

DAW Data Protection - Overview

Why do I need a data protection plan?

The essential task of data protection is to duplicate, and thus safeguard, your data. This process should be accomplished across multiple devices and/or media to be most effective. Each added redundancy increases assurance that the data you need can be made available if the original source of data becomes inaccessible. Your protection strategy should cover the full range of availability. From the occasional need for a “good” copy of a single corrupted file, to the loss of all files due to hard-drive failure or natural disaster.

Think of the time you spend in the process of creating computer-music. You sweat and toil, revising and updating, until your artistic vision is realized. As you work in a digital medium, all of this labor is saved in complex patterns of zeros and ones to computer files stored on hard-drive. These files can be easily retrieved; updated, relocated... it's a happy digital world until that hard-drive fails. This failure is not a distant possibility, but an inevitability. All hard drives, no matter how expensive and well designed, will fail. Your personal user files and the installed programs and data that help create them, can be gone in an instant, without warning. As traumatic as this scenario sounds, it is one you must contemplate in order to build a solid data protection strategy.

Without Data Protection....

If O/S drive failure: you will need to manually re-install your Operating System, re-do its customizations, then re-install all your application programs. This task could be multi-day. With an image backup DP strategy, your complete restore to the replacement drive should take little more than an hour.

If Streaming-Data (Virtual Instruments) drive failure: you will need to manually re-install all VI applications that do not allow for separate breakout of data-sets, then re-copy data-sets from cd/dvd supplied by the applications that do. This task could take a full day or more. With a drive-to-drive backup DP strategy, your complete restore to the replacement drive should take no more than a couple of hours.

If User-Files drive failure: your user-created personal files are in serious jeopardy of being irretrievably lost. The cost for drive data-recovery (from a professional service with “clean room”) is in the thousands of dollars and there is no guaranty of success! With a drive-to-drive backup DP strategy, your complete restore to the replacement drive should take no more than a few hours.

How much duplication is enough?

For protecting O/S or streaming-data, two levels of duplication should be sufficient. For user-created files though, there should be three levels of duplication, on different devices or media, as insurance in case one level fails. These levels represent: session, standard, and off-site duplication, with the last as emergency security. The off-site data set should be kept away from your studio at a friend's house or even a safe deposit box. If local fire, etc., wipes out your DAW and data – this is your ace to rebuild.

What basic tools are recommended?

Essential DP tools, regardless of DAW rig.

- DVD recordable device and “burn” software.
- External drive or multi-drive unit (USB, Firewire, SATA).
- Imaging and traditional backup software. (See Addendum).
- Backup Log: hardcopy record of all backup/restore operations, indicating date, media type, and relevant notes. Documenting to a visible log acts as both reference and reminder. Place this on a clipboard in your studio, within plain view.

PROTECTING YOUR DATA

A simple and universal solution for level-1 user-files duplication is to copy, at the end of each session, all newly created or modified files from their original locations to a single designated folder on a different hard-drive. When this folder is near DVD capacity, burn that folder’s files, then clear its contents to repeat the process with new data. This process is facilitated by naming your created files with a version# suffix, a good general practice that helps to insure more uniquely meaningful filenames. This procedure can be used for any DAW system’s user-data. An optional approach is to deploy file-synchronization software to automate this process as a one-way sync to another hard or flash drive, automatically updating before shutdown to minimize disruption to DAW processing.

The remaining solutions are specific to individual DAW configurations and are presented, sorted by cost, from least to the most expensive. A “tools” list follows each solution, with general cost estimate, but without recommendation for any particular manufacturer’s hardware or software – that is left up to you.

