Not having to reload and reconfigure a network operating system from scratch prior to being able to activate a restoration process, and
- * Absolutely minimizing downtime causing lawyers and staff to be disrupted and heavily inconvenienced by lack of access to their electronic files.

With local backup with a “disaster recovery” option that most backup/restore software includes, the process normally goes as follows:

- * Either using a loaner fileserver or on the newly repaired server, one boots the server from the disaster recovery CD made in advance by the backup/restore software (usually at the time of initial installation)
- * The recovery CD loads enough of the network operating system and the backup/restore software to begin the full system restore.
- * If restoring from hard drive media, this goes typically much faster than tape, but tape (especially faster Ultrium-class units) works too.
- * Typically with a smaller firm with perhaps 20-30 Gb of information, this process might require two or three hours, including operational testing time once the restore has completed before the fileserver is up and running as if nothing had happened.
- * In the meantime, with proper advance planning employing techniques such as local document syncing that apps like Worldox accomplish, local syncing of case managers, etc., lawyers and staff would have been working locally during this period so downtime could be virtually nil.

The point is that with planning and a proceduralized approach, local backup and restore is not complicated. Further, it can include a strong focus on downtime minimization to reduce the economic cost of idled people.

With online backup, how would a full restore be accomplished? What if you couldn't get a fast Internet connection that day - Murphy's Law tends to strike at the worst possible moments - situations just like this when the more you need high-speed access to survive, the less likely it is to be working well. Then how long would it take to transfer a full system? How would the fileserver be booted to even connect to do what we're talking about if there's a virgin drive in it? Those are questions that never seem to be directly addressed by online backup companies. Practically speaking, it may make sense for small firms to consider online backup services as secondary, rather than primary backup sources, protecting data folders and not attempting to protect entire network file servers, their operating systems and all applications.

Simple is better than complicated . . . If you apply that mantra to backup here's how one might see it play out:

- * Online Backup - simpler for daily functionality since it can happen without any user intervention in most cases
- * Local Backup - more complex for daily functionality since someone needs to change the media and periodically run a test restore
- * Online Restore - potentially much more complex, if not practically impossible in a reasonable timeframe (for full system restores)
- * Local Restore - potentially much simpler when a disaster recovery option is in place and a bootable CD is used that will automate the process without any dependence on an available high-speed net connection

So there's the balance: Online Backup = simpler, Online Restore = more complex Local Backup = more complex, Local Restore = Simpler

With that said, if the priority is to have an easier backup process, by all means, opt for online backup. But if the priority is restoring systems, then there is a clear advantage in favor of local backup (with the option always available to use online backup, if ethically permissible, as a secondary backup methodology for data sets).

Special Note - Intriguing Secondary Backup Approach:

There is a specific new category of product that deserves special mention as a secondary backup methodology - low-cost shared network storage devices. Offered by a number of companies including Western Digital, Maxtor and others, these devices are relatively simple conceptually. They consist of an external hard drive with an Ethernet network connection and data backup software. Plug the devices into a network and they appear as a shared network drive. Data backup software can then perform full, incremental or data-only backups of the network server or any PCs on the network system.

Why should these devices be considered only as a secondary backup methodology? The reason should be clear - no off-site storage. These devices remain attached to the network system so they

would not meet the “best practices” standard for off-site daily storage of the most recent media that is a cornerstone of a primary backup routine.

One device rises above the rest in this category - the Mirra Personal Server (www.mirra.com) is no mere external hard drive with networking pretenses. Rather, it is a complete “embedded” PC system, with its own processor and a cleverly hidden Linux-based operating system in a chassis only slightly larger than a typical external hard drive. What this device and its proprietary software accomplishes is not just “scheduled” backups, but rather, continuous automatic system backups that protect files as they are created or altered in real-time. Once an initial full backup is done, the Mirra software keeps the backup up to date. Multiple PCs on a network can be protected in this manner.

Further, the Mirra software allows for secure web-based access to the files it holds. It also allows data sharing - great for applications like access to a large discovery database. There is also a multiple PC synchronization function - providing the ability to keep key folders synchronized on multiple PCs. An interesting application would be ensuring that multiple law firm laptops all have a complete document library before they disconnect and hit the road. This product is currently being tested - it is successfully backing up four different PC systems, automatically, through the Kodner clan’s wireless home network.

Available in capacities including 80, 160, 250 and 400 Gb, the Mirra Personal Server 2.0 costs between \$200 and \$675 “street pricing.”

