

BACKUP, FRONT-TO-BACK

PENNSTATE



MACADMINS
CONFERENCE
2013



Allister Banks *abanks@318.com*  *@sacrilicious*

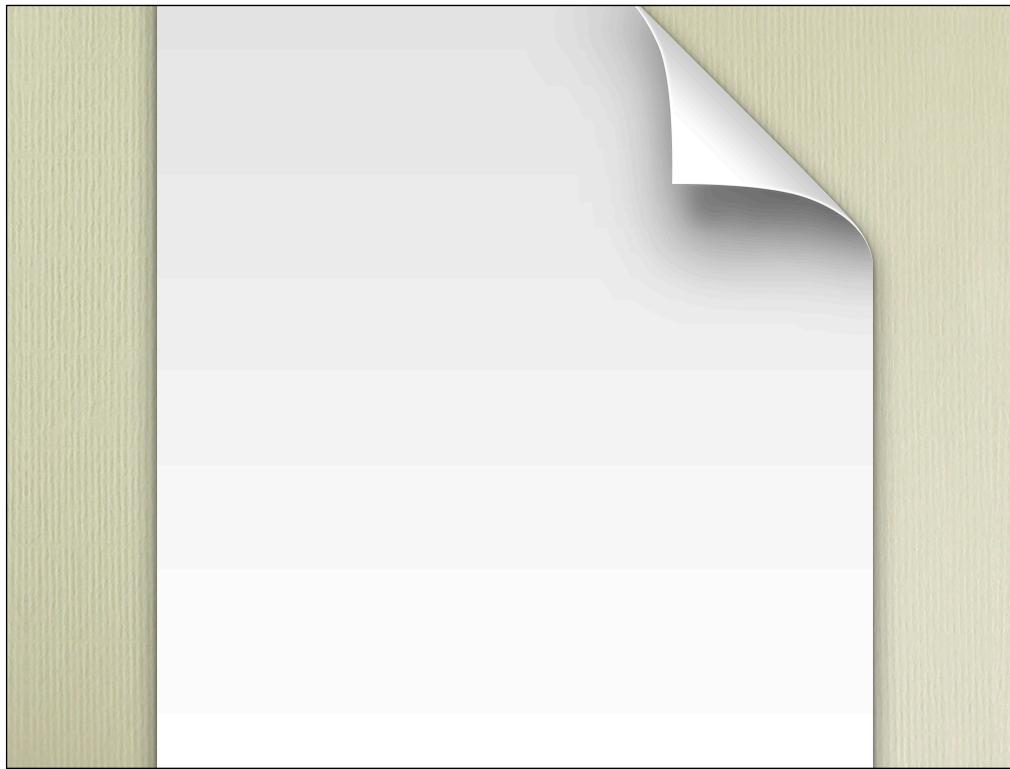
- What works to protect data
- Who are the players
- How it's made

1

In my session description I said we'll be going over what models work to help protect data,

who are the players,

and in general 'how it's made' Hopefully we'll leave this discussion with a firm agreement on the terms to discuss solutions in the future, and a plan to implement not just appropriate collections of archived data, but to verify the data regularly to ensure gravity works.



2

So we start off talking about why you'd even need to implement a project like backup is that you've done an architectural investigation into the common unit that backups operate on for your organizations: data –

Information Lifecycle Management

3

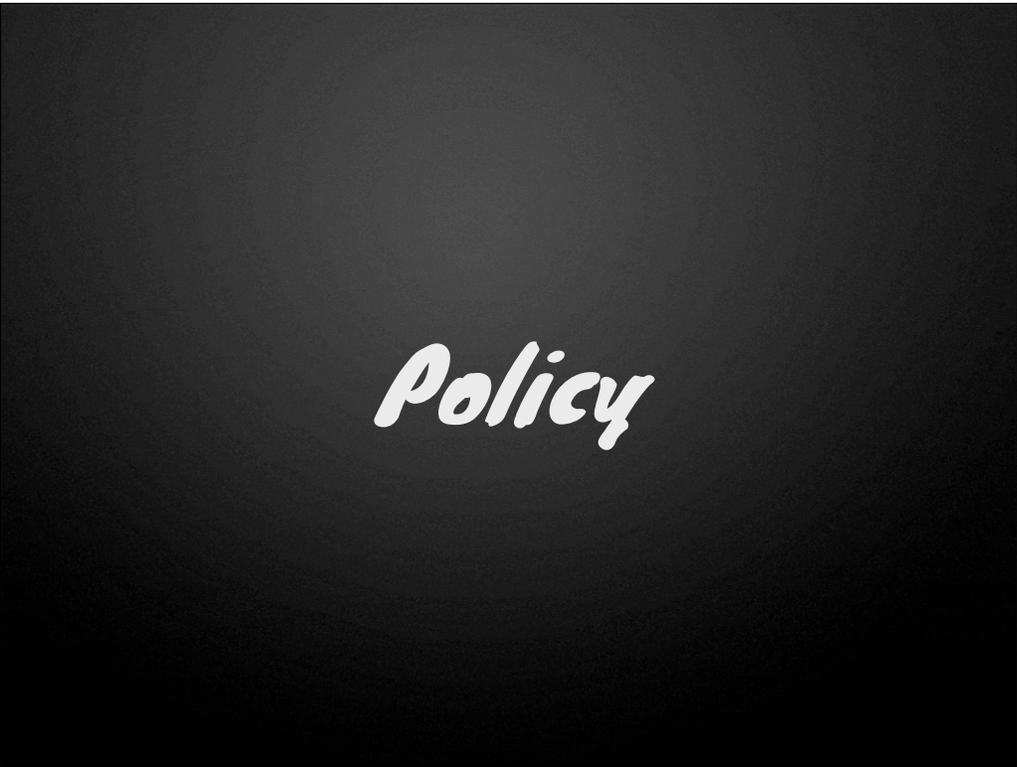
and the study of what it's all about is semi-buzzword termed information lifecycle management. There's asset management and workflow systems that have set a precedent, but I decided to move forward with a somewhat independent method, although it's nothing special. It's just about getting answers to the common questions across the system you support, in what I'll term the investigation stage:

Investigate

4

investigate

- How is data grouped, by whom, for what purpose and how long is it actively worked on?
- What file types, sizes and fidelity/metadata is in use
- does data production rate exceed available storage on a ongoing basis
- what expectation for whole copies/changes to retain in nearline/offline



Policy

5

draft policy

– SLA's, understanding of worst case scenarios



Plan

6

plan

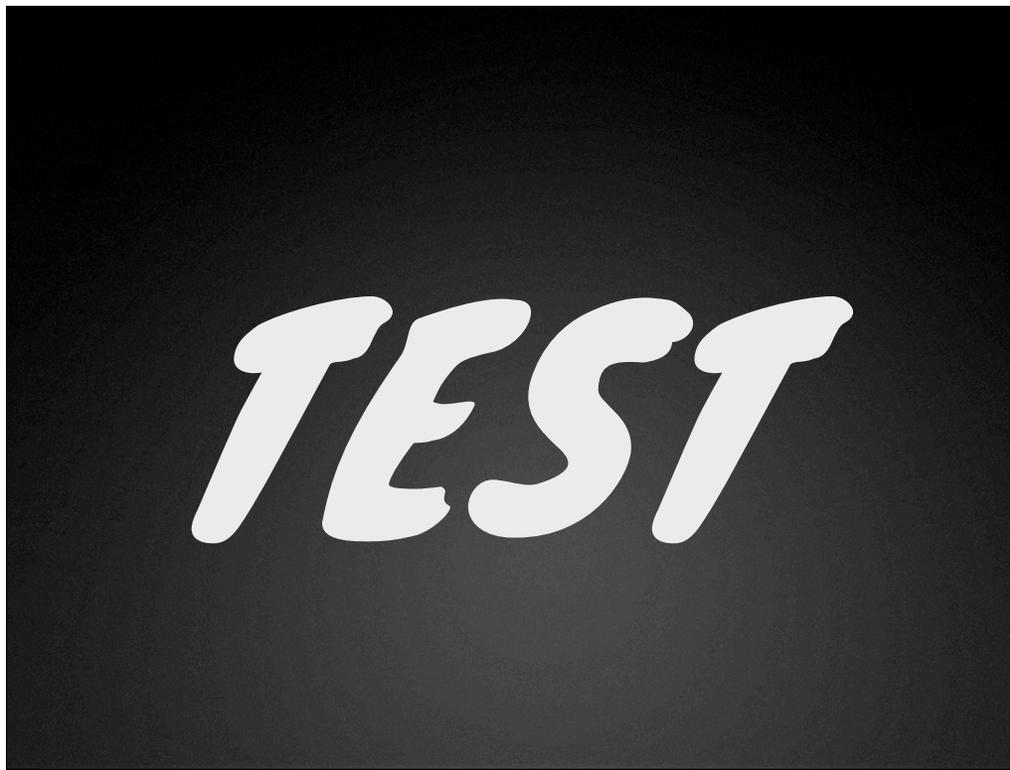
– arrange/evaluate/order necessary equipment/system

