

ENTERPRISE BACKUP SOLUTIONS



A LOOK AT SOLUTIONS FOR THE MAC OS X PLATFORM

WRITTEN BY: MICHAEL DHALIWAL

SIXTH EDITION - EARLY 2009

miked@district13computing.com

Table of Contents

Preface	2
Introduction to backup	3
IBM's Tivoli Storage Manager	8
Tolis Group's BRU	16
Atempo's Time Navigator	21
BakBone's NetVault	29
Archiware's PresSTORE 3	36
Conclusions	43

Disclaimer: Apple Inc. is not responsible for and does not endorse the content of this document. This document is solely the work of Michael Dhaliwal

Preface

Who should read this document...

This document is intended to help anyone who is in need of a backup solution for Apple products and servers. This document is in no way meant to be exhaustive, but merely a detailed starting point for those of us out there who are feeling a little overwhelmed by the idea of implementing a data recovery plan. As always, your own experiences take precedence over any findings in this document. There is no wrong answer to data backup and recovery, just some that may work better for some of us than others.

Why did the author choose these products over others?

The products described in this document are the dark horses, if you will, in backing up Apple enterprise solutions. For years, we've all heard that Retrospect is the only way to backup your data. Due to the fact that Retrospect is so widely used and known, it is being omitted from this document, for now, at least. Also, the omission of a command line interface in Retrospect has led many administrators to look elsewhere, as they wish to exploit all of the benefits of the Mac OS X Server platform. This, in no way, is meant to say that Retrospect is a bad product, but being a long time solution in Apple environments, it's hardly a dark horse.



About the author...

Michael Dhaliwal is a systems engineer in the Chicago area. Having worked in mixed platform environments in enterprise, government, non-profit and educational settings, Michael has a breadth of experience in creating not only cost effective solutions, but highly integrated and scalable ones as well. Michael is an Apple Certified System Administrator and runs his own consulting firm, District13 Computing (www.district13computing.com), specializing in global Mac OS X deployments. You can contact Michael at miked@district13computing.com.

Thanks...

To all of the software vendors, my friends, family and colleagues. Special thanks to Karen for being so incredibly supportive and understanding of exactly why our home looks more like a data center than a home and my use of every waking moment working on projects like these, my mom for getting me my first Apple all those years ago and Smitty for reminding me to take breaks by parking herself on my keyboard.

Introduction to backup

Where to start and some things to consider.

Mac OS X Server coupled with Xserve, Xserve RAID and Promise RAID technologies have thrust Apple into a place that many people never expected to see them; the enterprise. The days of seeing Apple exclusively in educational, creative and consumer markets are gone, as the platform that promised to “Think Different” has forged ahead into IT departments around the world.

As Apple continues to raise the bar of performance and packs more storage into each generation of their products, the question continues to come up time and time again: how am I going to back up all of this data? After all, your data is only as reliable as your backup strategy.

This article is not meant to be the end all guide to backing up Mac OS X Server, nor does its author guarantee that your results will match his own. Use this article to help guide you through this process. Consider this as a way to jump start your own solution. Of course, your comments are more than welcome, as they could be used in future versions of this document.

When you begin to devise a backup strategy, there are many questions you must start with. First, you must evaluate how much data you have, how much you need to backup regularly and how much time you have to do it in. These questions will help guide you through this process. You then must decide on an effective strategy to keep your backup safe. Most companies decide that keeping data offsite is critical, in case of disaster. Some people feel that just having duplicate copies of critical data is enough, though this doesn’t really guard against all types of disasters that could occur.

Now that you’ve discovered your backup needs, congratulations, you’re done! No, unfortunately, you’re not. In fact, you’re actually just starting, so get comfortable! You have to decide what type of media you would like to back up to. This will be decided by numerous factors. You have to consider cost, media durability, the amount of data you need to write and even your timeframe, again. If you are a small company, with perhaps one server running Mac OS X Server and you have less than 500GB of data, you might be able to get away with using Mike Bombich’s Carbon Copy Cloner writing to large FireWire disk. CCC has a very easy to understand interface and, in my tests, has been very reliable in creating daily backups of data. The new version even includes synchronization capabilities built into the product. In my testing, CCC was used to backup about 1.5TB of live data from Xserve RAID to LaCie 1TB disk, over FireWire 800. For data servers of this size, you won’t want to settle on this solution, as the price of the FireWire drives will very quickly outweigh the cost of a tape solution, if you want to properly rotate media. Of course, that doesn’t even begin to discuss the abuse media takes being moved by traditional, off-site data storage companies.

So, you've now realized that you purchased a 10TB RAID, that's already full, and you have Windows on your network as well. Well, Carbon Copy Cloner, while a great product, just won't cut it. You are now in the enterprise and you have to play with the big boys. You probably want to consider writing to tape for off site recovery. Note that I said for off site recovery. While many people in the industry want to discuss such terms as "online", "nearline", and "offline", the cost of disk is getting so low, that, unless it will cost you many terabytes of storage, you may want to consider keeping your data online. We did learn something very important in our first example: copying to a second disk is a good thing! I'd like to note that backing up to disk isn't necessary, but it is in good interest. There is no recovery faster than being able to simply replace a file from a backup RAID than finding and loading tapes.

Since we've decided to move out of the FireWire disks and CCC, we'll have to look elsewhere for our backup device and our software. Here's where you'll find some tough choices to make.

For most of us, we'll want to back up to tape and we'll think that the tape drive is the most critical part of our decision process. After all, hardware is more expensive than software and most companies will want to get more longevity out of their hardware solutions than software, which can easily be upgraded or patched. While you'll find numerous devices that range in scale from a few thousand dollars to many thousands, you'll also learn one interesting thing about Mac OS X Server along the way: Mac OS X Server does not have any built in support for tape devices. You'll find that the backup software manufacturers are actually responsible for all of the backup and tape library functions, which will be a very different concept for many of you. In this scenario, the drive will only behave as well as the software will let it.

Currently, the main players in the Mac OS X Server backup game are Bru by the Tolis Group, Time Navigator by Atempo, NetVault by Bakbone Software, Tivoli Storage Manager by IBM, PresSTORE by Archiware and Retrospect by Dantz. The most well known of this group to the Mac community is probably the Retrospect product, by Dantz, which was acquired by EMC.

Since you now understand a bit of the background of how backup takes place in Mac OS X, let's explore a bit about media rotation strategies as well as data and information lifecycles. Since we've moved to tapes, you'll want to keep these concepts in mind as we delve the intricacies of the software that will drive your solution.

As part of sizing your solution, you need to make a decision on how long you must retain your data. This concept is part of Information Lifecycle Management and helps you assign a value to your data, but requires specific application knowledge of the data as well. You might ask yourself why you might reduce the value of your data at all. Your thought being, if I backup the data, it must be important to me and my company. While this is true, over the lifetime of any given set of data, the value of it can begin to diminish. Many different factors can contribute to adjusting the value of data. In the most simple of ways, you might simply no longer need the data. Maybe the data is now obsolete or you have used that data for a revision and now no longer need those

original drafts. This concept is often confused with Data Lifecycle Management, which is more geared towards where your data is stored and how to move it between your tiers of storage. The Data Lifecycle will help guide you through placement of your data based on hardware and performance factors, as well as availability. Confused? You are not alone. These two terms have been used interchangeably for quite a while and its only been recently that folks have tried to differentiate between the two terms and assign them their own meanings. A good way to envision this is that DLM is a more simplistic version of ILM. Think about the actual words being used, meaning “Data” versus “Information”. The term “information” would be a more finite, granular knowledge, while “data” is more general. You can have terabytes of data, without any knowledge of the actual information about it. DLM is simply talking about files, while ILM is based more around the content of those files.

With all of that hopefully understood, we now will actually look at media rotation strategies. You’ll want a strategy that will properly protect your data and, hopefully, one that is able to protect multiple versions of your data as well. This is the first dividing line of backup strategies, in my opinion. Just as we discussed Carbon Copy Cloner as a way to take a live backup of your data, in the truest sense of a clone, we must also discuss how you plan to keep versions of your data, in backup form. What you need to ask yourself is if your environment requires more than just a weekly backup media set. If you’re in a smaller shop, maybe you are able to simply grab a clone volume at the end of the week, or at the end of each day, simply to guard against disaster or accidental deletion and your set. For most of us, though, this isn’t robust enough.

The first media rotation we’ll cover is Grandfather-Father-Son. There are many variations that you can use in all tape rotation schemes. The variant of Grandfather-Father-Son that I prefer involves a daily set of media for each business day, a set of media for the end of each business week and a set for the end of each month. You can choose whether you use incremental or differential backups, but the idea is that you run one of those two types of backups every business night, besides for the last night of your business week. These daily sets are stored preferably offsite to be used during the following week. At the end of your business week, use one of your end of week backup sets, to write a full backup of your network, which will be stored safely for a full month. At the end of the last week of every month, use one of your end of month sets to write a full backup to be stored offsite for a full year. Of course, if you’re using disk to disk to tape style backups, you can write to disk first and then to tape, to keep a ready set to backup from on site at all times.

The advantages for this rotation is that it’s very easy to follow. You won’t have to question which media set to use and when to use it. It has a nice balance of data retention and media that you have to dedicate to the strategy. You’ll still be purchasing a good deal of tape media, but you’ll always have a full year of full monthly backups to rely on for the long term.

For those of you who like ancient Chinese games, we have a tape rotation strategy for you! The Tower of Hanoi, which is based off of an ancient Chinese game, is just a slight bit more intricate than the Grandfather-Father-Son scheme. In the Tower of Hanoi, you start out with four media

sets, which we'll call A, B, C and D. Your first backup is an A backup. You repeat using A's tapes every second day, so if you used them on the first of the month, you'd re-write that media on the third and so forth. On the first day that you don't use media set A, which would be the second of the month, from our previous example, you'll use media set B. This media set, B, is reused every fourth day. In our example, after you write media set B on the second of the month, you'll re-write to that set again on the sixth of the month and every fourth day afterwards. Still with us? Good! Now you're ready to start using media set C. You start using media set C on the first day that you don't use media set A or B, which would be the fourth of the month, from our example. Media set C is used every eighth backup day. Finally, there's media set D, which begins on the first day when sets A, B and C aren't used, which, again in our hypothetical first month, would be on the 16th and will be reused every 16th session. You also use an extra media set, E, to alternate with set D. Of course, this is just giving you an idea of how to rotate your tapes. You can adjust this to meet your own data retention needs.

Now that you know a bit about the data rotations and media sets, we should evaluate one common issue that many folks forget about - client machines. How do you assure that all of the intellectual property that is being produced by your end users is properly secured? Some folks elect to put this task in the hands of their end users, providing server space in the form of a share that they are asked to place their important documents on the server, to assure proper backup. People forget, though and people tend to eliminate tasks they don't enjoy and, truthfully, who really enjoys backup?

To provide end users a more uniform experience, many system administrators have decided to use Network Home Directories. This technology has been around for a while and allows you to keep all of your end user data on the centralized servers. The advantage to this solution is that the end users are assured of having all of their data at any workstation they sit down at. Since the data is kept 100% remotely, this is very popular in educational and lab environments, where a single person may not use the same machine more than once. This keeps extraneous information from the machine's hard drive and makes backup simple, as all of the user data is centralized.

Added in Apple's Tiger release and further refined in Leopard, is the ability to synchronize between a local home directory and a Network Home Directory. Apple calls this a Portable Home Directory and provides additional levels of protection and data availability that the Network Homes cannot provide. First, Portable Homes keep local and remote copies of your data. Note, this is not backup, this is improved availability. Since the data is synchronized based on managed client policy, the network and local homes are fairly identical. Consider it a mirror. If you delete something locally, it'll be gone remotely at the next synchronization interval. Where this helps your data availability is that you can "lose" the client or the server and still have your data available. In reality, a Portable Home is just a fancy Netinfo account with special provisions included for items such as the Authentication Authority. A good example of use of this solution is on laptops and portables. You take your portable with the Portable Home with you on a business trip. On the way back home, it's stolen. Granted, the work you did during the trip was probably not synchronized, but upon arrival in the office, all you would need to do is grab a new laptop

with the appropriate software installed and login using your account that's Portable Home enabled. All of the data that you last synchronized will be back in place. The same thought could be used on a desktop, if you were to have a surge or an act of a caffeinated beverage destroy the local machine. In some of these examples, you could make an argument that a PHD is a backup, but remember, it doesn't replace traditional methods. With PHDs, you've consolidated your backup scope to the servers, eliminating the need to install agents on workstations/portables and you've provided the end users with an added layer of redundancy. Combine that with a strong backup solution and you've got a winning plan to properly secure the data that's all around your organization!

