If I keep files backed up on an external SSD, is it safe to assume that without damage to it, they'll last forever, are they safe there, permanently?

Dave Haynie

.

Updated February 6, 2019

Hardware Systems Engineer (Rajant + Commodore-Amiga + 20 years in startups)

Nope. All forms of flash memory degrade in a relatively short time. Could be ten years, could be less, particularly if the devices are somewhat worn out.

The Architecture of a Flash Cell

Flash memory is a relatively recent invention, so we don’t have as much history as you might think about longevity. But enough to know that flash is a temporary thing. This wasn’t a shock or even a problem, for several reasons, but largely because in its initial use, flash memory replaced EPROM (erasable programmable read-only memory), which itself had a limited longevity.

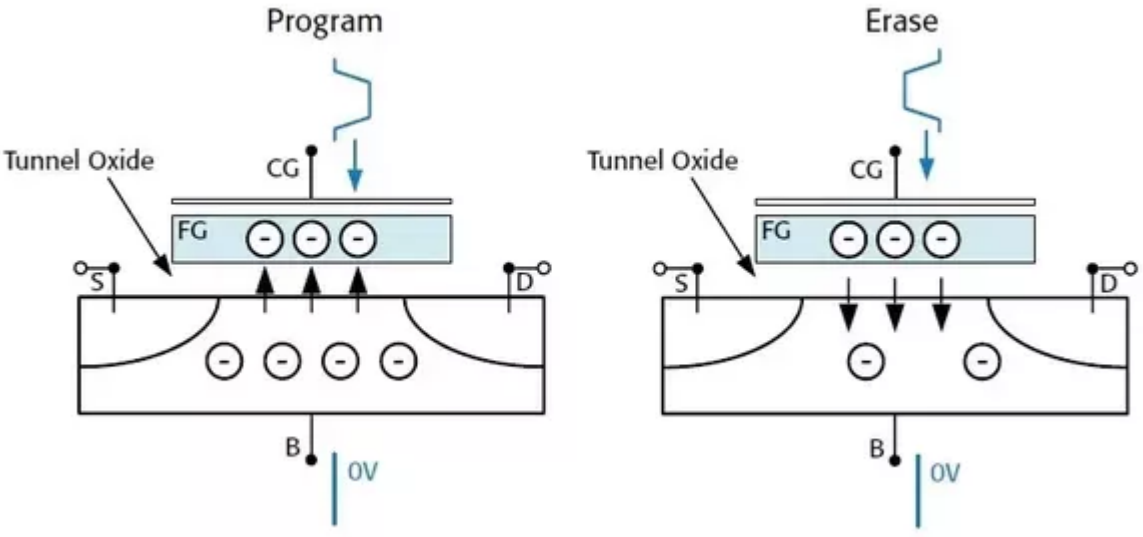

A flash memory cell is made of a MOSFET transistor with a “floating” gate. This is a gate surrounded by insulators. To write a bit, electrons tunnel through the insulator and from then are are stored on the floating gate in a process called channel hot-electron injection. To erase a cell, Fowler-Nordheim tunnelling is used.

The original sort of flash was retronymed “SLC”, for single level charge. This is what you’d expect out of a binary thing: there’s either a programmed charge on the gate (0) or no charge (1). An SLC flash cell should last about 90,000–100,000 program/erase (P/E) cycles…. some specialized parts can last 1,000,000 P/E cycles.

However, flash designers got more clever, and came up with “MLC” flash, for multi-level charge. Now they could write three different levels of charge, plus no charge, to deliver two bits per cell. However, this made the cells more fragile, less robust against corruption over time. More robust “enterprise” qualified MLC flash memories used in more expensive SSDs have a write/erase lifetime of about 20,000–30,000 write/erase cycles. Normal consumer oriented MLC flash memories deliver about 5,000-10,000 P/E cycles.

And even more recently, we have “TLC” flash, which puts in seven levels of charge or none, to store three bits per cell. That’s one of the reasons flash memory got so cheap. But you’re going see cells failing after 3,000–5,000 P/E cycles.

And here’s the other thing… flash has, for awhile, been getting worse. The highest density TLC chips may only deliver 1,000 P/E cycles these days. The wear mechanism in flash memory is primarily due to stress on the tunnel oxide layer in each cell. The oxide layer degrades as electrons, over time, become trapped in the oxide. And as we’ve been shrinking the chip geometries of these cells, their longevity has been dropping. Curiously, when a drive runs hotter, there’s less stress on the insulator, so the overall effect is a longer lived drive… but shorter data life.

