Why Hard Drives make more sense than CDs or DVDs; or in the end, you get burned more than the discs do

A Commentary

by Justin C. Lauzet

The purpose of this commentary is to analyze cost and time - and because *time is cost*, the "cost and cost" - of burning to optical media. It then discusses why the trustable reliability of optical media should be no longer than 2 years.

I. Cost

How much does it cost to burn to CDs?
For a spindle of 100 CDs you're probably going to pay from $10 to $30 at the corner computer outlet store - depending upon if you buy a name brand or not.

Let's say your data is important, so you go ahead and buy a $30, 100 disc spindle of 80 minute, 700MB CDs. You just bought 70GB of storage space for $30 - so, cost translates to about 43 cents per Gigabyte. Sounds great, right?

Let's assume that you're not going to be able to optimally fill 700MB each disc burn - let's say you're using these as archive discs, and your average is actually pretty high, and you manage to fit 650MB per disc - so the cost is now at 46 cents per Gigabyte. It still sounds OK up to this point.

Archiving 240GB of data will consume 370 discs at a cost of $111 (or an even $120 if you had to buy all 4 spindles).

Let's assume a 250GB drive costs $133. It sounds nearly dead even, costwise, at this point, right?

Not. Even. Close.
Even if you're willing to keep the discs on spindles after you're done and not buy fancy leather books for them - it's still not close. And we all know the cool leather cases are the best part of a massive optical disc collection.

If you have files greater than 700MB, you already know hard drives are the way to go - so let's say you have data collections of smaller, various sizes and you begin the arduous task of sorting your data into 700MB or smaller folders (370 of them). This doesn't - even for a moment - sound like fun.

Now I simply must enter a fantasy world of averages and assumptions - in the end, only you can decide how much your time is worth. Even if your time is worth only $5/hr, a 20 minute per disc average would cost you $616 in time. That estimate was just to sort the data in a sane and efficient way, and didn't include the time it takes to burn the discs - so increase the time by another 5 minutes per disc for burning and cost goes to $771. Add that to the $111 and cost per Gigabyte climbs to $3.68.

To be honest, I'm being overly generous with both time and cost estimates here. First, your time is probably worth more than $5/hr. Second, anyone would get bored and wander away (with each disc) during the burning process, making this job take weeks if not months. If your time is worth $20 per hour, then that 240 Gigabytes would cost you $3194 in CDs and time (again, without the cool leather cases), or $13.31 per Gigabyte.

You can buy a very cool WiebeTech hardware RAID enclosure like the RT5™ for less than that, expect it to be more portable, more accessible, have more than a TB of storage AND give you RAID 5 data redundancy.

Even if you "keep on top of it" by burning a 700MB CD every time it's convenient, it doesn't make the costs go away - you're still spending the time, and you still must store each of these CDs somewhere. And if you want to change any disc, it means burning it over again.

Talk about bad seek time...
Just think - what do you do when you want to actually access one of those 370 discs? Where is the file you really want?

DVDs - do they stand a chance?
Back in the day, there for a while, DVDs had a case. Declining hard drive prices, however, caught up to them as well.

A 100 disc spindle of DVDs gets us 400GB and a name brand will cost you $60. All of the above problems still exist.

With hard math, 240GB could mean 60 discs - but there's overhead and unburned space, so let's say it's 66 discs, or roughly $40 in blank DVDs.

The time costs are similar to above: 20 minutes for each disk and an estimated 10 minute burn. That means 33 hours, or $165 at our slave labor price of 5 bucks an hour. A total of about $205 to burn off those 66 discs and it would take just a little under a 40 hour work week to pull it off. I won't bother figuring at $20/hr - but go ahead, rearrange those numbers any way you want to.

It's expensive, and in the end you're still left with 66 discs and the inability to memorize where anything is on any of them.

For around $200, you can get a WiebeTech ToughTech™ 250GB storage enclosure with FlexMount™ shock absorbers. It'll have a single 250GB drive, be orders of magnitude faster

than DVDs, and you'll find your file by browsing to it rather than guessing which in the multitude of marker laden discs it could be on.

II. Reliability

When CDs first arrived they were hailed as having 100 year shelf lives. It's simply not true - especially for burned media. Consumers have demanded lower and lower prices on CDs, and have gotten what they paid for.

The media you buy with silver surfaces (instead of blue to burn on) aren't burned. They're pressed similar to old vinyl records - with pits and bumps. The processes used since the dawning of CDs have diminished in quality drastically - more data on a CD the same size means that we've made the track width even smaller, and new, lower cost techniques have pushed the 100 year number down. Admittedly, pressed discs are considerably more reliable than burnable optical media.

A quick look at how CDs work
When the reading laser passes over a bump, light reflects away from the sensor and reads a zero. When the reading laser passes over a pit (or flat area) the light reflects towards the sensor and it reads a one.

CD-Rs, on the other hand, work with a layer of dye instead of pits and bumps - and the type of dye varies by manufacturer. During a burn, the dye is darkened by being burned (to keep the reading laser from hitting the sensor) to create a zero, or left transparent to create a one.

Each dye type has different properties. For example, how affected they are by the environment such as ultraviolet light and temperature. Scratches to either side of the disc, or even using the wrong pen to label it can prove fatal and every brand of disc will behave differently when it comes to these factors.

The aging process begins when the disc is created, and is accelerated when you take it out of the package. Dust particulates on the disc during burn greatly effect the burn of a disc and change the life of the disc. A dust particle only a few microns wide will have an impact. The data track on a CD is half a micron wide, and each bit is less than 1 micron long. Any dust on the disc prevents bits from being properly written. Too many bad bits burned, and the disc is instantly a coaster. If it's on the threshold, which is all too often the case without us ever knowing, then it's only a matter of time before the disc is unreadable.

Yes, there are types of error correction (interleaving and parity checksums) that are used on CDs. But the truth is that any one of us can damage a disc with normal use - and even careful use. Google is rife with posts from users waking up one morning to find their well taken care of disc collection little more than a leather bound coaster museum when their discs have simply "gone bad" over time.

The length of time that I keep seeing is 2 years. There may be plenty of people who have

greater success, but I am not one of them. Many of my older, burned data CDs fail to read and I honestly haven't done much more than burn them and store them in what I considered a safe manner. Even if some discs can be read after 2 years, it's exceedingly clear that some cannot. Is it wise to expect the functional discs in this collection to survive even another year? I'm forced to say no.

Saving our data to potentially unreadable media is akin to playing russian roulette with our ones and zeros - you don't get to choose which discs go bad. The *trusted reliability* of important data burned to CDs should not be longer than 2 years.

There's no denying - hard drives have limited life, too. Nevertheless, it's much easier to locate, use, and protect data on hard drives than on optical media.

III. Conclusion

Optical media is expensive in the long run, a chore to burn to, scatters files across a plethora of painfully slow discs which must be accessed individually, and provides questionable (at best) reliability for our important files.

What is the solution? Back up your data to hard drives. Because humans prove constantly that backing up is hard to do (even me, I'm terrible at it), the solution I choose is a mirrored RAID with 2 hard drives. There are plenty of RAID solutions on our website, but SilverSATA™ II is the sweet spot. It mirrors in real time - all data written to the device actually writes to two drives. If one goes bad, you swap it out with no down time. Even better, it allows you to pull a drive tray and back up to a 3rd drive - the drive (and data) you pulled out is totally usable in a different enclosure.

If you're interested in reading about Drive Reliability, read our white paper titled Storage Enclosure Reliability. It's an in depth white paper from our CEO James Wiebe with technical mathematical analyses on many types of storage enclosures, including RAIDs.