Implement

7

implement

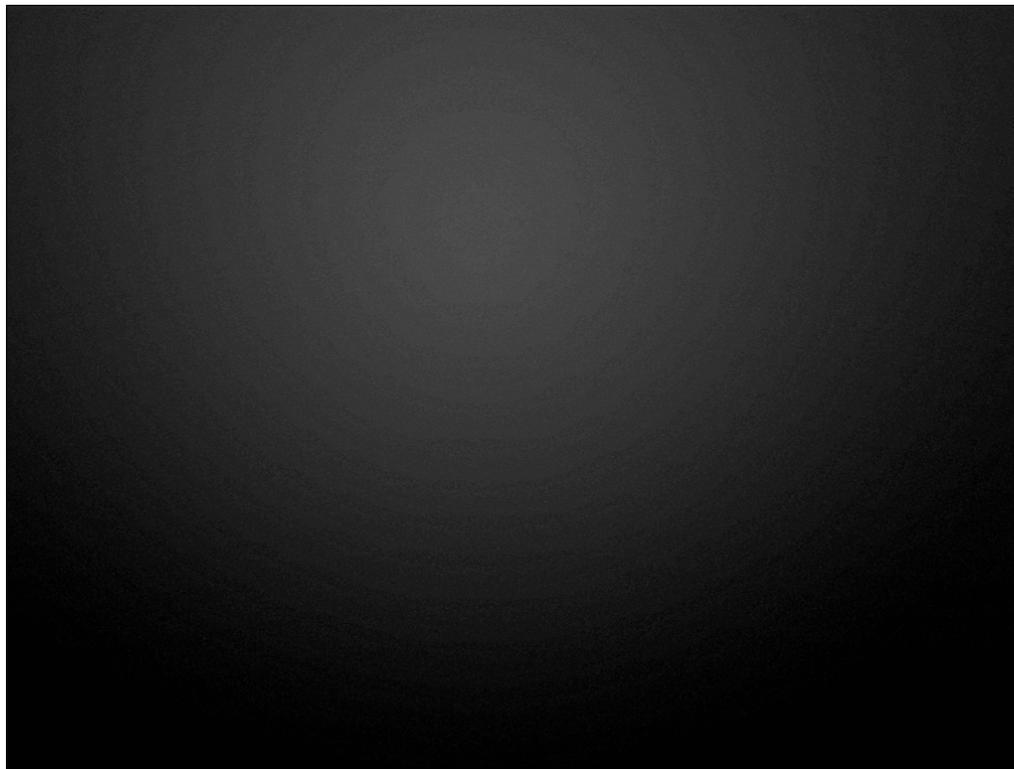
– put system in place, including scheduling and automation – humans are the severely limiting factor, so as much that you can keep in mind for out-engineering humans without building up insurmountable technical debt, the better



8

test

– verify proper retention, availability, fidelity, affected parties can have documentation/training pushed to them. At LOPSA earlier this month the local-to-Manhattan sysadmin for the Stack Overflow suite of sites talked about how lucky they were to have their DR site up and nearly automated when Hurricane Sandy hit and the lessons they learned about getting their people comfortable and communicating the disaster recovery process. It's like when a PR nightmare happens to a company – you do not hire a firm to handle the fallout after your CEO or Dean refrained from putting his foot in his mouth and offended all your customers, the PR firm you already trust has gone through those scenarios so they're ready to extinguish the flames before they burn your organization down.



9

As a result of a going through with the pre-flight steps and get to milestones like the above, which is usually when you have real buy-in and quite possibly after what's phrased in sales as an ahem

gnu.org/fun/jokes/error-haiku.html

11

"Three things are certain: Death, taxes, and lost data. Guess which has occurred" <http://www.gnu.org/fun/jokes/error-haiku.html>

The organizations we work for produce goods or services, and information revolving around those products is considered worth holding onto for at least a reasonable amount of time, which may be closely related to human memory as the practices of an organization may evolve over time. There was a great, thought-provoking post by Mark Pilgrim about how just upgrading software rendered previous formats incompatible and data unrecoverable, so he adopted Linux, in particular Ubuntu and hoped to invest in longevity by moving all his work into open formats. And there are other people having to adapt drivers inside emulators of past systems to hook up arcane tape players to first extract data and then go through the process of converting the recovered data to a format that can be manipulated by the tools we have today. So I think about things like the lifecycle of data and the worth in what we're creating.

In the Beginning...

12

But to get back on a mac-focused path, let me do a brief rundown of my experiences for what I consider the credentials I hold to be able to talk on the subject. Like a lot of us, in early OS X days I recommended Retrospect and as a general practice still use Carbon Copy Cloner, fast forward to recent times and I've been tinkering with the backup products on the windows and linux server platforms, with other products for the mac appearing with fits and starts to address archival and the challenge of backing up laptops at scale. And I'll end my talk with three stories of success to tell, and the appropriate caveats.

One-Offs

13

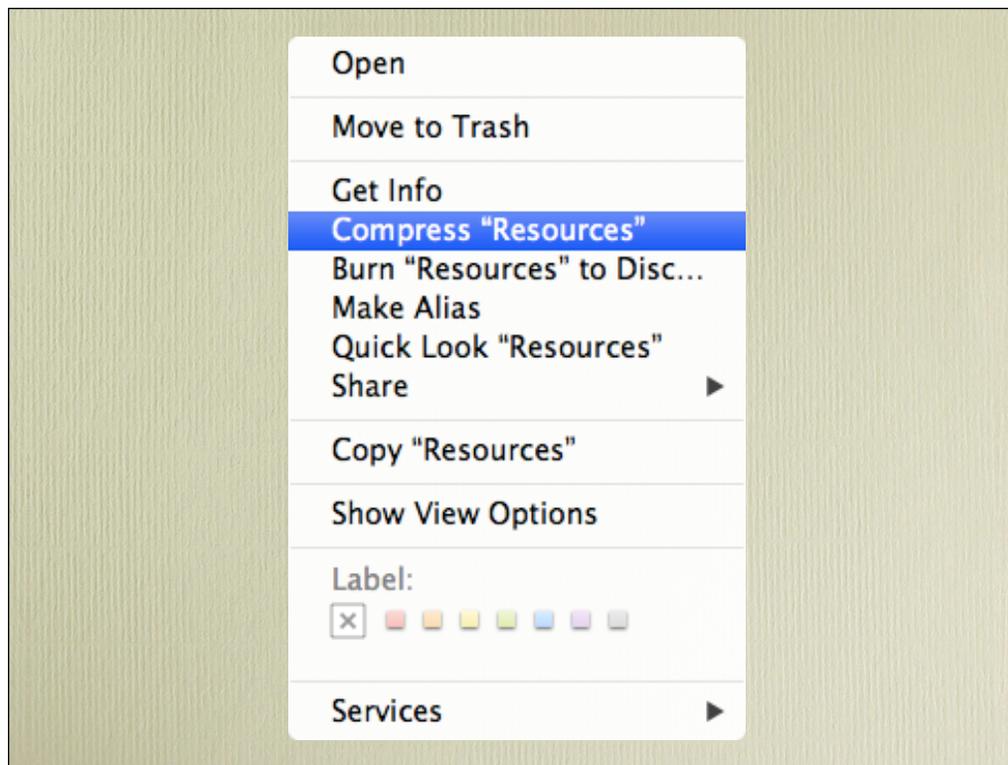
Instead of going straight into those products, let's see what we use day to day from the command line and one-off tasks in the GUI and build from there.