Now, onto the software!

IBM's Tivoli Storage Manager

A complete storage solution for the enterprise.

If you have a Windows server, or you just don't care if your backup is running natively in OS X, IBM's Tivoli Storage Manager offers many features and has a full featured integration that many Apple fans will enjoy.

As you might surmise from the name of this software package, IBM's TSM software isn't just about getting your data from point A to point B, it's a complete data management package. TSM is able to completely manage all of your data, migrate data to tape, catalog what data resides where and even expire older data, based upon rules that you set up through the TSM Management Console.

The initial installation of TSM was fairly painless. IBM has done a good job of including fairly intuitive installers that will guide you through this process, step by step. You have the option of doing a complete installation, or customizing it, as you see with most installers these days. For our demonstration, we are going to be doing a complete installation. The installation of the code for TSM took about 35 seconds on an 800MHz Pentium 3 server with 512MB of RAM. The second part of this process is to then install the server license for TSM. All the clients for TSM are freely available at all times for all platforms that are supported, the cost is incurred when you actually want the server software to do something with those clients. Like the code installation, installing the licenses was quite speedy, literally taking 5 seconds in our tests. Overall, the entire installation process, including answering any questions to tailor our install took about 10 minutes.

Upon completion, we have the code installed, but it has not been initialized yet. To accomplish this, we will use TSM's Management Console. This piece of the TSM software will actually begin our initialization and is much more involved than the first part of our setup. Again, you have options as to how you would like to set up your TSM server, this time called "standard" and "minimal". Note that minimal will take all default values for any parameters that TSM wants to configure. For our purposes, we are going to go ahead and do a standard set up. You will need to complete a handful of wizards during this process, the first of which being the Environment Wizard. The Environment Wizard is fairly straight forward and simple. You simply inform your TSM server of what type of environment it will be running in. By this, TSM wants to know if it is a network server, which is capable of backing up numerous nodes on a network, or a standalone server, which will only back itself up. Once you've made this decision, you're whisked off to the Performance Wizard. The Performance Wizard is one of the more important tools you will use in your TSM setup. This will be strikingly evident in our eventual outcome. The task this wizard completes is a search of local drives to determine where TSM objects belong. It will then optimize its performance based on the size of the files that you want to backup. You will be faced with choosing between three different optimizations. The first is to optimize TSM for

backup and management of small files, the second being to optimize for large files and the third for a mixture of large and small files. Note, TSM considers any file over 1MB in size to be a large file and anything under to be a small file. You will want to weigh out the overall amounts of data you have to decide which optimization levels you want. While it seems like simply saying to optimize for both large and small files would make the most sense, if a large majority of your data is residing in files over 1MB, you will probably see increased performance by optimizing for those files and taking a performance hit for those pesky text documents on your server. This choice is yours, of course. The Performance Wizard also optimizes based on how many TSM servers you will be running. Optimizations in this area are based on three server node ranges; 1 server, 2-49 servers and 50+ servers. It took TSM about one minute to apply all of these settings to our server, a one drive, one server environment.

Next up comes the Server Initialization Wizard. Yes, I know I said we were already doing the initialization process, but that was just for the code, this is now for the actual server. This wizard will create several large and small files that are vital to the operation of the software. By default, these files will be created in a sub directory of where the base code has been installed. You can decide to move them to another area on your system, which is a good idea, as these vital files can run very large in size. So, what are these magical files you ask? In the large variety of files, you will find a database, a recovery log and storage pool files. These are files that can be redirected and it is recommended by IBM that you place these files on different drives due to how they operate. In order to truly have an ideal TSM set up, you'll want to consider the roles that each part of the system plays in the whole. When you evaluate placement of these three files, you should understand how they access their data. The database is highly random access oriented, while the log and storage pool volumes are sequential access oriented. For that reason, it is recommended that you separate the database from the other two files, to achieve better performance. Now that you understand what the files do, you should evaluate the storage needs of the files. Luckily, TSM has the flexibility to leave the actual size of these critical files up to you. The database can scale all the way up to 520GB in size and the logs can be sized up to 13GB. The storage pool will be sized based on how much data is being backed up to that pool in a day, it does not consider how much data is being staged off to tape during that same period. Of course, you will want to give the storage pools ample space to work with, as you will see added performance by allotting proper space. If you drastically undercut your storage pools in size, you will take a drastic performance hit as the pools will be staged off and filled at the same time, once full. If you create the pools dramatically too large for your backup needs, well, then you're simply wasting disk. As with all other aspects of backup and disaster recovery, plan accordingly.

Now you're up to setting up your security settings. Of course, you don't want to just simply let anyone run wild with your TSM server and luckily there are some very nice granular controls that you can apply to your user accounts. You can run the server under System, or you can also configure it to use specific created user accounts in TSM. With each account, you can specify specific tasks that a user is allowed to perform. You can, and should, set a password for your server. This is the password that will be used for host to host communications, such as multiple

TSM servers or LAN free clients. At the end of this wizard, it will build all of the specific customized settings we have been describing thus far. This took approximately five minutes to apply on the hardware that we were using for this demo. Going back to the specs of our server, you may want to note that TSM likes to have about 2GB of RAM available to it. On top of that 2GB of RAM, the Admin Center will want an additional 1GB of RAM. I was informed that ideally, you want to have 3.5GB of RAM, or more, if one server is housing all of these components for you. Obviously, this is not for the weak of heart, or server.

Now that we've configured the actual server, its time to configure our devices. The Device Configuration Wizard is designed to work with TSM device drivers. Now, please note, if you are using an IBM drive, you do not have to use the TSM bundled drivers and, in fact, we did not install the TSM device driver for our trials, since we are using an IBM drive. This wizard is still of use, though. It will find all the names of your devices, which comes in very handy, but overall, we're skipping this configuration section due to our choice in drives.

Time for the Client Node Configuration Wizard. This will help you set up all of the devices that you want to back up to your TSM server. You can configure your client nodes from this wizard or from the admin wizard. You can even do a little of both, which is what we are going to do, so we can get a flavor for both. First, you will set up your policy information, which is your storage management. It is a good idea to perform this set up in this wizard as it is a little more verbose than the similar set up features in the admin console. You also create your storage hierarchy in this step, setting up your back up pool, where your backups are stored, and your disk pool, which is simply a generic pool of disk. After your hierarchy has been established, you can create your nodes. Your nodes are anything that you want to back up. For each node you wish to back up, you can assign a name and a password. TSM does support Active Directory integration for node registration. Note, this password is not related to the server password. These are client to server connections and requires a unique connection and password, as opposed to the server password we created earlier on, which was used specifically for server to server connections. You can decide also to brave it without these encrypted passwords, but it is not recommended.

During this process, you will find some of the features that TSM excels at; primarily being a very granular, highly customizable data solution. You will be establishing four basic data parameters. You will govern your data by the amount of versions you wish to allow and by days of retention. All of your data residing in TSM will be governed by both. If the data on your server is continually changing, you can choose to keep anywhere from two versions all the way up to an unlimited amount, or even disable versions altogether. You can even specify how many versions to keep on the server after a user has deleted the live file, in case they change their mind, as users do at times. Along those lines, you also can specify how many days to keep your data live in TSM. Note, TSM never throws away the last backup file of a deleted file, for recovery purposes, it only deletes previous versions in the TSM storage manager for restoration. Your last version will be left on tape for you. That doesn't mean that your data will reside there forever. You have yet another customizable option in TSM where you can specify when the last copy of deleted

files in TSM expire. If you like, you can set this to be ‘no limit’, but if you choose to specify a given amount of time, note that after that expiration threshold has passed, that file is gone for good. This places a large burden on the user to inform the administrator of deleted information. While this policy is able to reclaim valuable tape space for you, it is worrisome to consider the consequences of a user who didn’t realize that a specific file was deleted and the time threshold had been passed. With the good, you always have a tradeoff. This advanced policy is no different.

We’ll now begin to explore the Intelligent Server Console (ISC). The ISC can be run from the GUI, but also has a complete set of command line tools. The server comes with some built in accounts with administrative rights and you also create accounts for each additional client that you join to your server. Upon starting the ISC you will want to set up all of your admin accounts and get that out of the way. Here is where you can also create additional users to administer your TSM server. This area, again, shows the painstaking granularity to the controls that you can set in TSM. In a style that will seem very familiar to any LDAP administrator out there, your additional users can be configured to have password policies that you specify, as well as complete access control to allow your users to perform more, or less, on the server. There are three types of users that you can configure in TSM; Storage, System and Policy. The Storage user can only administer the storage pools on the server. The System user is the all mighty user, with no limitations, like root. Finally, the Policy user is only allowed to administer policies on the server. Another nice feature of the system is the fact that you set one group of security settings for all Tivoli functions. You also must set up each user that you have defined to have its own set of server connections. This specifies which servers they are permitted to administrate.

The next step in setting up TSM is creating storage connections. You will begin by adding a storage device, which launches yet another wizard. Simply pick the device you want to add to your storage list and click “Add a Device”. The wizard will guide you through defining the type of storage device you have selected (LTO, etc), setting up names for your devices and defining the type of library you have connected (changer, manual, etc). Note, TSM does not list fibre channel as a connection option. You will have to designate your device as a SCSI device in this case. Also, if you have more than one TSM server, you can simply direct your new TSM server to use a storage device that’s been defined previously on another TSM server.

After you have completed adding your devices, you will continue in the wizard to add a drive to your TSM server. You use a string to name your drive. You can also add a drive element address, if you like, but TSM is able to do this for you, so it is preferred to allow the system to handle this option. Continuing in the wizard, you will now be prompted to add volumes to your TSM server. Volumes can be discovered by barcodes, list, range or almost any other imaginable way. TSM is able to label the tapes for you, once they have been discovered. Upon completion of naming your tapes, TSM will now want to create Tape Storage Pools. A Tape Storage Pool is simply another stop for your data in the storage hierarchy of TSM. When your data migrates out of the disk pool, they flow right onto tape in the Tape Storage Pool.

Speaking of data migration, you can modify your pools to specify the order that TSM uses their resources. For example, you can modify your disk pool to point directly to the tape pool as its next device. What this accomplishes is, for example, when your backup pool becomes full, your data will automatically migrate to your disk pool, when that disk pool becomes full, the data will move from the disk pool to tape. So, what determines how much data should migrate to the next pool? Much like everything else in TSM, you do. You can set migration levels for all of your pools and specify when they begin to migrate data. For example, you can specify that a pool migrates to the next when 90% of the capacity of the pool has been used. You then also specify when the pool stops migrating data, for example, 70%. As a rule of thumb, you'll want to remember to give your migration percentage enough leeway to be sure you have adequate storage in the next pools available. If you consider that your pools may be quite full, eventually your data migrations will involve your final piece, migration to tape. You'll want to be sure that you have given TSM enough tape to complete its migrations, or TSM may not be able to complete these rule based migrations for you. Also, on the topic of migration, there is data freezing. You, as the TSM administrator, have the ability to freeze data on disk so that it does not migrate to the next pool. You specify the retention of this data in an amount of days. If your pool fills up during your specified time period, TSM will then look at the advanced property settings, that allow migration to begin if you have not hit your 'freeze' threshold. Again, another beautiful example of the in depth, granular control that you have over TSM.