Solution1 - Local & External Hard-Drive

If you have a multiple hard-drive (non-raid) DAW you can use some of its storage on one drive to hold backup data from another drive. Combined with an external “Supersized” dedicated backup drive (SATA/ATA), you can put together an effective and comprehensive data protection solution. You will need Imaging software to “clone” your O/S drive, storing its data-set on another internal hard-drive. Use traditional backup software to periodically record its new and modified files, backing them up to the same location as the image set. Most DAW applications store user-configurable settings and presets at “C:\Program Files” or “\Documents and Settings”. These customizations should be safeguarded. There are single applications currently available that can manage both imaging and traditional tasks. Having both types of tools saves the time and trouble of exclusively running full images. Traditional backup can secure all the data on your remaining drives by archiving to your dedicated external backup drive. If your DAW has a high capacity DVD writer, use it for off-site backup of user-data. Otherwise, use another single-drive external enclosure to make full periodic copies. Place the entire enclosure in a sealed moisture-barrier/electronics bag along with a de-moisture pack such as “silica gel”. This sealed, enclosed, drive should be then kept off-site. This DP

approach is the least expensive reviewed in this guide, with total hardware/software costs as low as a few hundred dollars.

Tools: Image & traditional backup software, USB/Firewire drive enclosure + supersized SATA drive. (Optional) 2nd drive enclosure + standard drive (see Addendum).

I have authored a detailed guide based on the above data protection solution, specific to a 3-hard-drive XP DAW, available as a free PDF at my website "<http://www.daw-solo.com/FYI>".

Solution2 – Removable Drive

A single drive-unit is effective for the full range of data protection as long as its media is removable. Tape systems, once dominant, are still excellent backup/restore solutions. They have the advantage of inexpensive media and mobility. A single USB DAT/DLT backup can service multiple computers. The disadvantages are relatively slow performance compared to hard-disk based systems, unless you invest in high-end DLT that can cost thousands of dollars. High-capacity DVD, with Blu-Ray at the lead, has similar performance to standard tape, can also be mobile, but has equivalent-capacity media priced at twice that of tape. One advantage to Blu-Ray is that a single device can burn media from CD-R to BD-RE. Additionally, the price for Blu-Ray devices and media should continue to fall. A general disadvantage to single-drive tape/DVD systems is that, because of their relatively smaller capacity, larger backup jobs will be "spanned" across multiple media, requiring periodic manual eject and load.

A third removable option is hard-drive chassis, where a hot-swappable drive is inserted into a designated USB/Firewire or ATA enclosure, allowing for high-speed backup and restore. Sync software could be used to effectively "mirror" files from your user-files drive to such an enclosure, providing your second level of data protection. A big advantage to any removable media is that it greatly simplifies "off-site" backup, as you can periodically rotate out a media-set for this purpose. You will still need both imaging and traditional backup software for the protection of your other DAW drives, storing the image-sets to removable media. This DP approach can service multiple DAWs but is moderately expensive. Total hardware/software costs can run from under one thousand to around fifteen hundred dollars.

Tools: Image & traditional backup software, USB/Firewire DAT or DLT tape, Blu-Ray DVD external drive + blank media, or hot-swap chassis + prepared drives (see Addendum).

Solution3 - External Disk Array

Previously only found in corporate environments, affordable disk arrays are now readily available, offering high-speed backup/restore via USB/Firewire/SATA connection. Disk Arrays are housed in a multi-drive enclosure that appears to the O/S as a single logical hard-drive via RAID implementation (Redundant Array of Independent Disks). Two of the most popular uses of this technology are RAID 1, also known as mirroring and RAID 5 or striping with distributed parity. Mirror RAID is a powerful tool for the data protection of a mid-to-large capacity DAW. It automatically duplicates the latest data identically between the two drives involved and so continually manages one level of data protection in real time. In addition, array boxes are "hot swappable" in that a single

failed drive can be replaced, while the array is running, and its data automatically rebuilt. This makes for easy off-site backup, with the option to simply cycle out one of the mirrored drives in its hot-swap bay. Because of the mirror process the actual storage capacity of RAID 1 is half that of the drives used (which should be identical in specification).