The 2 Problems:

- Cache Invalidation
- Naming Things
- Off-by-one errors

14

One-Off's don't scale, and if it's not part of an automated system that's suitable for use by humans you're not particularly better off, but at least these things help with data fidelity, keeping it from corrupting, which is one of those computer science issues that come up from time to time



15

When you're at a break-fix level, you'll hear about split-half search and other troubleshooting techniques for example to clear up a bad plist file in the library folder. One way to keep data fidelity when trying to isolate an issue is to move half of the contents of the library folder out of the way and do a logout/and back in to see if the problem is resolved, or if it's system wide. A practice I got in to was to right click the Library folder and 'compress' it, by which the Finder meant throw it in a zip file, and that essentially takes a snapshot of the entire folder before I go and touch anything.

CLI FTW

16

More geek cred is to, from the command line, use `cp` to copy the file you'll modify to `file.bak`, and that's the simplest baseline we can start with.



cp

17

More geek cred is to, from the command line, use `cp` to copy the file you'll modify to `file.bak`, and that's the simplest baseline we can start with.



18

Next is the obviously old-school tape archive utility tar, which apple has kept updated to some extent. It can compress the files it touches like the more modern zip and gzip, and mirror a hierarchy inside a container with various other options. It was obviously tuned for specific media in a time that has passed, but when you just need a more reliable method of moving data into a container it or the zip binary comes in handy.

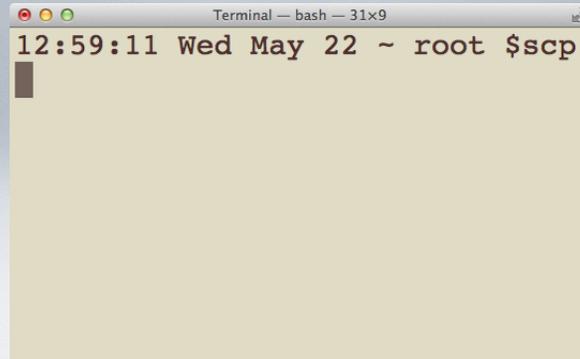


ditto

19

And next is ditto, which is a more optimized way to call cp with the a flag, and pull across the directory structure. Metadata handling was fixed for most of the unix utils in 10.4, but had always worked with ditto to preserve modification and other extended attributes for Mac files. One of the other things that would trip command line beginners up (and still gets me) is that ditto creates folders that aren't there when you try to copy the source to a destination path that only partially exists, whereas unix utils care if there's a slash at the end of the path.

Network

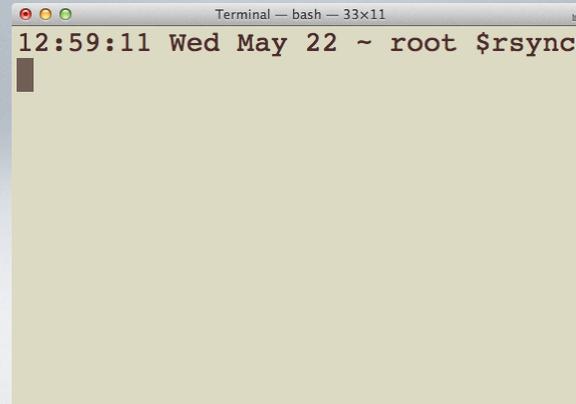
A terminal window titled "Terminal — bash — 31x9" is shown. The prompt is "12:59:11 Wed May 22 ~ root \$" and the command "scp" has been entered. A cursor is visible at the end of the command.

```
Terminal — bash — 31x9
12:59:11 Wed May 22 ~ root $scp
```

20

For moving files over the network an early standout was scp, which you could feed a host that you can ssh to for either source or destination and it would prompt you as needed to get the data moved 'quick and dirty' as they say.

Pater Omnium



```
Terminal — bash — 33x11
12:59:11 Wed May 22 ~ root $rsync
```

21

But the real mack daddy of the unix utils is rsync, which is totally network aware and optimized, compressing the data it sends if possible and giving more options to resume transfers and handle things like exclusions more intelligently. The one drawback for an out-of-box mac is

Current: 3.0.7

```
01:10:05 Wed May 22 ~ root $rsync --version  
rsync version 2.6.9 protocol version 29  
Copyright (C) 1996-2006 by Andrew Tridgell,
```

22

the horridly ancient version bundled, from when apple was still primarily shipping power pc – does that sound like a long enough time ago? Building your own requires a few patches to handle apple metadata, but that's been well documented by the creator of Carbon Copy Cloner, Mike Bombich,

Alternatively: Use C.C.C.'s



23

and his own custom-patched version is accessible if you specifically use the entire path to call his rsync, although it can have real verbose output.

A terminal window titled "Terminal -- less -- 79x15" displays the manual page for the 'tmutil' utility. The page is titled "TMUTIL(8) BSD System Manager's Manual TMUTIL(8)". It includes sections for NAME, SYNOPSIS, and DESCRIPTION. The NAME section states "tmutil -- Time Machine utility". The SYNOPSIS section shows "tmutil verb [options]". The DESCRIPTION section explains that tmutil provides methods for controlling and interacting with Time Machine, including restoring data, editing exclusions, and comparing backups. A cursor is visible at the end of the text in the DESCRIPTION section.

```
TMUTIL(8) BSD System Manager's Manual TMUTIL(8)
NAME
  tmutil -- Time Machine utility
SYNOPSIS
  tmutil verb [options]
DESCRIPTION
  tmutil provides methods of controlling and interacting with Time Machine,
  as well as examining and manipulating Time Machine backups. Common abili-
  ties include restoring data from backups, editing exclusions, and compar-
  ing backups.
:
```

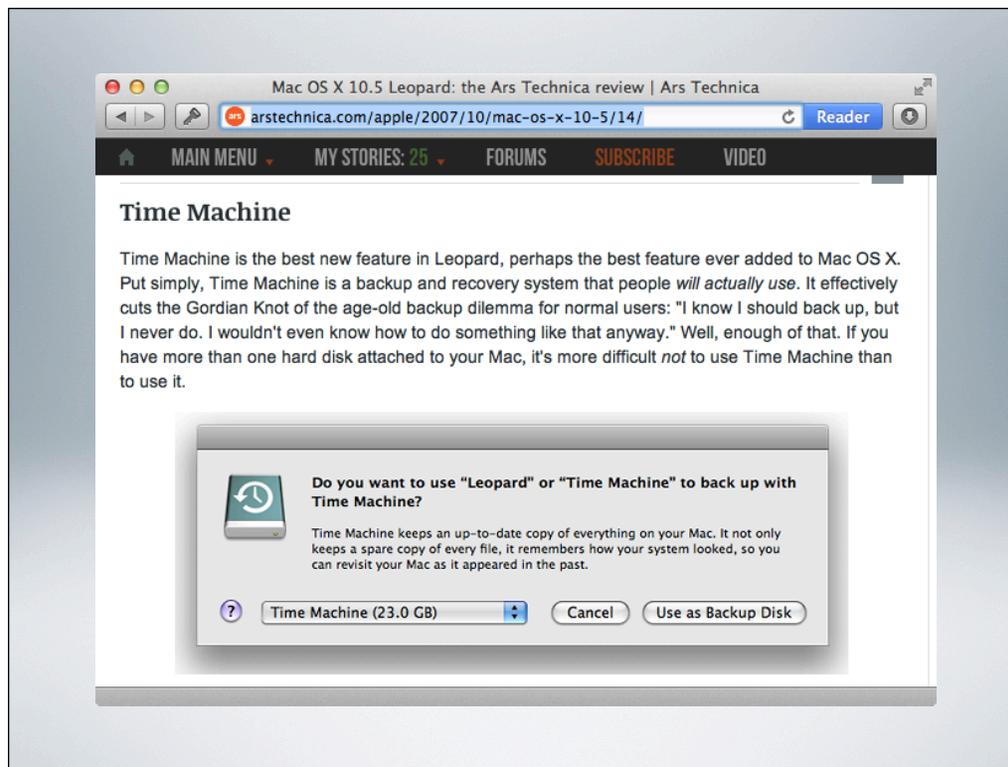
24

And I wouldn't be covering all the bases unless I mentioned tmutil, which has somewhat limited utility for interacting with