Finally, you're ready to perform a back up, right? Almost! Finally, you have to register your clients on the server for back up. Compared to all of the rest of the set up, this is a breeze. You first must decide if you want to use closed registration or open registration. Closed registration means that you must create a client account in order for the server to back it up. Open registration means that, if you simply connect to the TSM server, TSM will generate a username, password and account for your client to back up to the TSM server. You probably want to use closed registration to guarantee that you are backing up only data that you want to be backing up. Also, TSM only cares about the clock in the TSM server. This means that TSM is responsible for kicking off each scheduled client/server back up, it's not a matter of what the clock says on the client.

Beyond your daily backups, there are many other daily tasks and maintenance that you can, and will, undertake in TSM. Daily maintenance includes telling your different pools where they should off load their data to. For example, you may want to have your disk pool and tape pool and copy them down to your copy pool to move off to tape and be taken off site.

TSM excels in its ability to save space on your different media types for you. You can set up variables to allow the TSM server to go through your pools and tapes to free up different space for you, saving you disk and money in the long term. All of these services are scriptable and have complete Daemons that can run to keep different scheduled tasks for you. All of these scripts are kept as text files, which you can view and edit in the Integrated Solutions Console. You can set specific volumes, or network drives that are attached to a client computer to specifi-

cally be backed up as well. There are group level policies that you can create to back up specific areas on a group of computers. You can also enforce policies that exclude specific drives, directories or files on any client machine, or group of machines. For example, you may allow your users to keep a library of their favorite tracks local on their machine in iTunes. You can exclude specific directories from your backup, or you can also simply filter out all of the music files, by file type, allowing you to not have to guess where your users may be keeping their tracks, but still being sure that you're not wasting company tape on backing them up. All of this is accessible through the new Maintenance Wizard included in TSM.

Disaster Recovery management is still not included in the GUI, so it must be done completely by command line. For those of you who do not feel comfortable with the command line, this is probably where you may want to jump off the TSM bandwagon. On the other hand, you may decide that TSM has so many great, integrated features that you will begin to learn some command line. You can create administrative scripts to check out tapes that you wish to send off-site and another script to keep track of tapes coming back from offsite. Otherwise, you will have to use the TSM command line interface to manage this section. This feature is due to be released sometime in the near future.

TSM has one main activity log, which is held in the database. You can dump it out into a file, if you would like. There are searching and filtering capabilities for the log entries through the GUI, so it is not completely necessary to dump all the data out into a file to get meaningful information from it.

Tivoli uses threads to search for and back up data. You must run at least two threads at once. One thread is used to search for the data you want to back up and the second thread is used for the backup itself. Tivoli can use up to 10 threads for backup, at one time.

All reports in TSM can be emailed or even posted to a website, so that others can see the status of your jobs and help you monitor them. The standard report has numerous sections that you can turn on or turn off. There are also additional monitoring features that you can add-in to keep track of specific information that you may require. You can arrange to keep multiple versions of your back up reports on the web and you are able to configure how many reports to keep before deleting them off the site. You can also configure web pop-ups that will alert you to different status changes and events on your TSM server. A TSM administrator can also set the threshold of error messages you receive, meaning you can decide what a 'failure' or 'success' really means. Along with numerous log files, there are many graphical reports you can view as well, allowing you to visually see what's going on with your backups. For example, you can see mounts per hour, bytes per hour and so forth, to help you evaluate your schedule for your backups. In the Activity Log Details in the GUI, you can suppress different messages, so any specific types of messages, you can inform the system to no longer show you. The server is able to send hourly update reports on how your jobs are running, informing you of any successes or failures in your back up job.

Now, finally, on to what we found when running the TSM system on some live data.

In our test runs, we used a 1.25GHz PowerBook G4 with a 1TB LaCie disk attached over Fire-Wire 800. Obviously, this is not the most ideal set up for a back up, but being that these tests were taking place off site, they should have sufficed. I found the back up and restore console for TSM in Mac OS X to be very easy to navigate and use. If you've seen almost any back up and restore application, you pretty much have an idea of what this is comprised of.

What I was not expecting, were the test results. We allowed TSM to run on a 100Mb network for approximately 90 minutes, pulling data through the PowerBook from the LaCie drive. When we came back to check on the progress, we found that TSM was able to only pull approximately 3.5GB of information. Obviously, this was far from what we expected and we cannot expect anyone to use this demonstration as a true measuring stick to how TSM can perform. According to our hosts for the day, IBM quotes that TSM can pull about 100GB an hour in a tuned backup scenario. This also goes as a good example as to the ramifications to improperly allocating space for your storage pools. Most of the best practices that were preached about earlier in this document were only partially observed and now you can see the result. On the other hand, it's very hard for me to be able to confirm that TSM is able to truly pull 100GB an hour, simply if we tuned the network a bit better. Obviously, upgrading the back-end of the network from 100Mb to 1000Mb, or even to 2000Mb fibre would multiply our data backup by 10x-20x, in theory. That's moving you up to 35-70GB an hour, with the high end assuming you have a SAN in place. I did find that TSM is able to pickup a job incrementally from where it has been cancelled. The system was able to identify the 10,117 files that were to be skipped in about 15 seconds on my system and then began the job where it had left off without issue.

Speaking of incremental data, incremental backups will stack on one tape, until that tape is full. When the tape is full, it will move on to the next tape. TSM takes the list of things you want to restore, figure out which tapes they're on, load up the tapes, largest first, and begin restoring. As data expires, it will take the remains from a tape and combine incremental backups and store them on a new tape, freeing up multiple tapes. This is the TSM's ability to reclaim storage space back for you. The reclamation process is automatic to TSM, so an administrator does not have to perform this manually, or kick off such a job.

More importantly than the backup of the data is the restoration. We all know of various ways to copy data from one way to another, the question is, if we had to restore from that copy, could we without great amounts of effort? TSM was able to restore numerous files very quickly from both disk and tape, which was indicative of the robust database in TSM, keeping proper track of where your data is available. I found this feature to be quite impressive. More importantly, TSM was able to restore the Mac specific data, including resource fork and even the orange label that I placed on one of the documents.

That's the good, pertaining to data restoration, now here's the bad. TSM includes two different restoration consoles. One is accessible by regular end users and one is specifically for administrators. This is a great feature for anyone who wants to allow their end users the rights to restore their own files. My issue with this system is the permission sets that were applied to the documents upon restoration. When I restored a document as a normal user, the document was placed on my Desktop and I was given ownership to the document. No real problems there. As an administrator, I again specified to place it in the active user's Desktop directory. TSM's use of the term "administrator" was taken quite literally by Mac OS X, as my documents were restored to the root user Desktop directory (/var/root/Desktop). I attempted this restore again with the same result.

Overall, provided that the performance and permissions issues we experienced are worked out, Tivoli Storage Manager is a very high grade, versatile, comprehensive solution for the enterprise. It should also be noted that many features that are available to the clients of other platforms are still not implemented for Mac OS X. In that respect, we can expect TSM to only grow in popularity, in the Mac OS X community. When you consider that many high end design, graphic and video shops use Xserve hardware, the ability to create rules to migrate these huge files off your live disk and truly use nearline and offline backup solutions, you can see how TSM could save your company money over time.

Pros: Complete solution, great controls and reporting tools, ability for end users to restore their own data, quick and easy restores.

Cons: Permissions applied to restored data are suspect, cannot run server on OS X, weren't able to see faster backup speeds, may be too involved for most IT folks.

Tolis Group's BRU

Backups and restores. Plain and simple.

In stark contrast to IBM's enterprise offering, Tolis group offers a package that specializes in just data backup and restoration, which you can gather by its name, Backup and Restore Utility (BRU).

For folks who have come over from various Linux distributions, BRU is a name that you are probably very familiar with. Due to the reduced amount of features to BRU, there is far less to set up and maintain, which is something that many smaller shops will find very attractive.

The BRU installation comes in four different packages that accomplish four different tasks. The first component is the actual BRU server. This package is to be installed on the server that hosts your backup device. Set up was refined in the 1.1.1 version of BRU server as the software has an improved GUI and friendlier interface. With that said, I originally used the term "sushi-grade" for the BRU GUI; its edible, but raw. Installation of the server component is no more than a few clicks through a normal looking OS X Installer package. The process takes no more than a few seconds on a dual 2.0GHz G5 Xserve and is seemingly as fast on a 1.25GHz Power-Book. With the newer 1.2.x builds of BRU, you can continue to see BRU becoming more evolved, user friendly and generally easier to pickup and run with.

When you start up the Server Config Tool for the first time, BRU will ask you to authenticate as an administrative user and will then immediately begin to search for attached storage devices to your computer. Upon completion, BRU will ask you to provide a password for your BRU Server. This password will be used by your administrators to connect to your BRU Server to create backup jobs. You'll then need to install your server license, as the BRU Server Agent Config Tool will automatically be opened for you. Successful completion licenses the product for full use, which to BRU means that your ability to backup data won't expire and the restarts all of the daemons associated with the product. Licensing for BRU allows you to install multiple license keys in order to add extra clients and so forth, so you never really outgrow any single license, you just add additional keys to your installation. That's really all there is to the Server Config Tool. The whole process takes just a couple of moments of time and not a whole lot of thought, at all. One omission from Tolis' set up guide is that you must have a tape loaded into any tape device you are setting up during these steps.

Moving on, you now will continue to open up the BRU Server Console, which is the main GUI you will use to run your BRU server. On the first run, and Initial Config Wizard will launch and will guide you through some additional setup steps. You can elect to skip the wizard and configure your BRU environment manually, if you are familiar enough with the product. We'll begin by using the wizard.

Clicking the Next button will present you with a list of tape and disk devices available to your system. You will then be asked to select a location for BRU to stage your backups. This is a location that the server will write its data to. While the default is to place this location on your root volume, in the Applications folder, best practice would dictate using a secondary disk, such as an Xserve RAID unit or other such high capacity, high availability disk. You also will have to select a timeframe as to how long these backups will exist on your disk before they are removed. This is where pre-planning the ILM starts to come into play. Not only will you have to factor in any storage constraints you have, but you also may be guided based on corporate policy or any other such mandates or regulatory compliance. For our testing purposes, I'm going to use a secondary partition in my local system and allow these backups to age seven days before being recycled.

Next you'll be asked to configure any known clients to the system. Since this is a new installation and I haven't pushed the client agents out yet, this field is blank for me. For fun, I installed the agent on my PowerBook as I was typing this and fired up the Server Agent Config Tool. Note, the BRU server won't just simply find and connect clients that have the proper agents installed, you need to configure the agents themselves. Setting this up is fairly straight forward and easy. You can simply click the Start button in the GUI tool to startup the daemon, then specify the server that you wish to have authorization to backup this specific client. You'll need to provide the IP/DNS name of the BRU server and the password that we setup for that server earlier on. Now, stepping back in the wizard displays my PowerBook as a valid client node. Selecting a client system allows you to select the client specific settings of compression and/or encryption. It should be noted that this is done on a client per client basis and only affects the way that the data is transmitted between the client and the server. This does not specify in what manner the data itself is backed up on the server. Since this is a portable and is mostly using wireless networking, I'm going to use both compression and encryption for my backup jobs. Finally, you are asked to include an email address for the server to report to. Note, this setting is required. If you do not configure an email address for use, this wizard will be run each time you attempt to use the Server Console. Before completion, the BRU wizard will inform you which applicable ports you need to have available for proper use, which is very helpful to know and eliminates a great deal of guessing and administrative frustrations.

Now you can begin defining your backup jobs. This process is much improved over previous versions. The GUI appears to be much more responsive and easier to navigate. Select New to begin creating a new job and give the job a name. By default, this job will be a full job. After naming the job, make your selections in the data tree of the systems you wish to backup of the directories you wish to include in this job. Upon saving your job, you will be asked if you wish to schedule your new backup job. Clicking Yes allows you to select between Never, Once, Hourly, Daily, Weekly and Monthly. Any selection you make, barring Never, gives you advanced capabilities to fine tune exactly when the backups will take place. If you wish to base incremental or differential jobs on this full backup, simply select the job in the list and click on Modify. You then can choose what type of backup you wish to base off of this original job, including another full, if you wish. This allows you to create a single base job and then refine it

multiple times to meet specific needs, while eliminating the need to make the same backup selections over again.