Data Retention

So as mentioned, there’s a charge — a bunch of electons — placed on an insulated gate. That’s your storage mechanism. There are only so many electrons, and over time, they’re gradually going to leak off the gate. At the high end, the specialized, small flash memories in embedded microcontrollers can be expected to last 20 years to 100 years at 25C. On the other hand, Intel’s recommends unpowered consumer SSDs for data retention of only a year. And this is exacerbated as the drive takes on wear.

Now, it’s a bit different when the drive is active. Between wear levelling and data enhancing routines, an active drive will last considerably longer than a year, even the static data stored on that drive. There’s quite a bit going on within a modern SSD than you never actually see, which is designed to spread wear out across the drive.

And the trend toward less reliable flash has started to reverse itself, at least for now. The 14nm or so node is about as small as a practical NAND cell can go today, even though we have smaller chip geometries. They just get too unreliable. So the larger memories coming out today, 256GB per chip and so, are actually 3D chips using larger geometries. Some of these stack 48 bits worth of data on separate layers, and the reliablity is up.

I’m still expecting some future memory technology to get competitive with flash on storage and speed, but eliminate the wear and longevity issues. But for now, I’d look elsewhere for archival.

So, Where to Archive?

Well, I'd pick a hard drive over an SSD any day of the week. The magnetic storage on an HDD is extremely robust. That's likely to be good for well over a decade. The problem is, you're dependent on the mechanical bits of the drive tolerating that 10 year sit. You may be fine, or you may be subject to "stiction" problems, where lubricants get sticky over the years and prevent a normal spin-up.

My long-term medium of choice is Blu-ray, with qualifications. The minimum you want is "HTL" Blu-ray. The original Blu-ray formulation starts out shiny, and the laser melts a bit of silicon and copper together, forming a less reflective material that's stable against sunlight. Avoid "LTH" discs, which use dyes similar to those used on CD-R and DVD-R.

Even better is M-Disc, a proprietary non-organic HTL disc using a material engineered specifically for longevity. The disc is even HTL for the DVD-R version, one reason it needs a special drive to burn it… regular CD-R and DVD-R is always LTH…. the laser is making the dye layer clear, allowing the reflection layer through. M-Disc is described as "glassy carbon", it's unaffected by oxygen — the enemy of the aluminum layer in a CD-R or DVD-R — and rated for 1,000 years life. Even if that's an exaggeration, if that's not effectively "forever" for me, then I'm certain all my old media will be stored in 0.001% of my positronic brain upgrade. But this also points out a universal truth — nothing is forever and under every circumstance. There's no such thing as "waterproof", you only have degrees of water resistance. There's no such thing as "timeproof", you only have degrees of time/entropy resistance.

You can also consider managed archival… online storage. Storing thingsin the cloud, you will get the slight benefit of enterprise-level hardware, perhaps. But what you're really getting is active data management. You're paying that company to mind your data, replace drives (on a RAID, of course) before they fail, etc. I never would have considered this a year ago when I was on satellite internet at 12Mb/s and low data caps, but today I'm on a gigabit connection with no data caps. So it's an actual option for some.

Rot and Redundancy

No matter what you do, you're subject to data loss. The best archival medium is redundancy. When I store my photos on those pricey BDXL M-Discs, I write overlapped… so half of the last backup goes onto the next backup. Every photo is on at least two discs. And my RAID. And my backup HDD.

While SSDs can show data rot in a few years, HDDs and tape certainly can as well. I shot analog video in 1994 on high quality tape that was showing rot by 2008 (multiple playings on different 8mm decks got it all converted to digital). Some of the old CDs we made in the early 1990s are still good, some of the CDs and DVDs I made in the 1990s are dead. Particularly for anything using organic dyes, both light and heat are your long-term enemies.

Then there's format rot. If I backed up onto TR-4 tape in the 1990s, it might not matter if the tape's good if I can't find a compatible drive. Or a PATA hard drive for that matter… you can probably still find a PATA interface, but your modern PC may not have one. One of the best reasons for CD/DVD/BD is the fact it's a consumer format, not just a computer format… that's the primary reason the discs stayed upward compatible over three decades and counting. But will that last? UHD Blu-ray is out, with 66GB and 100GB discs, based on the same tech used in BDXL, so that is

a consumer format. But it’ll never be as popular as HD Blu-ray, which isn’t as popular as DVD. Will we have another consumer optical upgrade? But I digress…

So archive often, on different media, and in different places. That’s the only way to prevent data loss for certain.

Read More

So what's up with the M-Disc? Part 1

So what's up with the M-Disc? Part 2

Wanna keep your data for 1,000 YEARS? No? Hard luck, HDS wants you to anyway