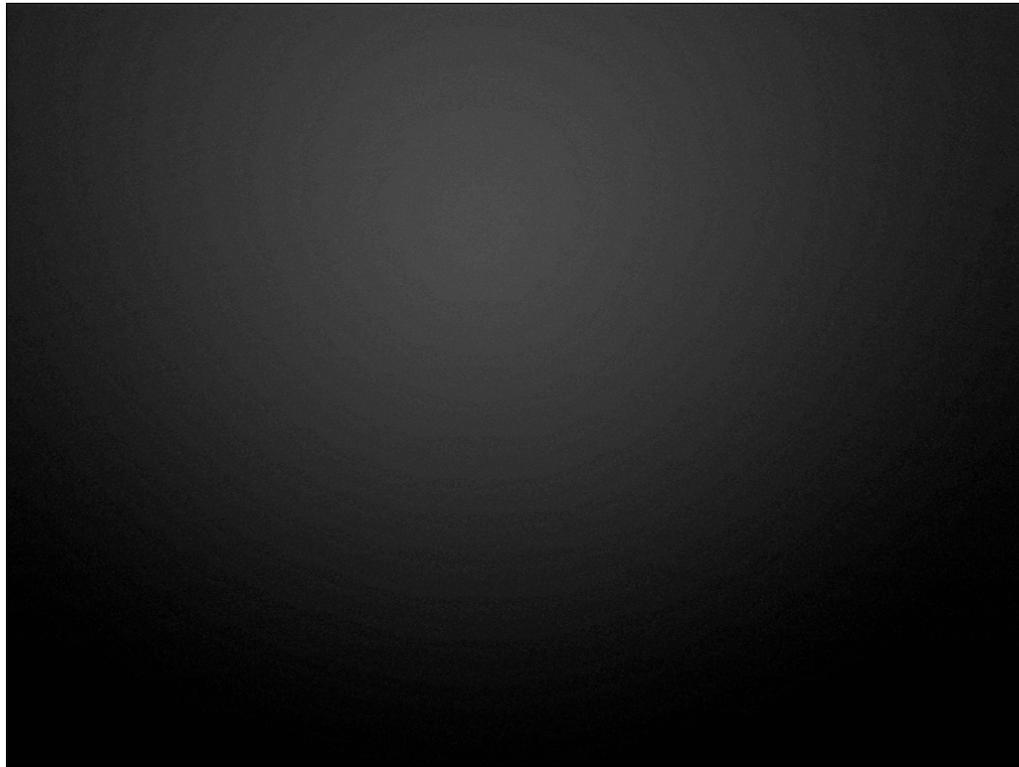
25

time machine. Ah, time machine, and the last time since Filevault 2 that I can recall a major plumbing change turning into a user-facing, marquee feature that radically optimized how the OS keeps track of itself, with spotlight_s mds, the metadata server, making waves before that. Filesystem event tracking, driven by dtrace, was lovingly combined with the old school unix ln commands hard link feature. And it has been... an interesting experience to try to deploy for anything besides direct-attached, single-workstation use, to say the least.



26

Since we're talking about a geeky-under-the-covers OS X feature it would be heretical not to mention the article where this was explained to a young padawan like myself, page fourteen. <http://arstechnica.com/apple/2007/10/mac-os-x-10-5/14/> And really, he warms you up with both dtrace and fsevents in earlier pages, it's understandable with something complex at its heart to take a bit and cover what Time Machine does, but I'm a natural pessimist: I'll cover what it doesn't do:



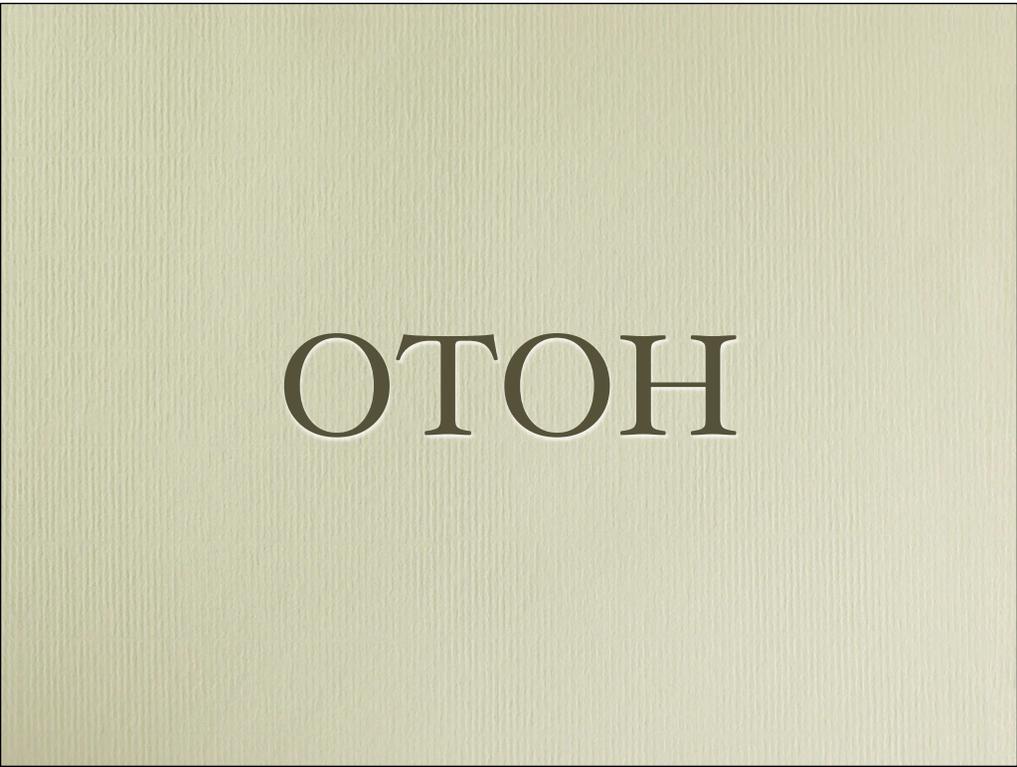
27

In practice, it doesn't treat your network connection well, nor does it deal with less than optimal networking conditions if you want to send those backups to a server. Heck, they've never even exposed quota's to subdivide the data a user can back up before older versions are tossed. At an even more low level, it's also dependent on what fsevents tell it, which can be troublesome. Even when directly attached, there's been countless reports of

“Gee, I hope...”

28

'I get to guess if this drive is failing at the source or destination and if time machine was doing the right thing not just in the past but recently enough that I don't lose data.' Or the even more common other symptom of 'I'm tired of fighting with this thing, dump my versions and start over because I'm too fed up to get another drive to start fresh on'.



OTOH

29

On the other hand, the simplicity of the implementation means you can pull out contents of past backups quite granularly, and there's a new optimization that if you run it on a mountain Lion server it triggers a script to archive the directory and postgresql databases daily if needed. Oh, and if it works, because until 2.2.1 it had a improperly configured launchd trigger and didn't consistently either 1. after first being triggered on setup, ever run a second time, kindof important if you , you know, add of change users,



[flickr.com/photos/treehouse1977/615766222](https://www.flickr.com/photos/treehouse1977/615766222)

30

or 2. updated the associated log file, I mean, if a backup runs in the night and no log file hears it did it make a sound?

To analyze the failures we can take away from TimeMachine as a part of our data management strategy it would be good to find a proportionately simple goal, and when I started out I'd explain backup to customers by saying it's just like Noahs Ark, we want two of every animal.

3-2-1

*[hanselman.com/blog/
TheComputerBackupRuleOfThree
.aspx](http://hanselman.com/blog/TheComputerBackupRuleOfThree.aspx)*