Kicking off a job provides a much more refined GUI screen, displaying real time statistics about your backup job. While previous versions allowed you to monitor your jobs, I found that monitoring a job would be erratic and unpredictable, many times becoming non-responsive.

Of special note is the minimal amount of system resources the job took to run. On the client machine, the backup job took about as much of a toll on the system as a single Dashboard widget. The server component seemed to take about 15% of my Quad G5's CPU capabilities.

You'll find that the GUI to the new version of BRU has been refined over the previous offerings for Mac OS X. You'll find very easy to understand drop down menus that will help you set up a backup job and save your setup for later use. The GUI is a single window that has different tabs designating the type of information you may want from your BRU server, such as Backup, Restore, Verify, Data Manager and Tools. Under each category you'll find pretty much what you'd expect, with more features being neatly tucked away by drop down menus. Since Backup, Restore and Verify are all fairly standard, we'll take a moment to overview Data Manager and Tools.

Data Manager holds all the information about your data, as the title says. In here, you will be able to set up different variables about your data retention. The first option you have is to configure specific information about your backup administrators. You can specify a storage limit and the maximum amount of days your users can keep their data in storage for. Of note is the inability of this release to create additional users. According to Tolis Group, this feature will be implemented in version 2 of the BRU product. Continuing on, as you can specify specific parameters for your users in this area, you can continue on to the Client drop down and specify if you would like your data to be moved with encryption, or compressed, for each client node in this area. You also can find out when the last time a backup was performed on specific client nodes on your network. You can receive similar information about the tapes that you are using, such as the longevity of a specific tape and how much data has been written to a said tape as well. Archives work very much in the same mold as tapes do, except with archive specific information instead of tape, like a Machine Name instead of a Barcode. History behaves as it should, displaying a list of previously attempted jobs on your BRU server. You can also check in on your destinations, meaning where your data is being written to, and your schedules to see what is going to take place on your server and when.

The Tools tab contains mostly system information, but you'll find it especially useful when troubleshooting. The Job Monitor does as it says and will report back information about currently running jobs. The Device Configuration drop down will show you specific information about your backup device, such as an LTO2 drive, or such. It is important to verify that your device information listed in the BRU software matches what it is in real life. If it does not, you will

have to reconfigure your device through the provided BRU assistant. You accomplish this under the System Hardware Scan, the next drop down in the Tools tab. I found this hardware scanner to be slightly problematic at times and did not always properly update the BRU server with the proper information about my specific tape drive, despite discovering it. The most important part of this area is the Server Commands area. Besides being able to stop your BRU server, you can also run your Housekeeping here, though BRU will do this for you automatically and you can import tapes to the BRU system. More important than all of these functions is the Debug Dump command. This is the first place you want to turn when you are having troubles with your BRU server and can provide many valuable tidbits about how your system is really performing.

In previous builds, such as the 1.1.x series, I was originally unable to get BRU to properly work with my tape library, I since have set up the software properly and haven't looked back since. The software itself has been rock solid in my experience with it. BRU has its quirks, but it also has a very nice scripting ability, allowing you to really customize your experience. For example, you might find that BRU's GUI doesn't handle importing and exporting tapes from your library's mailbox slots to the drive that BRU controls. If you look through the command line administration features of BRU, you'll find that BRU's libctl command can be used for this purpose, so over an SSH connection, you'd be able to import and export your media sets. This is my preferred method for rotating tapes in and out of my BRU controlled drive. Performance hasn't been an issue either. BRU doesn't seem to have much trouble writing on average 80GB/hr to tape, depending on which servers I'm backing up. I've seen a range from 65GB/hr all the way to 90GB/hr.

In my newer testing, in the 1.2.x series, the BRU server was able to backup 2.57GB of data, spread over 4656 files, from my portable machine. This took about 35 minutes to complete the backup of my PowerBook, which was connected over wireless 54G networking. Performing a restore of a single text file seemed to never end in the GUI. Looking at the client machine showed that the data had been successfully restored, with proper permissions and without corruption, so we'll chalk that up to a GUI error. Looking at the disk backup, the file shows up as a generic UNIX executable that's 2.45GB in size, so slightly smaller than the actual data that's stored within it.

Overall, I've had to restore data and have never had an issue restoring data back to my servers, while using the product. Permissions were always preserved and the restores were always done quite quickly and painlessly. A worry that most folks have with software that replicates Mac OS data is preservation of Mac specific data. BRU is fully aware of such data, like resource forks, and handled this data properly.

Moving back to the overall user experience, I will note that this solution has a native Aqua interface. This allows native installers, which make BRU very attractive to those of us who are squeamish of the command line. With that said, the BRU GUI feels more like a hinderance on the great command line tools that BRU comes with, rather than an enhancement, at times. This

is not necessarily a bad thing and could come from my previous experiences with earlier builds of the product. As they say, old habits are hard to break and when you are used to using a specific method to administrating a system, you tend to stick that way, even if GUI features are fixed and enhanced, which many are in the newer BRU builds. I would think that BRU might make it very easy for many users to become more comfortable with command line tools, as you will see much of the same output in the GUI that you will in the CLI. Also, when you consider that BRU is able to schedule your backups for you, you won't have to work with BRU too often, aside for loading and unloading your media.

Pros: Native to Mac OS X, continually tested against newest OS X builds, very attractive price tag, very quick to pick up and simple to use, no fluff.

Cons: Can require some CLI prodding and familiarity to feel comfortable with the product, lacks some advanced features some enterprises may desire.

Atempo's Time Navigator

The Flux Capacitor for your backup needs.

Ever lost a volume? There's always that sinking feeling. No matter how much planning you put into this scenario, you're never just quite prepared for when it happens. A million thoughts rush through your head. Do you have a good backup on site? How good was that 'good' backup? What if it wasn't as good as you hoped good to be? Suddenly, the thought of your server, in all of its glory. The good ol' days, when your data was safe, your users were happy and you were happily sipping on your freshly brewed, morning cup of coffee, out of sight in your office. Don't you wish you could go back to those days; to that feeling of security? Enter Atempo's Time Navigator software and its unique look at your backup needs.

The first thing you will notice about the new Time Navigator software is that you no longer need to install the X11 application! Atempo is now an Aqua based application. With that said, if you have used previous, X11 windowed Time Navigator installations, the GUI looks pretty familiar, with some nice cleanup and polish that may have been lacking a bit in the past. Odds are, if you disliked the interface before, you'll probably still dislike the new interface a bit less, but remember, backup software is meant to serve a purpose and be reliable. No one cares what the GUI looks like if you can't recover lost data with it. Of course, I've seen products in other areas where the GUI is so horrible, that you simply can't work with it. This is not the case with the Time Navigator product, at all. Time Navigator's GUI is pretty nice looking, intuitive and even colorful! Remember, just because the GUI is now in Aqua doesn't mean there are wholesale changes to how you run the application. Time Navigator is a different breed of backup, synchronization and archival solution, some people may prefer this flavor over others or may dislike it for the same reason.

Back to your installation. So, we've already eliminated the need for X11, but Atempo's done more than that. This update is much more than a skin deep refresh. If you've previously read this document or run TiNa yourself, this passage may seem familiar:

"You've installed X11, but you have more system configuration to complete first, before you begin moving executables to your drive. Get ready to use your command line skills, you'll need them! The first step is to manually edit your hosts file, located in etc/hosts to reflect your host's IP address and your host's hostname. You'll now want to move on to edit your services file, located at etc/services. By default, Apple has included services to run on ports 2525 and 2526. Time Navigator needs these ports available to it to run its own services, including the tina_daemon. You'll want to delete the entries that are included by default. At this point, its time to move on to change the shared system memory settings for your system. This setting is kept in /etc/rc. Once again, use the editor of your choice and change the value of shmmax to 268435456. That wasn't too bad, right? Ok, now you can reboot your server and get ready to start installing your Time Navigator software." - Apple Enterprise Backup Solutions vol 1-3

While that may have seemed like a laborious task, remember, that was just being setup once, when you installed the server. With that said, Atempo has made this easier. Licensing, one of the most painful parts of my original experience with TiNa is now refined. When installing the server you can choose to use a License Manager Server, a License File or a simple 30 Day Evaluation. It is as simple as that and works as advertised. The License File was a big feature I wanted to see in this product and I'm thrilled to see that there are simpler ways to set this up. Of course, most folks weren't re-installing the license over and over again as I was!

Continuing on, the need to edit the `/etc/services` file is no longer done by hand. When installing the server component, you are simply presented with the TCP and UDP ports to use in the GUI. Again, just another refinement of this product that shows how Atempo truly wants to continue to improve upon what already was a strong solution. You'll need to reboot when you are done, as there are plenty of configurations going on, but that's to be expected.

As you continue, the installer will then ask what type of installation you would like to perform, such as Server, Storage Node or Agent - meaning one installer is used for any of your choices. Previously, Atempo's included installer is able to install locally on the machine that it's running on, which is no surprise, but it used to be able to install on one or more machines through an IP network from the GUI as well. The network installation functionality appears to be gone now. They are using a standard Apple installer now, so the process is a bit easier and more intuitive. Next, you will be prompted to decide what type of installation you would like to perform, except this time, instead of wanting to know about the methods used, Time Navigator wants to know if you are trying to set up this machine as a Server, Storage Node or Agent. For this case, we'll be using Server, as this is our primary backup server that is attached to our LTO2 library. You'll be prompted for some environmental parameters, some of which is the information that we added into your system in the previously mentioned steps and refinements. Next, you have the options of which components you install on your Time Navigator server. You can decide to omit items such as the web restore console, web supervision and advanced database indexing. You do have the option of coming back later and installing these optional components and tools. Finally, you are prompted to specify a path to where the Time Navigator web interface directory will be stored.

While the network installation option right in the installer may now be gone, Atempo has included a new application as part of the TiNa installation that allows you to create an ARD package. I found this to be a great addition to the solution, as most administrators would prefer to have a simple package that they can then push out to their clients as they wish, as opposed to having to use the specific Atempo tools each time.

Again, since we're looking at Time Navigator and how the product has progressed, let's take another trip down memory lane. If you previously read one of the series of these documents, you may remember the following passage:

“You’ll now continue on to install client software onto the systems you wish to backup to your Time Navigator server. This is a fairly painless process. You will have to edit your hosts file to include the IP address of the client you want to back up and you also have to include your hostname. One quirk I ran into was in the naming conventions here. I originally used a FQDN, but I was unable to establish a connection between my client and server. Atempo includes a great utility called tina_ping that can help you resolve such networking issues. In this case, on a hunch, I decided to change my hostname from a FQDN to a .local address, which then seemed to work perfectly. This may be different in your environment, though. Also, remember, that editing hosts can cause different network services to change behavior. For example, if your server is bound to an Active Directory to provide Kerberos support for SMB sharing, adding your server in hosts in a link-local naming scheme, instead of FQDN will cause your Kerberos to break. Atempo thought of this and included its own hosts file, kept separately from the system hosts file, where you can create your own aliases, or any other records you may need. You will also have to edit /etc/services, as you did on the server, to delete the services that Apple includes running on ports 2525 and 2526. Time Navigator will fill in the information it requires in this file for you, just as long as you free up the space for it to write to.” - Apple Enterprise Backup vol. 1-3

While you still may need to work with a hosts file for configuration reasons, one thing that I pointed out to most of my clients is that Time Navigator included their own files where you could easily configure these options without having to disrupt the files used by the OS. If you take a look in /Applications/tina/Conf, you’ll now find a handy hosts.sample file, which you can use to associate different localhost name with the machine the file is stored on or if you wish to reach a client that’s associated with a different domain name. Atempo also states that using this configuration is helpful in clustered configurations. If you poke around in this area a bit more, you can also find the Security Compliance Modules that Atempo includes, which covers SEC rule 17a-4, SOX, HIPAA, and Basel II.