A mirror array can be a very effective dedicated backup device, but if used exclusively as your user-files “drive”, with a spare hot-swap drive/bay, it can be nothing short of revolutionary. RAID 1 used in this manner allows for all three levels of data-protection to be achieved without the use of any software backup application! Level-1 backup would be met by the previously described copy technique at the end of each session. Level-2 would represent the real-time duplication offered by your mirror. Level-3 would be achieved by periodic swapping of a mirrored drive with an off-site “spare”. You’re safe as long as you keep one good “mirror” off-site, capable of auto-rebuilding to a replacement array in a disaster recovery scenario.

- 1) Keep all original user-created files on the RAID-1 enclosure, providing real-time duplication of its data via "mirroring".**
- 2) On at least a quarterly basis, manually eject one of the mirrored drives, replacing it with your "spare". Let the data auto-rebuild overnight. ***
- 3) Place the recently ejected drive/bay in a sealed moisture-barrier/electronics bag along with a de-moisture pack such as "silica gel". Take this (spare) to your off-site location. Periodically repeat steps 2 - 3.**

This radical approach to user-files backup/restore works only with RAID 1, because only a “mirrored” array can allow complete regeneration via a single drive. Respected DAW builders: PCAudioLabs and Sonica, replied affirmatively to an inquiry sent them regarding the feasibility of this DP approach and the suitability of a RAID 1 array as a recording drive. Caution was advised though in use of anything less than RAID 5 in a studio environment where time is (client) money. A more conservative off-site backup procedure could be employed by copying or syncing to an external hard-drive or removable. Alternately, a “RAID 1 + Backup” enclosure comes with a 3rd bay specifically designed for separate hardware backup of the mirror. By periodic swapping of this backup drive you can achieve a more failsafe, though more expensive, means of “software-less” user-files protection. A general limitation to RAID 1 is its capacity, with current ceiling at 1tb.


** **Note:** rebuild speed at about 100gb per hour. Length of time between “spares” represents your window of acceptable data loss if worst-case disaster requires rebuild of array solely via off-site spare. Invest in a quality raid unit! Its hot-swap and auto-rebuild capabilities are critical to DP functionality and must be reliably first rate.*

If RAID 1 capacity is an issue, you can move to RAID 5, requiring at least three and usually five disks for the array. This level of RAID provides both data and error correction striping across all drives, resulting in good performance and fault tolerance. If any single drive fails, all data is preserved and automatically rebuilt to the hot-swappable replacement. With identical drives in the array, capacity is determined by total drive-size minus one drive. If using 500gb drives, then a three-drive array yields 1 terabyte, while a

five-drive array yields 2 terabytes of storage! This should be sufficient capacity for a small network's recorded files or as its dedicated backup. For a RAID 5 system, you will still need high-capacity DVD, DLT tape, or additional external drive(s) for your off-site backup. For quick historical access to old-version user-files, you have your "session" DVDs, date-marked and referenced by backup log. This DP approach offers huge storage and can service multiple DAWs, but can also be relatively expensive. Total hardware/software costs are well under one thousand dollars for RAID 1, but well over two thousand for RAID 5 and tape.

With your user-files protected by a RAID solution, you will still need strategies for protecting your other data. Your O/S and programs are best handled by an imaging application, doing a complete sector-by-sector "clone". Streaming-data "samples" can be safeguarded effectively by any traditional backup software running a Full/Differential plan (see addendum). Even the basic backup application offered on your Operating System should handle this easily, and at no additional charge.

Tools: Image & traditional backup software, USB/Firewire/SATA RAID (1 or 5) disk array enclosure + matched SATA hard-drives. (Optional) 2nd drive enclosure + drive (see Addendum).