31

More recently for me I heard the concept people refer to as 3-2-1,

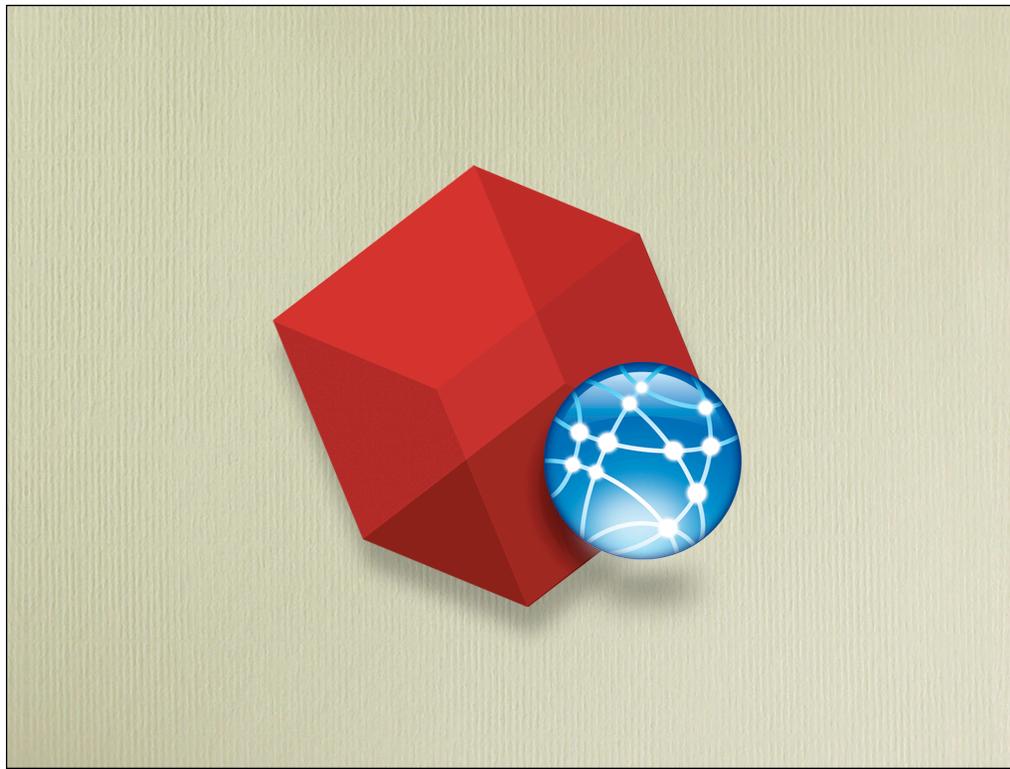
IF IT'S IMPORTANT

- Three copies of the data
- On two types of 'media'
- One of which is offsite

32

three copies of the important data using two different types of dedicated 'media', one of which is definitely offsite.

How does time machine stack up? While you can hack it, it only recently allowed you to have a second local backup destination, which wouldn't take over from a pre-existing one, on hard drive media(or if you can get OS X server backup to work at scale... good luck with that), and unless you get over-the-network to work over the WAN(which, if anyone in here has that working at this moment anywhere on non-google fiber internet connections, or even over regular wireless on the LAN I will eat my hat). So okay, maybe to be generous it can satisfy one out of three



33

Alright another craptacular product for backup at scale was Retrospect – but the failure isn't what it purported to do, it's the model it was designed with. Remember when every desktop was leashed to a desk and ethernet connection, and digital cameras weren't popular and business data was the size of text files? That's when it was designed, and an agent was installed on clients so that when the server went and knocked on its door a sending session could be initiated from the client and packed away in proprietary containers governed by catalog files because we liked one extra thing to have to juggle. The indexing implementation detail, man, backing that up as well was fun,

A file that big?
It might be very useful.
But now it is gone.

34

because if that got corrupted, as the tech haiku says A(n archive) file that big? It might be very useful. But now it is gone.



35

Oh and hoping clients could send their backups on a schedule lacked foresight, because over the past decade laptops got very popular, and they have this pesky habit of not being reachable when you close the lid. So – ha hah, Retrospect outsmarted those pesky laptops – they came up with something called Proactive backup! Does anyone remember that if you're unlucky? It would poll clients and slow them down for hopefully long enough, having all manner of issues pulling that backup over wireless... yeah, the whole model is just a kludge. It could also push backups offsite with wonderfully secure FTP... before they were bought either the second or third time and tried to integrate Mozy offsite backup, which is the one reason I've never trusted Mozy, sorry if folks in here use them it's probably fine...

Alright, I'm going to leave my bitterness behind now and move on to success stories of products that have worked for my company for information lifecycle management



36

One thing I had to get over about the first solution I'm espousing, Archiware's P5 or the artist formerly known as Presstore, is I learned to stop worrying and love tape as media – with the 3-2-1 strategy it isn't the ONLY media I rely on, and there are capabilities that the hardware engineering present that can be economical which outweigh some negatives like time til initial start of retrieval or writing, and I've just gotten over the ubiquity of SATA drives giving me a sense of security regarding turnaround time interacting with backup media, it's not enough of a worry in practice, although it is still a bit too much manual interaction, even with tape libraries.



37

Archiware also tries to be a bit of a jack-of-all trades, everything to everyone when other products do some portions better, but I don't fault them too much for it. And if you don't license it for the features you're not going to use it's greyed out in the interface and doesn't cost you anything extra. The only other real Caveat which is the same with CrashPlan is why in heck does it have to be so confusing to buy – Archiware goes through resellers and is based on slots in your tape library, or if your media is disks which works perfectly fine, it's licensed by the size of the destination volume you're writing to, and CrashPlan... I'll get into Code42 in a second. Archiware's flagship product is broken into four sections, Synchronize, Backup, Backup2Go and Archive.



38

I can't really speak to the laptop-focused 2Go offering although I've heard good things, but we rely heavily on all the others for our customers. We had one instance synchronize helped push a server's assets to distribution points, which is of course one way, but the ubiquitous file system like alexander FS for having a true DFS of multihomed but synchronized SAN storage always has seemed like a problem nobody had figured out without either considerable expense or incurring significant technical debt. Still, a sync process makes a lot of sense in a disaster recovery site. Backup has been employed to great effect when you need to move a lot of data by running to a set of tapes in parallel to increase throughput which has long been the most efficient way to move a lot of big contiguous data, and the great thing about Archiwares products are you can drive all of it from the command line, and setting it up without hand-holding hasn't been an issue, and support has responded with accuracy and in a timely fashion – in our experience.



BRU

39

Compare that with BRU or atempo or backbone. Those are not highly intuitive products to get going and Mac is not necessarily the platform you want to tell them you're using when you call. Backupexec is still our recent choice for windows because they swear physical to virtual and restoring a bare-metal backup to non-similar hardware now REALLY works. But seriously, it's because they have the modules to backup Exchange and MSSQL and it's a fine integrated experience.



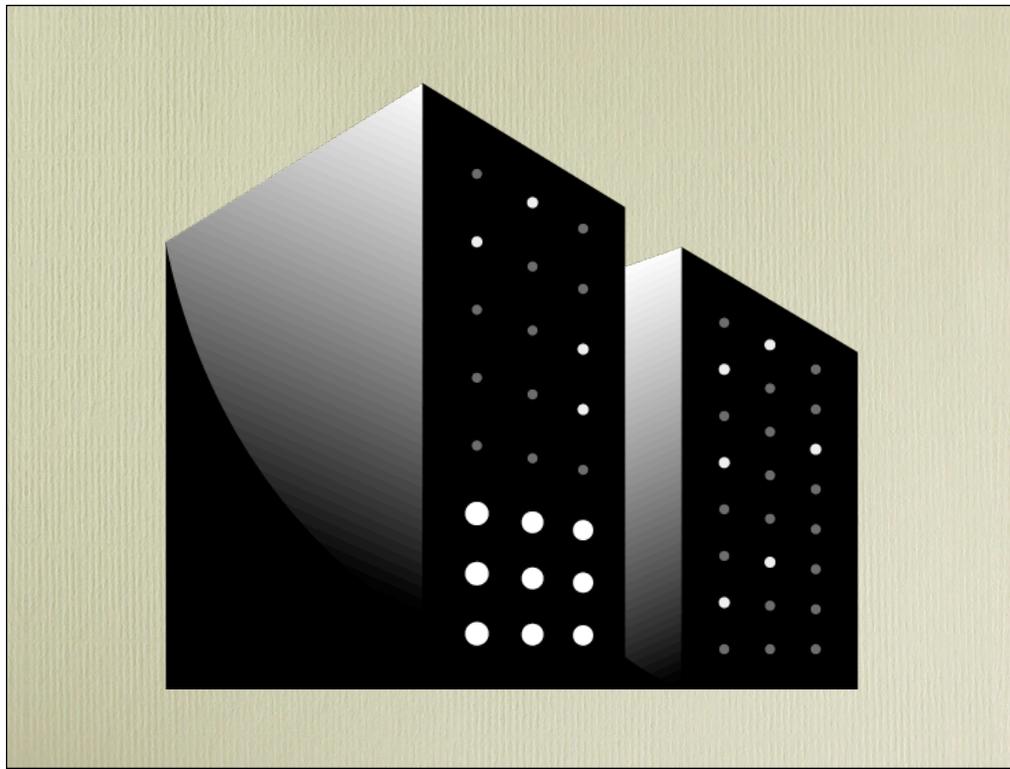
40

Which brings me to the point that we're only talking about Macs that are workstations for information lifecycle management, products that address a servers services will be limited – I won't be discussing restoring postgresql databases even though we've done it – there aren't popular quote-unquote off-the-shelf products I've come across for some of the database-driven services you see running on macs.