Atempo includes the ability to install Time Navigator in Batch Mode. In Batch Mode, you use the editor of your choice, I’m a Vim person, and define specific environmental variables to help configure your agent. This is the way to go if you have multiple Mac OS X Servers that you want to set up as agents, as you can reuse this batch file over and over. It’s tailored to be specific to your Time Navigator catalog, backup server and type of installation you want to perform, not specifically to the agent you are trying to install the software on. Atempo states that you cannot perform this action over the network, but a little ingenuity can get around that. Very simply, you can push out disk image of the Time Navigator software and initiate the installation process. Since it will look at your batch file for all user data, it won’t need any live user interaction. This whole process could be scripted and then simply automate a process to delete the disk image when your done. When you have finished with the installation, which took no time at all, simply initiate a connection from your backup server and create your jobs, classes and schedules for your new agent. A good tip is to be sure that you have connectivity from your Time Navigator server to your new agent, which can be done two ways. First, you can run top on the server you wish to back up to see if tina_daemon is running, which would facilitate all of your backup proc-

esses on the give client. Second, is to use the Time Navigator specific `tina_ping` mechanism as follows: `tina_ping -host clientName`. If you perform this action, you not only will be able to see if your Time Navigator server is properly seeing your newly installed agent, but you can also get versioning information, which can be helpful when performing updates of your software.

Also of note is the new administrative console in Time Navigator. You'll notice that you now have a Time Navigator Launcher application, which has numerous additional applications and functionality available through it. Keep in mind that the Launcher application doesn't include every bit of TiNa functionality. You can still look into the Applications folder to find additional tools at your disposal. In the main window, you'll find the familiar Administrative Console and Restore & Archive Manager; the tools you'd expect to be your main living area in TiNa.

In addition to the main administrative tools and your monitoring tools, you'll also find the configuration tools, the Virtual Library System, Environmental Reporter and Certificate Installer. One nifty tool here is the Command Line Interface. This button will simply start up your Terminal.app, but will have you set and ready to use the TiNa environment. Just a bit of a quick time saver, if nothing else, but appreciated none the less. We'll get into nifty TiNa Terminal tricks in just a bit.

Creating a connection to your library, or tape device, goes much as you would expect. Some things to note during this process. You first define your library and despite Time Navigator's ability to detect the make and model of your library, you will still have to manually select a driver to use with the device. Not the biggest deal in the world, but something that stuck in my head. The second part of this process is to then define the drive technology that is being used in your library. This part may confuse some folks out there. You run into the same deal of having to pick a drive to be used. Now, recall that I am using an IBM 3583 with LTO2 drive technologies for this test. You would assume to look for your driver under the IBM heading again, as I did for the specific 3583 library in the previous step. This is not the case. This time, despite the GUI asking for the vendor, it's really looking for the type of technology first, which sends me to the LTO heading and then I can find the specific driver I need in this category. Again, not the biggest deal in the world, but something that some folks may get stuck on and there's no worse feeling than going through all these steps, thinking you'll have a backup done in the morning to only come to the office to find nothing wrote to tape.

Your basic Time Navigator configuration is in place. You have your clients installed, along with your server and your tape device, you now should look at some of the other features of the Time Navigator software and create some storage pools.

In reality, Time Navigator doesn't really do anything earth shattering in this arena. They use Cartridge Pools that are created to manage cartridges in a catalog and are governed by a management policy. This allows for automatic movement of data, much like other products offer and much like many backup administrators, or non-backup administrators crave and demand. The

difference here is the process is very painless in Time Navigator and seems extremely intuitive, allowing you to set up real policy and pools without much effort. Along these lines, Time Navigator is also able to set up separate pools for automatic backup and archiving.

Time Navigator also has a Virtual Library System. A Virtual Library System, or VLS, is a way of creating a disk to disk backup that resembles the characteristics and has the same functionalities of traditional tape libraries. Since it is treated just like traditional, physical tapes, you can apply the same types of management and data expiration. Think of it as a new way to perform disk to disk backups, but with more granular control and an easier way of managing it.

Creating backup jobs is fairly straight forward. Time Navigator includes a Wizard to walk you through creating a Class and Strategy, which are the two major components of your backup task. Of course, you can forego the Wizard and simply do this by hand, either way works. There doesn't appear to be much of a time savings to do this one way or another through the GUI. To further differentiate here, your backup Strategies are associated with Classes. These classes are used to determine which directories you want to back up for specific platforms that have the Time Navigator agent software installed. Classes also control the Time Phases that are allowed for incremental backups, the format in which your data is being backed up in and any filters you would like to set up, to eliminate specific file types or files with specific name attributes.

While Classes designate which directories are backed up, Strategies determine the rest of the critical information needed for a back up in Time Navigator . Strategies allow you to schedule dates and times for automatic backups, specify which cartridge pools to be used, if you would like to write to two destinations at once and other job specific preferences. In Time Navigator, strategies run independent of each other, so if a particular file is backed during a full backup during Strategy A, is then modified, backed up by Strategy B, the next incremental backup that runs with Strategy A will back up this file again.

One of my favorite features of Time Navigator is the ability to create quick, easy, complete backups of your catalog. You can specify that your Time Navigator server back up its own catalog, but sometimes you would rather do it yourself. From your terminal window, you can accomplish a full backup of your catalog by simply typing:

```
tina_odbsave -file save.cod
```

This creates a backup of your Time Navigator catalog, with the name save.cod. If you have more than one catalog, you only have to extend this slightly:

```
tina_odbsave -file save.cod -catalog tina1
```

When you add -catalog, you can specify any catalog in your Time Navigator system. What makes this so nice is the ability to capture these saves without stopping the catalog itself. You

can keep your catalog live, with the catalog connected and quickly grab a backup of it. The only convention here that needs to be followed is that you must follow the name of your catalog with the .cod suffix. You can specify an absolute path as to where to save the .cod file, for example, if you wanted a copy on a FireWire disk to take off site, or you can simply specify a name for the catalog backup and it will be placed in the directory where you ran the command.

Some additional strengths I found in Time Navigator also include its Synthetic full backup technique. A Synthetic full backup is a full backup that is built from a previous full and incremental backups. Time Navigator is able to locate the last backed up instance of each object in the catalog to build a new full backup out of your incremental data. Additionally, I also found that Atempo's tech support team to be exceptionally speedy and knowledgeable. They fully understand their product and use the built in tools to actually diagnose your problems. For example, I had sent in a question regarding the catalog and attached a copy of Time Navigator's Environmental Report. Their tech support started our conversation with a full understanding of my environment and a solution already in mind. Equally as impressive, they took the time to use my own environment's host names, which I know doesn't necessarily solve a problem, but it was a nice touch that I thought I should note.

Now, what you all have been waiting for; some real world results! Since I had more access to data and tape devices previously, the performance numbers are going to stay the same, as the product performed well and continues to perform well in each installation I've seen it in.

I ran Time Navigator for about ten days. During that time, I ran two different types of backups. My first set of backups were completely isolated to one test server. I replicated nearly 1TB of data from an Xserve RAID to a LaCie FireWire 800 drive, connected this to a PowerMac G5, which had a Fibre Channel card installed and was connected to a single drive in an IBM 3583 library. I ran this trial four times and found Time Navigator to perform very nicely in this scenario. I averaged 105GB/hour of data, moving from FireWire 800 drive to tape, over my four trials. After a few days of doing this, I decided to deploy the agent on my live data server and give it a whirl. This scenario involves backing up a dual 2GHz Xserve with 2GB of RAM, 80GB boot drive mirrored, connected to a 3.5TB Xserve RAID, configured with two RAID 5 sets. This server will connect to my Time Navigator server over gigabit ethernet. I found that Time Navigator took a large performance hit while backing up my live boot drive. Speeds were fairly slow during this process. When Time Navigator hit the Xserve RAID, performance went right back up and improved the overall score. Overall, Time Navigator backed up my 1.6TB of total data three times, without issue, averaging about 65GB/hour.

For early adopters, or folks new to the OS X Server party, you'll be happy to know that Atempo has added in support for ACL permissions and even a Widget for your Dashboard. While the Widget may seem like an extraneous feature, it is nice to know that you can keep this lightweight little app in your dashboard to get quick information on your backup jobs. It's not necessary, but these are the little refinements to software packages that help make your backup solution a little

more user friendly. This also reaffirms Atempo's continued support and development of their Time Navigator product for the Mac OS X platform.

For users who have invested in Final Cut Server, Apple's content management system, Atempo's Digital Archive solution may be the right solution for you. Atempo Digital Archive (ADA) is sold separately from Time Navigator and runs independently of the Time Navigator solution. What makes this solution so unique is the tight integration it offers to Final Cut Server installations. While Final Cut Server is able to track all of your digital content for you, Atempo Digital Archive makes restoring and migrating this data easy.

When ADA is added to your Final Cut Server installation, Final Cut Server users have the ability to archive material to pre-defined storage. Note that the archival migration of the clip from the live system to the ADA archive doesn't de-reference the clip data from the Final Cut Server itself, so users are still able to find the data, as if it was still on the live system. This is helpful in freeing up disk space on your primary, mission critical systems, while keeping the data accessible to the end users. When a Final Cut Server user attempts to view a clip that's been migrated off of the primary system, a special archive icon is displayed with the media asset. All the user has to do is click on the archive icon and select the "Restore" option. In the background, ADA will migrate the content from the archive and back onto the live Final Cut Server system. The archive icon will be removed from the asset and full use of the content will be restored. The ease of use of the ADA solution will make this a big player for environments that have brought in Final Cut Server, as it will take a lot of the administrative load off of IT and allow the creatives to seamlessly access their data, while keeping mission critical storage tamed. While the ADA Server itself does not run natively on Mac OS X today, this surely does give us a glimpse into the type of integration Atempo continues to provide on the platform. Atempo does intend on having the ADA software running on OS X in the near future.

A bonus on the Atempo website is a working demo release of their product for Mac OS X. This was a previously overlooked avenue for getting their product into the hands of perspective customers. You have to give Atempo some information about you to get the download, but that simply guarantees that you'll have a contact at the company during your trial period, rather than having to look around for one later. Now you don't have to take anyone's word besides your own, when judging if this product will work in your environment.

Overall, I found that Time Navigator performs very well in every instance. Having used this product with larger Xsan installations, I know first hand that TiNa scales very well. With a plethora of supported platforms, application plugins and replication features, Time Navigator can provide any sort of data retention, archival, backup and recovery options to fit your specific environmental needs. If you have a lot of data to backup, and possibly were turned off to Time Navigator with their earlier, pre-Aqua releases, you owe it to yourself to look at Time Navigator again. The system continues to live up to all the promise it showed in it's early days and makes implementing high end features fairly painless.

Pros: Command line tools, hot-save catalog feature, Xsan 1.x and 2.x support, fast restores, extensive feature set, great tech support, application specific plugins, innovative view on backups, replication capabilities, LAN free backup options, policy based administration, widget, automated archiving, encryption features, demo available for download. Final Cut Server integration via Atempo Digital Archive solution.

Cons: Lack of readily available documentation online (only available via installation media), can seem overwhelming to new users, some features not as intuitive as others, some modules are not Intel native yet, may be overkill for smaller shops (scales up better than down).

BakBone's NetVault

The Fort Knox of data protection.

Making the leap from the X11 window interface to Aqua is BakBone's own NetVault. Filling a vital need in the OS X community, NetVault offers many of the advanced enterprise class features that administrators desire coupled with a user interface they are accustomed to.

While some administrators will ignore a GUI upgrade as a feature of a specific software package, as the second major release for the product running on Mac OS X, the maturity of the product developments are notable.

Installing NetVault is fairly easy. You'll find that NetVault consists of two primary packages: one client installer and one server installer. The installation of the server component took practically no time on a 1.8GHz iMac G5, which means it should be incredibly fast on an Xserve. The installation does not require a restart.