Image-Based Backups - this approach involves using “bit-image” software (the best known of which is Symantec’s Norton Ghost series) to create an exact image of a hard drive. Modern iterations of such systems have the ability to restore entire drives in one fell swoop, or more flexibly, restore selected files. Can this work as a backup methodology? I see Image Backups as a secondary method - a great way to preserve the setups of individual workstation classes - a laptop setup and a desktop setup. This can allow you to quickly rebuild a “blown up” system after some errant software misbehaved or your partner’s 12 year old “optimized” Dad’s PC on a Saturday morning for playing “Half-Life.” Solid secondary approach - not intended for primary system backup.

Interesting, the primary Ghost competitor had been Powerquest’s Drive Image Pro - it’s now a Symantec product alongside Ghost. Another entrant into this market segment is True Image from Acronis (www.acronis.com) - it’s winning the favor of übertechies.

Taking Your Chances and Hoping for the Best - No. Just no. Don't even think about it. Unless of course you want to be the first lawyer in recorded techno.history to lose their license for failure to take ANY kind of care of confidential client information and electronic files.

It's All About the Software

None of the internal backup approaches works without software controlling the process. There are data backup products intended for individual computers and smaller server-less peer-to-peer networks. There are also server products intended to backup real networks.

Individual Backup Systems - there are a number of well-known products with good reputations for protecting and restoring individual PC systems. These include Backup MyPC 2005 (from Stomp Systems - www.stompsoft.com/backupmypc.html) for \$59.95. Backup MyPC was formerly the "personal" edition of vaunted network backup titan, Veritas Backup Exec. EMC Dantz's Retrospect Professional 7 (www.dantz.com/en/products/win_personal/index.shtml) for \$99.95 which is also available in a Macintosh edition. Otherwise lesser known but capable personal backup systems include Newtech Infosystems NTI Backup Now! 4.0 Suite (www.ntius.com/default.asp?p=backupnow/bun_main) for \$79.95 and Novastor's NovaBackup (www.novastor.com/pcbackup/backup/n_backup.html) for \$49.95.

All these systems support the important system registry backup and restore functions - the entire reason for using actual purpose-built data backup software instead of merely trying to manually copy files to another medium. All support backup to hard drive, tape, and removable optical drives.

Network Backup - for backing up Windows Server, Novell Server, Linux servers. These heavier-duty systems have the internal "smarts" to properly backup all the key core program settings and operating system components. With an array of additional "agents," these systems can also leap tall servers in a single bound - backing up open files, running Exchange Servers SQL databases, both Windows Active Directory and Novell Network Directory Services (NDS) files, Lotus Notes databases (okay, it's a bigger firm and corporate legal department thing, but it does exist) and even specified workstation hard drives.

The two dominant titans in this arena are Veritas Backup Exec, available in various iterations from the bargain bundle "Small Business Server" edition to the larger scale

Enterprise Edition. The former includes a number of “file agents” in a value-priced bundle for users of Microsoft’s Windows 2003 Small Business Server edition (\$499 street pricing). Info is at www.veritas.com/Products/www?c=category&refId=148.

The arch-rival is the venerable Computer Associates Arcserve (info at www3.ca.com/Solutions/ProductFamily.asp?ID=115). As with Backup Exec, there are versions for Small Businesses (Windows and Novell) as well as full-blown Windows, Netware and Linux/Unix servers. The Arcserve Small Business edition includes a number of file agents, as does its Veritas counterpart (also about \$499 street pricing).

Both are very capable products that have long track records of rescuing beleaguered law practices from the brink of total disaster. In many ways, the decision between the two products may very well be based on the available support from a local network integrator, who will invariably prefer one over the other. The capabilities are an effective “wash.”

There are other network backup products - Novastor, NTI and EMC Dantz all offer network backup systems as well - but all are far less well known products in this segment than either Arcserve or Backup Exec . . . and accordingly are relegated to second best status because of a general lack of local integrator support.

Summary

Backup isn’t exciting to think about or talk about, much less actually have to do. However, it is the one technology that will one day save your law practice as a business. Remember the basics: full backup - no shortcuts, regular test restores, alternate the media, replace the media periodically, never rely on just one backup method, think about downtime reduction also, think about multiple people being responsible. But by all means, think about it - regularly, prove to yourself it is actually working to protect your practice and your clients’ confidential information. Because ultimately, the “best” backup system is that you actually use. So start practicing safe computing today.

Ross L. Kodner, an attorney, is founder and senior consultant for 20 year old MicroLaw, Inc., helping law firms and legal departments understand, select and better use technology in the workflow of their practices. He can be reached at rkodner@microlaw.com and 414.540.9433.