Archive

41

Back to Archiware, I don't do as much work with Archiware's archive product as my coworkers, but we did a lot of custom development with the command line options, and for the use case of taking data offline and/or putting it on a shelf for 30 years I feel ever slightly more confident it made it onto the tape 'as advertised' by using Presstore/slash P5. And in practice it's the cheaper price point than alternatives and mainly for that archiving use case, which is big when I'm continually talking about that buzzword-y info lifecycle mgmt, that's why we've had success with it, and what really makes it stand apart.



42

Now my biggest Caveat for CrashPlan when talking about it at this conference should have been that there's someone here giving a presentation on it. And they're from Code42, makers of CrashPlan! But the person they sent is from sales... losinghorns.com, we got totally gypped by it not being an engineer like Andrew Renz, who is a super onit tech with plenty of street cred. My current employer has never been partners with CrashPlan but we can commiserate if you've been following their sales practices over the years if you want to talk about the company offline, but I'll just try to separate emotion from the logical facts about the product, let's just say I believe I'll be covering valuable information even though there will be some repetition since I want to cover things for people who don't attend the guy from Code42's talk.



43

Anyway, CrashPlan is a great example of a somewhat specific use case, a simple-seeming but surprisingly complex feature set, and a brilliant software deployment model: from way back, they seem to beta test early releases or newer features on the free and consumer tier I call Green for obvious reasons, then move it to the business backup they host for Small to medium sized businesses that I affectionately refer to as Blue, and then you see the improvements in the Enterprise or black version, not unlike the black label 12 year whiskey.

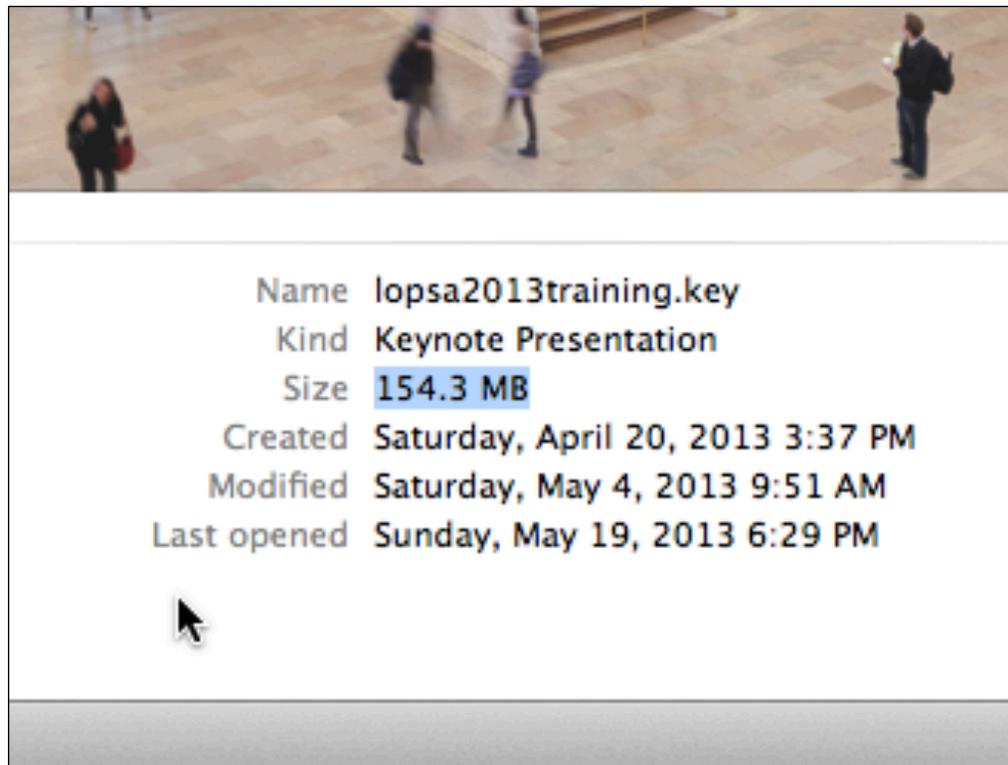
The five cent tour of the differentiating factors touted by all versions of the product are that it compresses, encrypts, and de-duplicates on the client end, but those are just a list of features, how does that benefit the process?

Marquee Features

- Less data to send
- Looks like gibberish
- Just what (*blocks*) changed

44

It sends a smaller version of the data, securely in a way that if intercepted or snooped on it wouldn't be easily unencrypted, and even inventories the data at the block level so it can look at other files in a users backup and ensure they don't send the same blocks of data twice –



45

what's the easiest way to describe that is, say I've got a 200MB powerpoint because I don't know how to lower image quality for a 1024 by 768 projector, and I give that presentation in Albuquerque on tuesday, and I say hello albuquerque on the first slide because I like saying albuquerque – and then later I go to oklahoma city thursday and I dupe the presentation with a 'save as' but have only modified the first slide, the promise is it will only backup the distinctly updated or different blocks of data.

Product Comparison - CrashPlan PRO - Backup for Home, Business & Enterprise

www.crashplan.com/business/compare.html

CRASHPLANPRO™

products | try / buy | support | company

There's a CrashPlan for everybody.
Home. Business. Enterprise.

Choose your CrashPlan backup software based on type of data, number of computers, backup storage requirements, and administration preferences.

FREE 30-DAY TRIAL
For Windows, Mac, Linux, Solaris



	CRASHPLAN+ Home	CRASHPLANPRO Business	CRASHPLANPROe Enterprise
Overview			
# of computers	up to 10	up to 200	100 - 100,000+
Licensed for	Personal data	Business use	Business use
Mac, Windows, Linux, Solaris	✓	✓	✓
Online mgmt. (end user)	✓	✓	✓
Real-time dashboard / reporting (administrator)	✓	✓	✓

46

And I said it was difficult to buy, like Archiware's artist formerly known as Presstore – here's the URL you want to go to if you're a consultant like me – the second row should cut down the selection somewhat, because you're violating EULA to send business data to CrashPlan's servers, so we're left with the choice between Blue and Black. Here's where it gets interesting – you cannot seed or get a drive with your archive on it back for a restore if you choose the blue pill. More interesting: Black is cheaper and assumes users will have multiple devices, up to 4 per user license. But, with black you need to start out bringing your own infrastructure – you can hook up with a code42 partner or code42 themselves to provide a second destination, but you need to own the 'master' server that links to LDAP and houses the accounting about your archives and all that.

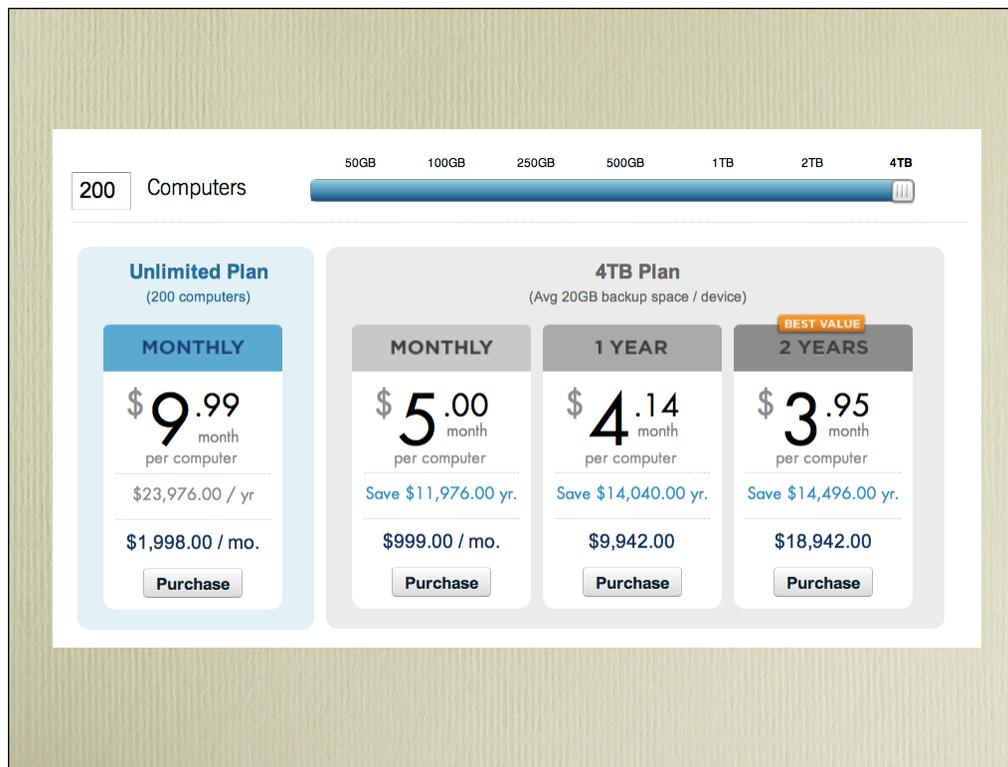
And you want a second destination – the server side is totally free, put 'em wherever you can afford the bandwidth and storage space, it's all encrypted, and their dead-simple seeding process makes that initial sending to a new destination a non-event.