Before you install your NetVault server (and you will need to install your server before any client software), you'll have some planning to do. Besides the normal considerations that you'd take with any backup solution, you should be aware of the storage needs of NetVault. Upon installation of the server component to NetVault, a netvault directory is created in the /usr directory. Inside of /usr/netvault, you'll find a db directory that includes four key components: MediaDatabase, ScheduleDatabase, install and keys. Understanding what these four directories are responsible for will help you understand the server component's storage needs. Keys is the simplest directory and simply stores any NetVault product licensing, making this a very small directory. The install directory is similarly sized, and contains a binary on which modules have been installed in your solution. The ScheduleDatabase, according to BakBone, is usually under 10MB in size. This directory holds the records for all of your backup and restore jobs. The largest directory that NetVault installs on your server is the MediaDatabase. The MediaDatabase manages all the records for your media and the backup indexes for all jobs that you've performed from your NetVault server. Liken this to the catalog that other companies use for keeping these records. BakBone provides a formula for calculating the space your MediaDatabase will grow to:

Number of files & directories	X	Number of generations	X	Number of	X	60
backed up per machine.		to keep using Backup Life		machines		

Now that you've decided on a directory for your installation, go ahead and double click that installer. It's a fairly common install process.

Now it's time to check out the server console. You'll find the NVAdmin application, which is the main console for NetVault, located in the Applications folder. Double click that icon and you'll be pleasantly surprised to see something very familiar staring back at you: Server Admin (of sorts). The folks at BakBone have done the Mac OS X administrators of the world a great favor. They have mimicked the GUI to Server Admin as their main console. You can add/remove servers just like Server Admin and you'll even find the status lights you've grown to rely on for instant status of your services included in your NetVault installation.

There are some "gotchas" to using the built in NVAdmin console included with this build. First, for NVAdmin to work from a client with your NetVault servers, the client software's version must be the same as your NetVault servers. Secondly, the NVAdmin console for OS X will only be able to Domain Manage, as BakBone refers to it, other OS X NetVault installations. On the flip side, all of the buttons in NVAdmin work as they do in Server Admin, with the results you would expect. You'll also find the same built-in Logs, Overview and other navigational buttons, installed right into the GUI, again making this almost seamless to integrate into your existing OS X Server infrastructure. The Settings button reveals a System Preferences like set of buttons to help you tune your NetVault installation.

Focusing in on some of the functionality of this area, clicking on the Jobs navigation button reveals a window that gives you the ability to view your currently running jobs, as well as the definitions that have been configured for the NetVault server. Along the bottom, you can choose to create a new backup or restore job. We'll go into more detail on job creation in a bit.

Next to the Jobs button, you'll find Reports, which also has more functionality than its name would indicate. Instead of just finding reports on your NetVault environment, you'll find a list of predefined reports that you can double-click on to run. Running a report causes NetVault to generate a dynamic web page displaying the results. You can also right-click (or ctrl-click) on any report to run and view the report, simply view the log, or delete the report altogether. The same functionality can be accessed by simply selecting the report name you wish to access and choosing the appropriate action out of the Item menu. This is all found under the Definitions tab in the Reports section. You can also click over to the Current Activity tab to find out what reports are currently being run, or have been run.

Before getting started with defining jobs and schedules, we'll need to add in all of your clients and storage devices. We'll assume that you installed the client software on the nodes of the network you wish to backup. BakBone's client software doesn't require any sort of configuration to announce itself to the network. To add in a client, simply go back into the NVAdmin console, under the name of your NetVault server, click on the Client Management heading. You'll then see a list of clients that have already been configured for use with this server. At the bottom of the screen, you'll see a button that says Add Client on it. Simply click that button and NetVault

will auto-discover other NetVault clients on your network. If your system does not appear in the list, you can specify an IP address or a network name to add. Note, the NetVault client software needs to be installed on the client for it to be discovered, either automatically or manually.

The Client area of NetVault allows a great deal of customization and configuration. Right-clicking on the client name reveals a contextual menu that gives you many different actions that you can take on your NetVault client node. These options include Properties, Remove, Check Address, Install License Key, Install Software, Remove Software, Set Description and Configure. Most of these are pretty straight forward, so we'll focus in on just a couple of these options. Properties offers an extremely comprehensive look at the client node on your network. This option provides your NetVault version and build information, the Machine ID (which you'll need for licensing purposes), the IP address of the node and the OS version of the node. Dig a little deeper into this sheet and you'll find the ability to list the installed NetVault plugins, including version numbers and installation dates, as well as the software that's been installed on your client for NetVault. Setting the description of the node, another feature of NetVault, allows you to add any sort of descriptive information that can help you identify a specified node on your network. This data will populate the description field in the Client Properties feature. You also can choose to configure your nodes from this console. Now, there are many software packages that allow some basic configuration. NetVault goes above and beyond and allows you to truly tune your environment. If you remember, I mentioned a System Preferences-like interface that's used to configure your NetVault server. The Configure option for your client nodes allows the same type of configuration capabilities. To begin to attempt to describe all of the configuration options you can control in this area of NetVault is best left to BakBone's own documentation, or a hands-on demo of the software. Yup, there's just that much you can configure! I will say that many features that a backup administrator might crave have been included. For example, you can specify an alert threshold to notify you when you have reached your storage capacity based on a percentage value you decide on. NetVault also has the ability to email you notifications, which is functionality that every administrator has become accustomed to. Functionality that perhaps some of us aren't used to includes the ability to specify how long NetVault should attempt to connect to the remote nodes and how much time to wait before dropping an inactive connection, as well. You also have the ability to specify in which directories NetVault will store all of its reporting, logging and statistical data. Again, this just scratches the surface of the advanced ability to tune your NetVault environment.

Ok, now that we have our clients connected and some basic configuration done, we need to move on to configure our storage and backup devices. Right under the Client Management, you'll find the Device Management heading. First, specify which NetVault node has your storage or media device attached to it. You'll then be asked to define what type of device you wish to add. You have the option of adding a single virtual disk based tape device, a virtual disk based library/media changer, a single physical tape device, or a physical tape library/media changer. When adding in a virtual device, you'll have to specify a target location to create the new device. From there, specify how you want the device configured, such as how large of a storage device you want, or how many slots to include in your virtual library and NetVault will go ahead and create

your virtual device, using a process named `nvdiskdevice`. Creating a single library with four slots, totaling at 1024MB on a 1.67GHz PowerBook G4 took about 10 minutes, so keep that in mind when creating larger libraries. Once your virtual device has been created, you must add it in to your BakBone system. Simply select your newly created device and click on the Next button. You can now go into the Devices navigational menu and see your virtual device listed under the Current Status tab. You are then able to right-click on any of your virtual slots and manipulate the environment, as if it were a physical library with actual tapes. Additionally, you can label your media, create a group label, include an offsite location, and even change the format of the virtual tape. Note, this all leads us to another feature of the BakBone product. While some solutions require the storage or media library to be attached to the backup server, BakBone is able to use any client node on your network to act as a storage node as well. NetVault, by default, will also leave the device under the control of the node that created it. You can also change this behavior to allow the device to be distributed (shared) amongst the NetVault server and any of the clients you specify in a single NetVault domain. Note, that an optional SmartClient license is required to attach devices to a client node.

(Unfortunately, I don't have access to a physical tape library right now, so reviewing setting one up will have to be reserved for a later revision.)

So far so good, right? Guess what? You've already configured your back-end and you haven't even broken a sweat! Well, now it's time for some more thinking. You now have to create your backup jobs. Click back on the name of your NetVault server that you are actively configuring in the list and click on the Jobs navigational button at the bottom. Now, it's time to get your server actually up and working, not just sitting and looking pretty!

You can start creating a new backup job by clicking the appropriately labeled New Backup Job button. This reveals all of your currently configured nodes on your network, under the Selection tab. Expanding a node reveals something that you might not expect. You'd expect to see a hierarchical view of the node, but instead you are presented with five options to dig further into: Consolidate Incremental Backups, Data Copy, File System, NetVault Databases and Raw Devices.

Expanding the File System subheading will reveal the hierarchical view that you are used to when configuring backup and restore jobs. You also have the ability to checkbox your way to a full backup job and submit it right away for completion. This works as expected and I was able to complete a backup to a remote node to disk quite quickly in my test. Before submitting your job, you can continue with the tabs across the top to set up your job, just the way you like. Clicking on the Backup Options tab allows you to decide what type of backup you wish to perform. Your choices are to choose a full or an incremental backup. You may immediately wonder, how can you perform a differential backup instead? NetVault doesn't appear to use the term "differential", instead, BakBone supports two types of incremental backups; data changed since the last full backup and data changed since the last backup, regardless of type. You are also able to select the option to backup data through an NFS mount and to check your files for modifica-

tions as you write your backup, two features that should be very popular with backup administrators.

The next tab over allows you to set the schedule for your configured backup job. Your options include immediate, once, repeating and triggered. Three of those four are pretty straight forward and function as their names would indicate, and allow for finite definitions of when to run the job. Triggered is an interesting concept. When you choose the triggered option, you are asked to supply a name for the trigger. What this does is it allows you to call a specific backup job, or operation, from an external script that you write. For example, you could use a script to dump out your Open Directory vitals and create a trigger called ODBackup that backs up those directories. In your script, you can call the nvtrigger command, which is located in /usr/netvault/util, and supply it the ODBackup name, which corresponds to the trigger you've assigned to schedule this job. When you make that call, your NetVault system will begin to execute the job you defined with that ODBackup trigger name. This is great functionality and could prove to be invaluable to many administrators, not just those who are into scripting.

Under the Target tab, you have the option to set up specifics on your output devices and media for this backup job. You can specify a particular device to backup to, such as the virtual library we set up earlier. You are also able to specify what media to use for the back up job. This is where the group label may come in handy, since you can specify a specific group of media to write to and even specify if and how the job can reuse your media. The general options allow you to specify how much space must be on the media before you begin to write to it and also allows you to lock the media from future writes after the job is complete. You can also specify that this particular job be placed as the first backup job on the media.

The Advanced Options tab allows you to include your own pre and post backup scripts, as well as specify any sort of duplication you wish to complete. You can elect to verify the integrity of your backup and use network compression. Another extremely helpful feature of NetVault is the ability to actually assign a life span for your backup. Not only can you set a length of time to keep this backup for, you can also specify when to compress and offline the index for it. These two options are vital in helping you control the amount of disk storage consumed by the media-database mentioned earlier. Compressing the backup index saves disk space but requires a little extra time to decompress when opening a backup index for a restore; if you set a time period for moving the index to an offline state, all space used by the index will be freed, but a media scan will be automatically triggered to rebuild the index prior to making restore selections. You can even base your backup lifetime on how many full backups for this job have been completed after this backup has finished.

So what's with the other subheadings? Do they do anything? Sure do!

If you read the section on Atempo's solution, you'll know that a strength that the product has is a Synthetic Full backup. Their Time Navigator product is able to take your last full backup and your differential backups, to create a new full backup, that's up to date. BakBone works much in

the same way. NetVault is able to take your most recent full backup and the most recent incremental backup(s) associated with that original full job, and consolidate them into a new full backup. This is done using BakBone's Consolidation Incremental Backups Plugin that was installed originally on your system. Note, the plugin doesn't actually backup any data, meaning there is no impact on your production servers. Its main purpose is to merge your full and incrementals together to generate the new, consolidated full backup image. Also note, you must use a full and an incremental to accomplish this. Two fulls don't make a consolidated full and it defeats the purpose of the consolidated full, anyway.

Actually performing a Consolidated backup is pretty simple to do, as well. You'll want to go back into the Jobs area and click on New Backup Job. When you see your client nodes, expand the client in question and expand Consolidate Incremental backups. You'll have two options, Backup Jobs and Backup Sets. While Backup Jobs is pretty straight forward, you might be wondering what Backup Sets are. NetVault understands that many different jobs share common backup selections. What you can do is create a Backup Selection Set and keep those selections available for future use, simply loading those selections by set name into a future job. We'll continue to work with a single job, for now. Simply select the incremental backup job you wish to use to create your consolidation, enter a name for the job and submit it. There are no specific backup options that you can select for this task, though you can use the other tabs to choose a target device and media, or schedule when this task should take place. Performing this with a set, instead of a job, runs pretty much the same way. The only real differences being that you would expand the sets folder, instead of jobs, and that you'll select the set that contains your incremental data in it. The rest stays the same. One point to call out is that NetVault will display this new consolidated backup as a standard full backup. If you wish to be able to identify this as a consolidated backup, you should consider that when naming the job, before submission.