Addendum: above (left) is a graphical representation of RAID 1 "mirroring" followed by RAID 5 "Striping with distributed parity". Below is a chart comparing various attributes of the media covered in this overview. I have tried to come up with "real world" statistics that could provide a solid basis for judgment. Capacity and speed ratings are for a single device without data compression.

| | <i>Capacity</i> | <i>Cost per GB</i> | <i>Backup Speed</i> | <i>Data Retention</i> |
|-------------------|-----------------|--------------------|---------------------|-----------------------|
| DLT | 160gb | 25 cents | 40gb per hour | 20 years |
| SSD | 128gb | 200 cents | 100+ gb per hour | 10 years |
| Blu-Ray | 50gb | 55 cents | 20gb per hour | 5 years |
| DAT | 30gb | 50 cents | 10gb per hour | 10 years |
| Flash | 16gb | 250 cents | 50+ gb per hour | 10 years |
| DVD+r | 4gb | 10 cents | 20gb per hour | 5 years |
| Hard-drive | 1tb | 15 cents | 100+ gb per hour | *(see notes) |

* modern large-capacity hard drives are designed primarily to store data while spinning, though If used in periodic backup rotation they can provide many years of service via USB enclosure, etc. Drive backup to SSD is preferred, especially as off-sight media, because of superior data retention "life span".

Using Both Imaging and Traditional Backup Software

Image backup is sector-by-sector, effectively enabling an exact “clone” of your hard-drive. While this provides the most reliable means of cloning or restoring an OS drive it does have the limitation of insisting that the restore target boot under similar hardware, especially motherboard/chipset. It is also unaware of the file “archive” bit which is essential to traditional backup software applications.

Traditional backup is file-by-file, utilizing the “archive” bit and other file-based info to check each file’s backup status. The O/S drive requires both imaging and traditional backup strategies to insure its complete data protection. The first provides a sector-by-sector base, while the second adds the ability to include modified files to insure the latest DAW customizations and backup catalog are restored atop the clone. Other DAW drives can be protected using traditional backup software alone.

The three basic types of backup/restore offered by traditional software.

1. Full: backup of all selected files (often targeting an entire drive).
2. Differential: backup of all selected files that have been modified or newly added since the last Full.
3. Incremental: backup of all selected files that have been modified or newly added since the last Full or Incremental.

Software Restore and Copy Protection

Copy protection mechanisms should only pose potential problems in restore of the O/S (Programs) drive.

BEST: USB Master-dongle, such as used by Steinberg or VSL. This one (Syncrosoft) USB dongle can authorize all compliant manufacturers software products. When invoked, these applications will look to the USB dongle for access authorization. In the event of OS&Programs drive failure, an image of this drive can be restored to a new drive and all the residing master-dongle apps should function perfectly without need for software re-installation.

ACCEPTABLE: Challenge/Response variant of "signature" based copy protection, such as deployed by Spectrasonics. With this method the "Response" code is the signature authorizer, based on the Challenge code issued during software install. This approach allows for cloned applications to function perfectly on their newly imaged drive, as long as the Challenge/Response procedure alone is redone.

WORST: Hard-drive "signature" (as used by “brand X”). This creates a unique reference from low-level hard-drive info concatenated with application serial# to create an authorization reference for each application. The purpose of this approach is to prevent software piracy. Unfortunately, it severely cripples cloning/imaging as a legitimate means of data protection for fully licensed users. In the event of OS&Programs drive failure, an image of this drive when restored to a new drive will disallow access to any "signature" based application. The user will be forced to uninstall and re-install such applications. A more humane, “lite”, version of this copy protection is used by Native Instruments, where application serial# is combined with hardware component info, such as motherboard, to provide a unique reference. Simple drive replacement should not be a problem here, as only major changes to low-level hardware require full re-authorization.

John O’Mahoney

DAW-solo.com © 2008

V1.1.3 (11/30/08)

Disclaimer: I offer this guide as a free service to the DAW community. Use it at your own risk. I make no guarantee or warranty about its results, though I have made every effort to carefully research its subject material.