47

But then there's blue, which is a cloud you don't control, You talk to some people about it and it's like a George Carlin routine, "where did you put it? I put it in the cloud!"

Blue is hard to sell customers on just like Simian hosted on google app engine is hard: capacity planning and initial usage estimates.



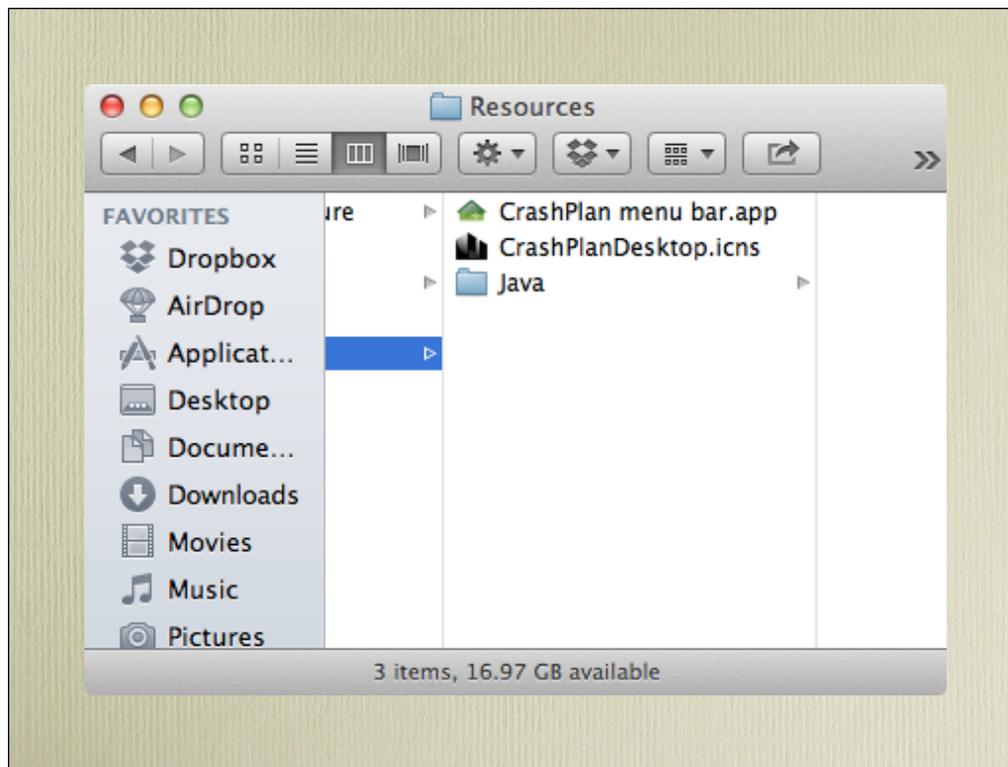
48

For Blue, you can either do the retail \$10 per month per workstation or unlimited computers that fit within a shared pool of up to 4TB. You can fluctuate over time and after a 30-day trial, you're being billed monthly so it isn't impossible to make sense of it, but it's hard to have a conversation about because you need a slide rule to calculate, and I don't know how to use a slide rule.

Caveats

49

Let's assume this audience and the folks playing along at home are all individuals that can get some infrastructure together, wouldn't that be safe to say? So black's not a challenge to implement, and you're wondering about other caveats and takeaways when approaching the Code42 solution:



50

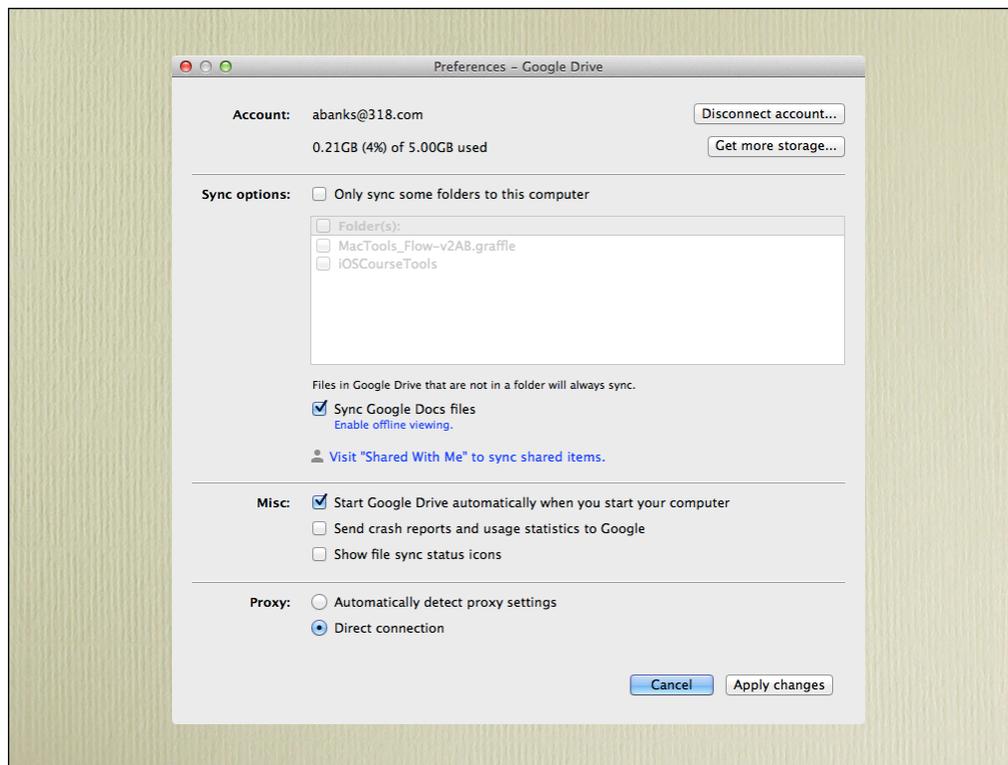
One that people keep, rightfully, bringing up is the fact the windows and mac client software is based on java. That probably won't go away any time soon server side, but they've committed to making some sort of progress with native clients. Being organizations that need solutions that work now, I'm not one to hold my breath, but it's nice to have a vendor acknowledge an issue and commit to a fix.



51

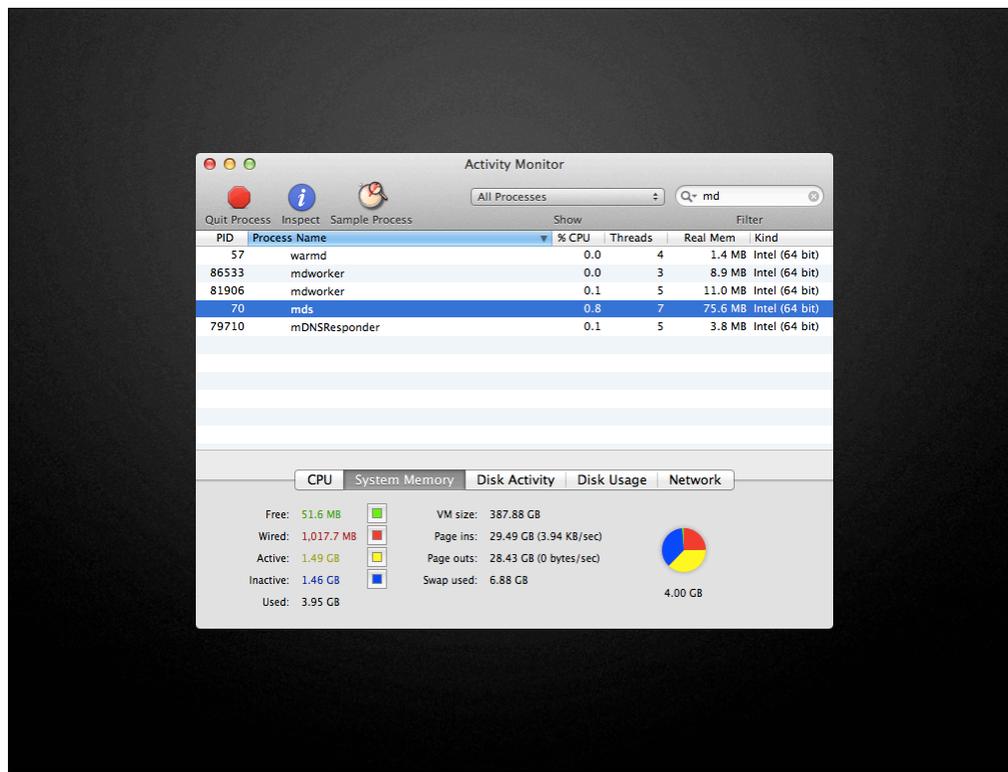
Another CProE Caveat is again, like Noah's Ark but a step further, you want two destinations to send to at almost all times, because the archives have maintenance performed on them server-side, during which you can't send. And it seriously burned someone I know who did not get in front of that before he had an ahem compelling event with the location his storage server was experienced indoor scattered showers. It used to be a lot more 'single threaded' metaphorically speaking to perform maintenance on archives from the server, so they have made efforts to lessen that failure scenario, but a second server with non-consecutive maintenance obviates the issue.