The Data Copy section allows you to create multiple copies of previously completed backup tasks. This is a great way to achieve redundancy of your data and is a big feature in shops concerned with high availability. For example, you could use this plugin to allow you to keep one media set on disk, onsite, but also to copy that same backup to tape, for offsite disaster recovery. Much like the previously talked about data consolidation, you have the option of copying data from a single backup job or from a set. The rest of the procedure is much like the consolidation process, which helps reduce the learning curve of this software, by keeping task set up uniform across many services. Use of the Data Copy option of the Duplication option that is executed as the second step of a backup job is desirable when backup windows must be kept to a minimum.

The next topic, and plugin, we'll discuss is the NetVault Databases plugin. This plugin specifically deals with the database that keeps track of all your critical NetVault data. Needless to say, this is a critical part of your NetVault server. If you were ever in need of restoring your NetVault server, this information will be invaluable to you, so it is highly recommended that you be sure to back this database up as often as you deem appropriate in your environment. The procedure for backing up the database works much like all other jobs in NetVault, however, you want to be

sure that this job is the only job running, when you perform the backup. This is vital to getting a clean backup of the database. A good practice is to schedule this job at the end of each days backup window to be sure the most recent information is included. Simply select the database out of the list of NetVault Server nodes available and submit the job. There will be no specific backup options available for this task, though you can schedule your task for later completion, or repetitive completion and pick out what target to use, etc.

The final plugin in this list is the Raw Device plugin. This plugin allows you to perform a physical block level backup, or restoration, of physical disk or partitions. The Raw Device plugin exhibits higher performance, hence shorter backups, than a File System plugin backing up the same set of data because the file system overhead associated with opening and closing individual files is eliminated. One drawback is that individual files are not able to be recovered from a raw backup. Instead, you must recover a physical disk or partition in its entirety. This plugin typically has two primary uses in practice: as part of a disaster recovery to restore your disks to a baseline image after a damaged server has been replaced and reconfigured, for example; or when needing to backup a very large number of relatively small files quickly and the ability to restore those files individually is not a paramount concern.

My overall impression of BakBone's NetVault is a good one. The product feels very mature and has many features for power hungry administrators. Beyond simply touting an extensive feature set, NetVault delivers impressive functionality and exceptional ease of use. There is ample documentation for the NetVault product available and all of the folks who I spoke with at BakBone were exceptionally well versed in the product and helpful; two traits that you'll appreciate, if you ever need support for your installation. The NetVault GUI is well refined and, if you use the Mac OS X Server tools regularly, you'll truly feel right at home in the NetVault environment visually. NetVault crosses many lines in the grid of Mac OS X backup solutions. The ability to setup the system fairly quickly will make it exceptionally attractive to most any SMB's and the scalability and features will make it viable in the enterprise.

Pros: Server Admin-like interface, strong feature set, various plugins available, replication features, quick to pickup and use, very good support features, documentation readily available.

Cons: Some GUI features can feel less than intuitive, media recycling can sometimes lead to confusing results, some settings/tasks feel repetitive.

Archiware's PresSTORE 3

A flexible solution that can fit many needs.

In what appears to be a blend of the features, looks and capabilities of all the solutions in this paper, Archiware's PresSTORE provides a very balanced ,versatile backup, archival and synchronization system, that now boasts an excellent solution for a mobile workplace too.

When you receive PresSTORE, you'll notice that it doesn't require a lot of space to install and the installation takes just a moment, at least on the Mac Pro we're using for this test. The software installs itself in /usr/local/aw and runs on port 8000 by default, as you connect to PresSTORE via your favorite web browser that's named IE, Safari or Firefox. If port 8000 doesn't work for you, you can move the application to a different port fairly easily. All you do is issue a simple move command in the Terminal to change the configuration file (located in /usr/local/aw/conf/) from lexsrv.8000 to lexsrv.DesiredPortNumberHere and restart the application. Restarting the application is easy and accessible via the command line, as well. Inside of /usr/local/aw are three handy tools. These self evident tools are called “stop-server”, “start-server” and “restart-server”, making it quite easy to control the system remotely. Upon restart, my application was properly running on port 8500. The included tools also provide feedback, letting you know as the application stop, starts and what port it is currently active on.

Connecting to the web interface was easy and as advertised. Logging into the application is accomplished by using one of the administrative accounts in Mac OS X itself. Upon connecting to the server, I'm immediately told that the license is not present, and that the services will be dormant until one is installed. While the previous version had a a very sort of Mac OS 9/Classic look, the new GUI is much more refined and pleasant to look at. This includes more subtle colors and better looking icons to help guide you through the system. Again present in version 3 of the system is the ability to use the "First Steps" option on the left hand side of the GUI. While the previous build used a less than obvious need for expansion of that heading to continue, the new GUI simply shows you an option to license the software as your #1 item in the First Steps list. Personally, I find this to be helpful. The system provides you with the options you need to continue and doesn't distract you with the items you'll need for later configuration yet.

Entering a license is pretty simple, you simply click on the button labeled New in the licensing screen. That brings up a new view in the browser, showing you the License Configuration Screen. The tool is pretty straight forward, looking for a serial number, license key and a selection of the product type you are registering. Previously this screen also asked for contact information, which I had some “trouble” with. In the older version of the product, there wasn't an option for selecting the United States as your country of origin, which caused me to have to take

up “residency” in England. This process has been completely removed from the new version 3, which helps get you into the actual configuration of your system quicker.

Speaking of licensing, PresSTORE has five basic types of licenses: Backup, Archive, Synchronize, Backup2Go and Professional. Most of these licenses are pretty evident in what they provide. The Professional license includes all three traditional backup modules (all but the Backup2Go module), as well as licensing for a changer and a medium drive. To be noted, the licensing includes a single client, so you will have to purchase additional client licenses for the nodes on your network. There is a Demo license that provides all of this functionality for you to try, as well.

Now that I have licensed the product, we're set to start. The right column has changed. While the First Steps originally only listed License the Product before, it now shows three primary areas of configuration under First Steps: Synchronize, Backup and Archive. Note, you'll also see the Backup2Go module if it was included in your licensing. Each of these primary topics of data protection includes two to three sub-topics. We'll begin with what I'd anticipate is the most common use of PresStore; Backup. Archiware anticipates that setting up a backup solution will take three main steps: configuring the storage devices, adding the clients and setting up the backup plans. As each step is numbered, this helps a first time user understand what they'll need to accomplish to configure each of the three primary functions. Adding the storage device was quite easy. You begin by selecting the type of storage device you wish to use, a jukebox or a disk. We'll continue with a disk based backup. When selecting a disk, you can elect to limit the space used for the PresSTORE system or you can elect to dedicate the entire disk for use. You can also label the volume as backup, archive or leave the volume unlabeled. Upon completion of these basic questions, a new library will show up in your storage devices. If you elect to add a jukebox, the system will automatically scan the system for the proper hardware to add. I don't have any such devices attached and the system reassured me of this fact.

Going along with the basic backup setup, we'll now add in our clients. By default, you'll see the local backup server as the first available node on the network. To add in additional hosts, you'll have to install the PresSTORE software on the additional nodes. Since the installation is done from a standard package, you can use Apple Remote Desktop to push the software out to your other clients and servers. The installation can be done without user interaction and the client is ready to use without any further setup, beyond installing the package on the node. All of the licensing information is taken care of on the server.

You can now continue to the final portion of setting up a backup; defining the backup plans. As with the rest of the setup, this is pretty straightforward. Upon clicking the New button to begin the initial configuration, you'll be presented with a pretty simple setup screen. Begin by naming your backup plan and click apply to unlock the rest of the configuration. You now continue to setup the backup task, which is the actual definition of the job. Select the host(s) you wish to backup and select the directories you wish to backup. By default, the system will backup all of

the directories on the hosts you select, but selecting your scope is as easy as navigating through a file browser and selecting the files/directories you wish to include. The only limitation that I found to this was an inability to select multiple directories at the same time. For example, I could select the directory that housed all of my shared user homes on my Leopard installation, but I was unable to select multiple specific directories, without adding them one at a time. Luckily, once you setup your selections, you shouldn't have to revisit the selection process. The process is pretty painless overall. You then proceed to define the backup even, which is the schedule and type of backup. You have three types of backup that you can choose, including Full, Incremental and Synthetic. There are numerous scheduling options you can select from. At the most basic settings, you can define a date and time to start your backup. In the true sense of a "backup window", you not only can specify when to kick off any backup job, but you can also specify how long you will allow a specific job to run. The use of this feature may not seem to be evident at first, but if you want to be sure that your backup doesn't continue into your work day, this is a handy feature to have. Moving forward, you have the option of defining a frequency for your backup, which can be daily, weekly, monthly or you can elect to not include a backup frequency. Each one of the options, besides "none", allows further refinement. You can elect to run a backup every X days or every X weeks on Y day of the week. You can also elect to run your backup every X months on a specified day on first/second/third/forth/last week of the month. This made setting up a basic schedule quite easy to follow. In the same way you define when you want to backup, you can define when you don't want to backup. The exclusion rules work much like the inclusion ones do. This allows you to setup much more complex schedules fairly painlessly. Finally, you select the storage you wish to use for the backup. Once that's complete, you can begin backing up data!

Since, we can, we will. Kicking off a backup job was easy. You can elect to go by the schedule you just defined or you can manually start a job at any time. The Job Monitor provided the information I would want to see, primarily the the start time, amount of data to be backed up, elapsed time, transfer rate and the current file that's being backed up. My job, which was fairly small (about 6GB of data) was done very quickly.

Now that I completed a small, full backup, I went back into my original backup plan and added an incremental and a synthetic full. The directories were already defined for me, since this is all part of that single plan, so I only had to define the additional jobs as a part of it. The incremental plan worked properly and properly picked up the files that I added and changed. Running the synthetic full, now that I had a prior full and at least one incremental backup was easy, as well. The synthetic kicked off two simultaneous jobs; one restoration job and one backup job. If you aren't familiar with synthetic backups, is the same concept as you'd see in both the BakBone and Atempo offerings, where the system is able to consolidate the incremental data you have taken along with the last full backup to create a new, point in time, full backup. This avoids any load on your client systems and network, as everything is performed locally on the secondary media you write your backups to. This was completed in a reasonable amount of time.

Like I've always said, it doesn't matter how many times you can backup, it really matters how many times you can restore. Since I have a few backups that supposedly were successful, it seems like a good time to try restoring some data. I tried deleting single files, full directories and everything in between. No matter what I deleted, PresSTORE brought back; all of the permissions and attributes intact.

You can also define more elaborate restoration jobs, you have more options of how and where your data is to be restored. First, you can select where you'd like the data to be restored, such as to the localhost or to a specific client. Second, you can select a path to restore to, in the case that you wish to restore your data to an alternate location, instead of to the original place it was backed up from. Third, you can choose how to resolve data conflicts, if the data you are restoring matches the data pre-existing in that location. These options include discarding the restored file, adding an `_R` to the recovered file, overwriting the existing file and overwriting the existing file only if the recovered file is newer. These options should satisfy most folks.

When you've made your selections, you're ready to restore. PresSTORE will find the file you requested and present you with an overview, such as the file size and where it plans to restore it to. Of importance is the Volume Listing information. This is where PresSTORE has found your file to be residing. If that volume is currently offline, you'll need to bring it back online before restoring. I had thought, since my media was setup to be automatically mounted when needed that this would be handled for me. Unfortunately, it was necessary for me to manually go back and mount the correct pool. Not a huge deal, but one more thing to consider, remember and navigate to. In the end, I was able to restore about a dozen photographs that I previously backed up without much trouble. Once I mounted the volume, the restore took a matter of a couple of seconds to complete.