Two smaller points policy and not necessarily tool related are maybe you want to help end users not backup their Dropbox folder, or any other 'folder that syncs' service, as they almost always provide some kind of versioned backup server-side.



52

You'd also be doing them a favor by not leaving the synced folder at the root of their home folder with the allow traverse permissions set which lets any other users on the local system view the contents of files at the root of that folder – and google drive, I mean where do I even change it if I make the wrong decision after the initial setup? Only box.com maps to a folder in documents by default, would be a little better if the box service was actually intuitive to use in general IMHO, but I guess that's besides the point.



53

Also, what some may consider a plus of crashplan's client on the mac is that it does get tickled to pick up on new files in realtime with the same service spotlight uses... which may be a negative for some that find spotlight... less than reliable.

CrashPlan API Viewer

https://10.1.1.10:4285/apidocviewer/

Search API **CrashPlan API**

- Archive
- AuthToken
- BackupReport
- Computer
- ComputerBlock
- ComputerDeactivation
- DataKeyToken
- Destination
- DestinationAssignment
- DeviceUpgrade
- Diagnostic
- LoginToken
- Org
- OrgBlock
- OrgDeactivation
- OrgMoveProcess

Archive

Methods

get

Examples

```
GET Archive?storePointId=1
```

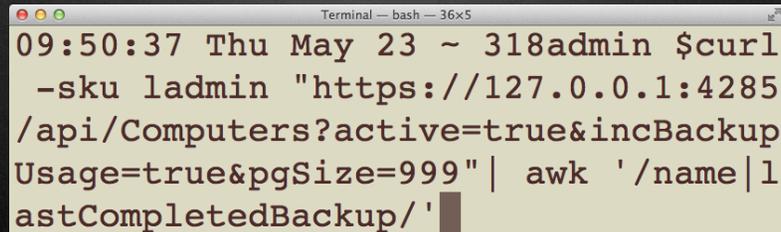
```
GET Archive?guid=9999999999999999&targetComputerId=1
```

Error Codes

Params

Name	Type	Default	Description
storePointId	number		returns archives in this store point
serverId	number		returns archives in this server

API updated as of: 2013-5-22 [Ref](#)

A terminal window with a dark background and a light-colored text area. The text shows a command being typed: '09:50:37 Thu May 23 ~ 318admin \$ curl -sku ladmin "https://127.0.0.1:4285 /api/Computers?active=true&incBackup Usage=true&pgSize=999" | awk '/name|lastCompletedBackup/''. The cursor is at the end of the command.

```
09:50:37 Thu May 23 ~ 318admin $ curl
-sku ladmin "https://127.0.0.1:4285
/api/Computers?active=true&incBackup
Usage=true&pgSize=999" | awk '/name|l
astCompletedBackup/'
```

55

on the server side it's all web based and ajax-y and web 2.0-y, and when you run Black you can take advantage of the API to query all the present data without logging in and waiting for AJAX, here's a one-liner I came up with just last week to get status across a bunch of clients without sending a report and tracking it down in email:

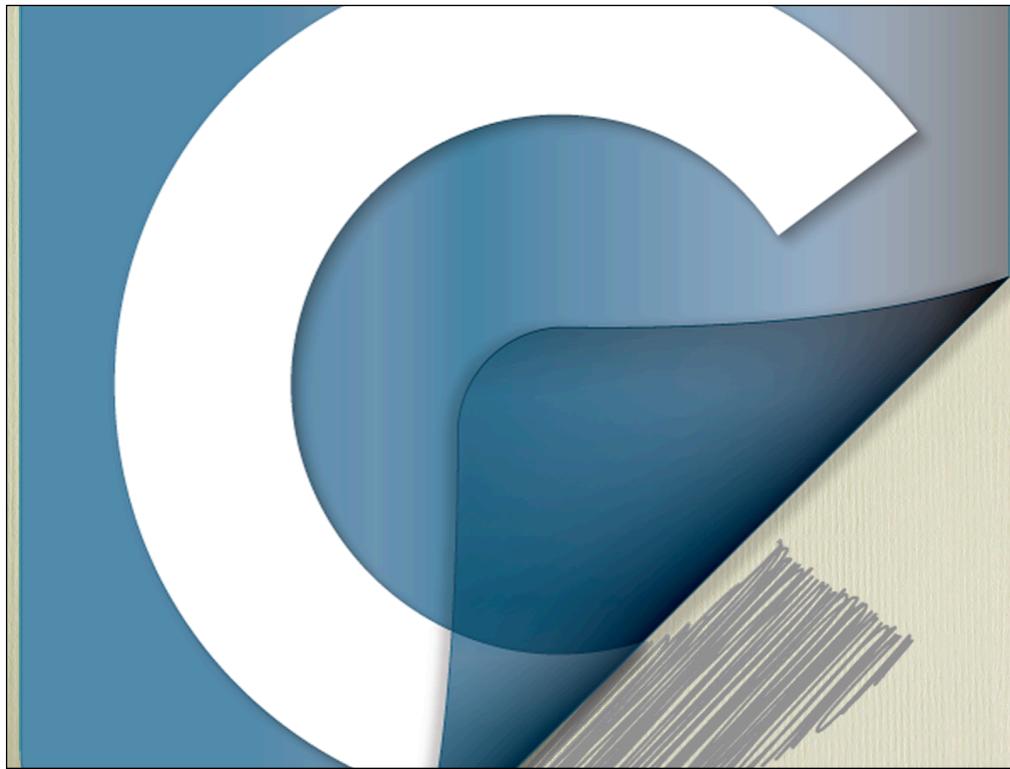
```
curl -sku ladmin "https://127.0.0.1:4285/api/Computers?
active=true&incBackupUsage=true&pgSize=999" | awk '/name|lastCompletedBackup/'
```

To take this from the end I'm awk'ing for both name as an identifier and last completed back, then from the front s is for supressing errors when running curl 'silently', k to suppress checking the SSL cert - and really, you should only be running this over the local network if you're paranoid about that, and we need to authenticate with u for a user with the access role of sysadmin.

API FTW

56

And this whole restful API interaction underlines a point I was making back when I mentioned laptops aren't leashed to a desk with an ethernet cable – you have to assume a web interface is more compatible for access and is WAN-optimized as a general design requirement, you interact with Archiware P5 almost completely in web pages, which may be strange to call up a web page for a local configuration, but you get over it.



57

And the last product we've had success with, just to wrap up and cover the bare-metal backup use case beside archiwares p5 backup is Carbon Copy Cloner. It's not the most slick, enterprise-optimized tool, but it does what it does better than anyone else in my opinion, and doesn't have an exclamation point in its name, no offense Super Duper. With CCC I commonly do what I call a lazy mirror instead of software mirrored raid because then, before I send corrupted data to another disk automatically, CCC will tell me, AND as a result of that setup there's less writes going to the destination drive which means if they were purchased at the same time or are of the same vintage you possibly have a lower chance of it failing at the exact same time as the original, because the duty cycle is going to be shorter, potentially much shorter on a Server's OS boot drive when you keep the service data separate.

BACKUP,
FRONT-TO-BACK
—
THANKS!



Allister Banks abanks@318.com  *@sacrilicious*

PENNSTATE



MACADMINS
CONFERENCE
2013