With that now established, we'll try a synchronization job. Setting this up was very easy. First you pick your source data. You then pick your target location, which can be locally attached or remote. Then you pick the days and times you wish to have the synchronization process take place. Conversely, you can also run your synchronizations on demand. I found that this worked as well as a traditional backup and restore did in PresSTORE. My data was synchronized in a reasonable amount of time with the proper permissions and attributes intact.

Moving away from the primary uses of the system, the Administration pane provided the tools I'd expect to have, all clearly labeled. The Client Manager included a new feature, allowing me to update my PresSTORE clients to the newest software builds directly from the administrative console. This tool also allowed me to set data encryption for backups of the individual nodes on my network. It also allowed me to setup restoration path options. In addition, it also includes a ping feature.

Another part of the Administration section is the Media Pools. You can create a Pool out of tape devices or disk. Creating a Pool is pretty easy and resembles creating the disk. You'll set a name

for the Pool, define if you want it to be used for backup or archiving and select your media type.. You can also restrict the Pool to only use specific drives, that you define. In addition, you can set data cloning right from this screen, which also allows you to pick the drives you wish to use for the cloning procedure. In the additional options, amongst other things, you can select your data recycling policy, allowing you to reuse your storage either when the backup has expired or allowing the system to thin out your backups as storage space is needed. There are also parallelizing options, in which you can choose to allow multiple data streams per drive in this area, allowing multiple jobs to write to your pool at the same time.

In addition to backup selections, you can also select to archive your data. This option allows you to set directories to archive and also allows you to select options that include the ability to remove the original files as well. The PresSTORE software also allows you to include a selected preview size to be generated of the data you are archiving. For video and audio uses, this allows you to select a preview time in seconds. The software will take care of including a unique date/time stamp, if you wish to have one included and includes the ability to assign pre and post flight scripts to execute before and after the archiving process.

While Synthetic fulls are able to reduce your backup window by consolidating your incremental and full data on your secondary disk, PresSTORE also includes the capability to define filters that can be used to reduce the scope of your backup. You can base your filters on items such as file name, modification date , size criteria and the like. This filter actually works both ways, by allowing you to specify files to select with specific expressions and also to exclude specific files based on the expressions you create. For example, you can specify to eliminate any files that end in .mp3, so that you are sure you are not backing up someone's music collection, even if they've picked up on the fact that you exclude the Music directory in their home. Just another feature that PresSTORE includes to help make your job easier.

New to the PresSTORE line of software is Backup2Go, Archiware's answer to the question of backing up a mobile workforce. As more and more enterprises are electing for the flexibility of portable computers, the need for flexible, robust backup solutions becomes exponentially more critical. Traditional backup solutions rely upon their own rules and schedules to backup clients. Portables bring their own requirements as they are, as the name would indicate, portable.

Backup2Go is licensed as an add-on to the existing PresSTORE software solution and works in the opposite way that regular backup does. Backup2Go enables the end client machine to initiate the backup process. If the client machine is not connected to the network, it will attempt to begin the backup later. If the client machine is removed from the network during backup, the PresSTORE software is able to maintain the history of the attempted backup and continue it from that point, when able to connect to the server again.

Setting up the Backup2Go is fairly simple and very straight forward. To start, the PresSTORE admin modifies a Backup2Go template, which is included in the installation. Like the rest of the

PresSTORE system works, you simply select “Edit Backup2Go template” from the administrative console, located in the left hand navigation bar. The Generic Template that’s included can be modified to meet your environmental needs. You are also able to define more than one template, if you so choose.

To get started, define the location where you’d like your Backup2Go backups to be stored. One note about this; you should keep your Backup2Go data separate from your traditional backup data. For example, don’t use the same directory where your PresSTORE backup disk storage is located; define a new directory or use a different disk for storage. I’ve found that mixing the two types of backup together can be problematic and isn’t within Archiware’s best practice. After that, you can define as much or as little as you like. You have the ability to specify directories that you require to be backed up, filter specific file types and set the interval at which backups will be taken of the client. You also have the ability to limit the bandwidth used, as well as compressing and encrypting the data. In addition, Backup2Go is able to use File System Events for backup. What’s File System Events? File System Events is an API that’s available for people developing for the Mac OS X platform to use to track file changes. PresSTORE uses this API to be able to quickly identify files that have been changed since the last backup time stamp. While you don’t necessarily have to use this option, it’s great to see developers fine tuning their code to use the resources made available to them on the platform.

On the client side, end users can use a shortcut to open their own Workstation Admin console inside of Safari. Entering the Backup2Go server’s IP or DNS name, along with a username and password that’s able to access the server completes the basic setup of the connection to the backup server itself. If the administrator hasn’t defined the backup scope, the end user is able to navigate into their machine to select what they’d like to include. In all of my testing, I included my own user folder, as I would anticipate many would. A simple click of the “Apply” button and you’re on your way!

Restoring your data is simple as well. Simply open the user interface in Safari on the client side, click on Restore and select the files you’d like to bring back. It’s just that simple. No need to call IT, no need to submit a help request. The solution works, is fast and is simple enough that non-technical people I’ve shown it to were able to use it without guidance.

On the server side, the admin is able to login to the traditional PresSTORE backup admin GUI and track the Backup2Go jobs. Simply click on the Workstation Manager under Administration and view all of the joined workstations and their current status, including results of the last backup attempts and time since the last backup was completed. End users can get similar information in the Workstation Admin console, being able to see the logs from the previous backup jobs and attempts.

While I’ve been able to use the software on Mac OS X 10.5 Leopard Client and Server, PresSTORE has really built this solution based on use of the ZFS filesystem. Running the soft-

ware on Leopard has provided me with a point in time mirror of my client, but only one of them is available to me in the live system. For example, if I backup my files at 8am and then again at 8pm and then want to restore a file at 9pm, the 8pm version is only available to me. I could theoretically migrate the earlier data to tape or some other storage as a versioning measure, but that will limit the usefulness of the end user console for restoration, as they'd then also need the admin to restore the previous Backup2Go data to the server for them to access. In essence, it would make the software less friendly for everyone to use. Through the use of ZFS, PresSTORE is going to be able to present the end user with a point in time view of their data, allowing for versioning of files that the end user can access. With ZFS, the end user will be able to get that 8am copy back without the need to call IT or for IT to devise a way to keep Backup2Go versions themselves. The end user will be able click on the date and time stamp of the snapshot they wish to view and will be able to view their data at that point in time, as well as restore.

Overall, the new PresSTORE 3 takes a product that was seemingly simplistic to use and truly makes it easy to use. The documentation of the solution has also improved greatly with this new build. I think there are a lot of pluses to this system. With the ability to backup Mac, Unix, Windows and Linux clients, the software is versatile and fits nicely into heterogeneous environments. The new Open File Backup for Windows provides support for the Microsoft Volume Shadow Copy Service allows the backup of Exchange servers and other databases. The new version also builds upon a fairly impressive list of supported file systems, now including support for Novell Storage Services for Linux, to complement previous support for ZFS, Xsan, NTFS, HFS and others. The web interface is fairly well done, very responsive and mostly feels like a local application. Being a web interface, this does allow you a lot of flexibility of being able to work with your PresSTORE installation from virtually anywhere in your enterprise. Support is extended to browsers including Safari, FireFox and IE, which should make most everyone happy. When thinking about the web interface, my mind wandered over to the next logical place to run my backups from: the iPhone. I wish I could say that the PresSTORE solution ran flawlessly on the iPhone, but I did run into a few issues. While the pages rendered properly, performing double clicks on the file browser to permit a specific restoration job was problematic. I'd hope that Archiware could incorporate a way to open directories from the File menu within the software. Otherwise, I was able to monitor my jobs and various other items in the system from my iPhone without too much trouble. From what I've seen, PresSTORE has a very powerful backup tool with a very easy to understand and administrate interface.

Pros: Full featured, fairly intuitive and easy to use, VSS integration, agents available for multiple platforms, flexible administration through the web interface, intuitive solution for backing up and restoring client data and portable computers.

Cons: Limited application/database support, web interface isn't quite as iPhone ready as I'd hoped, some people will still prefer a dedicated management application.

Conclusions

Overall impressions and final thoughts.

Illustrated throughout this document are numerous ways to view the same basic conceptual idea; data backup and recovery. In the long run, there is no right and wrong way to accomplish this, there are simply better ways to do this in different environments. The basic defining factor is, if you are able to accurately back up and restore your data.

All of the solutions discussed in this document will accomplish your basic backup goals for you, though my experiences obviously varied greatly from one product to the next. I was afforded the opportunity to speak to the individual who codes the OS X agent for TSM. According to him, he was unable to recreate my permissions issue, so this may have been a complete fluke event. He also suggested that he is able to pull about 4GB per 10-15 minutes over 100Mb ethernet, which was drastically higher than the 3GB per hour we saw during our live demo with TSM. Our TSM reps also suggest that there is at least a weeks worth of consulting necessary to properly install a TSM server, which comes at a fairly large cost. At the time of this writing, I had requested a copy of the software, as Tolis Group and Atempo provided, to test in our environment, but this would only be granted if I agreed to pay for the week's consulting charges. It can be viewed that the group had already granted me a day's time to review the software, so why should they allot any more time to my account, or you could consider the performance that their product displayed during their demonstration and wonder why they weren't more eager to show off the true capabilities of their product.

BRU, while starting off to a rough start, turned out to be a solid, simple solution. Tolis Group has had no qualms with providing tech support and fielding my questions. Their product is by far the simplest in writing to tape and the most cost effective, if you just want to simply write to tape and get your data off site. They have very attractive pricing that includes 25 agents to start with and a free download for you to test out.

Atempo's offering brings many enterprise level capabilities, fast backup and restore times and support to backup numerous platforms and applications, such as Exchange and MS SQL. If you are deploying Xsan in a large enterprise environment, you will probably want the added features and capabilities that Time Navigator provides over BRU and you'll appreciate the native OS X functionality and ease of use over TSM. With their continued support and development for the Mac OS X platform, Atempo has improved their product substantially since their X11 beginnings. Many of the many configuration options are now presented in more friendly, easy to understand ways or are completed for you, without the additional administrative attention. Their move to Aqua makes the solution feel more refined and the features and tools included further

justify that feeling. This package has come a very long way and, if you ran into glitches and annoyances with X11 on your first trial with Atempo's product, you owe it to yourself to look at this solution again. You'll find other solutions with a lower starting cost, but you'll find more scalability, features and support with this product than most of those.

BakBone's NetVault can truly fit into almost any environment. The GUI has a familiar look and BakBone packed a lot of substance behind it. NetVault won't be found in places where BRU might be overkill, but it will be found in any environment that would like more granular control over their backup solution than BRU can offer. Seeing as this product seems to have matured so quickly, you cannot question the dedication that BakBone has shown to the Mac OS X platform. This product could very easily make backup one of the most enjoyable parts of your day.

Archiware's PresSTORE is the only solution to provide a full web interface for the entire product. This gives you a portal to your backup and restore center from any supported web browser. While many may be leery of the overall performance and experience provided by this type of solution, it should be noted that this application feels much more like a locally installed binary than a simple web form. The administrative interface was found to be quick and consistent. For folks who have looked at PresSTORE in the past and may have run into some quirks, the new version proved to run smoother and with a more refined GUI and better documentation, that can help get you up and running fairly quickly. Add in the new Backup2Go features, a key component that I believe will make PresSTORE a more recognized name in the enterprise, and you have a very complete solution that can easily address nearly any backup need that an administrator may have. Overall I found this to be a very nice solution to work with and I was able to reliably run my my backup and restore jobs without much hassle or thought. The GUI is pretty simple, but the features and tools are fairly robust. I feel this product could fit very well into many different environments, though the lack of specific application plugins and open file handling may be a concern for certain enterprises.