

Why the Future of Data Storage is (Still) Magnetic Tape

Disk drives are reaching their limits, but magnetic tape just gets better and better

By **Mark Lantz**



Photo: Victor Prado

It should come as no surprise that recent advances in big-data analytics and artificial intelligence have created strong incentives for enterprises to amass information about every measurable aspect of their businesses. And financial regulations now require organizations to keep records for much longer periods than they had to in the past. So companies and institutions of all stripes are holding onto more and more.

Studies show [PDF] that the amount of data being recorded is increasing at 30 to 40 percent per year. At the same time, the capacity of modern hard drives, which are used to store most of this, is increasing at less than half that rate.

Fortunately, much of this information doesn't need to be accessed instantly. And for such things,

magnetic tape is the perfect solution.

Seriously? Tape? The very idea may evoke images of reels rotating fitfully next to a bulky mainframe in an old movie like *Desk Set* or *Dr. Strangelove*. So, a quick reality check: Tape has never gone away!

Indeed, much of the world's data is still kept on tape, including data for basic science, such as particle physics and radio astronomy, human heritage and national archives, major motion pictures, banking, insurance, oil exploration, and more. There is even a cadre of people (including me, trained in materials science, engineering, or physics) whose job it is to keep improving tape storage.

Tape has been around for a long while, yes, but the technology hasn't been frozen in time. Quite the contrary. Like the hard disk and the transistor, magnetic tape has advanced enormously over the decades.

The first commercial digital-tape storage system, IBM's Model 726, could store about 1.1 megabytes on one reel of tape. Today, a modern tape cartridge can hold 15 terabytes. And a single robotic tape library can contain up to 278 petabytes of data. Storing that much data on compact discs would require more than 397 million of them, which if stacked would form a tower more than 476 kilometers high.

It's true that tape doesn't offer the fast access speeds of hard disks or semiconductor memories. Still, the medium's advantages are many. To begin with, tape storage is more energy efficient: Once all the data has been recorded, a tape cartridge simply sits quietly in a slot in a robotic library and doesn't consume any power at all. Tape is also exceedingly reliable, with error rates that are four to five orders of magnitude lower than those of hard drives. And tape is very secure, with built-in, on-the-fly encryption and additional security provided by the nature of the medium itself. After all, if a cartridge isn't mounted in a drive, the data cannot be accessed or modified. This "air gap" is particularly attractive in light of the growing rate of data theft through cyberattacks.

The offline nature of tape also provides an additional line of defense against buggy software. For example, in 2011, a flaw in a software update caused Google to accidentally delete the saved email messages in about 40,000 Gmail accounts. That loss occurred despite there being several copies of the data stored on hard drives across multiple data centers. Fortunately, the data was also recorded on tape, and Google could eventually restore all the lost data from that backup.

The 2011 Gmail incident was one of the first disclosures that a cloud-service provider was using tape for its operations. More recently, Microsoft let it be known that its Azure Archive Storage uses IBM tape storage equipment.

1951: Magnetic tape first used to record data on a computer (Univac). *Photo: Bettmann/Getty Images*



-

All these pluses notwithstanding, the main reason why companies use tape is usually simple economics. Tape storage costs one-sixth the amount you'd have to pay to keep the same amount of data on disks, which is why you find tape systems almost anywhere where massive amounts of data are being stored. But because tape has now disappeared completely from consumer-level products, most people are unaware of its existence, let alone of the tremendous advances that tape recording technology has made in recent years and will continue to make for the foreseeable future.

All this is to say that tape has been with us for decades and will be here for decades to come. How can I be so sure? Read on.

Tape has survived for as long as it has for one fundamental reason: It's cheap. And it's getting cheaper all the time. But will that always be the case?

You might expect that if the ability to cram ever more data onto magnetic disks is diminishing, so too must this be true for tape, which uses the same basic technology but is even older. The surprising reality is that for tape, this scaling up in capacity is showing no signs of slowing. Indeed, it should continue for many more years at its historical rate of about 33 percent per year, meaning that you can expect a doubling in capacity roughly every two to three years. Think of it as a Moore's Law for magnetic tape.

That's great news for anyone who has to deal with the explosion in data on a storage budget that remains flat. To understand why tape still has so much potential relative to hard drives, consider the way tape and hard drives evolved.

Both rely on the same basic physical mechanisms to store digital data. They do so in the form of narrow tracks in a thin film of magnetic material in which the magnetism switches between two states of polarity. The information is encoded as a series of bits, represented by the presence or absence of a magnetic-polarity transition at specific points along a track. Since the introduction of tape and hard drives in the 1950s, the manufacturers of both have been driven by the mantra "denser, faster, cheaper." As a result, the cost of both, in terms of dollars per gigabyte of capacity, has fallen by many orders of magnitude.

These cost reductions are the result of exponential increases in the density of information that can be recorded on each square millimeter of the magnetic substrate. That areal density is the product of the recording density along the data tracks and the density of those tracks in the perpendicular direction.

Early on, the areal densities of tapes and hard drives were similar. But the much greater market size and revenue from the sale of hard drives provided funding for a much larger R&D effort, which enabled their makers to scale up more aggressively. As a result, the current areal density of high-capacity hard drives is about 100 times that of the most recent tape drives.

Nevertheless, because they have a much larger surface area available for recording, state-of-the-art tape systems provide a native cartridge capacity of up to 15 TB—greater than the highest-capacity hard drives on the market. That's true even though both kinds of equipment take up about the same amount of space.

With the exception of capacity, the performance characteristics of tape and hard drives are, of course, very different. The long length of the tape held in a cartridge—normally hundreds of meters—results in average data-access times of 50 to 60 seconds compared with just 5 to 10 milliseconds for hard drives. But the rate at which data can be written to tape is, surprisingly enough, more than twice the rate of writing to disk.



Photo: Victor Prado

Inside and Out: A modern Linear Tape-Open (LTO) tape cartridge consists of a single reel. After the cartridge is inserted, the tape is fed automatically to a reel built into the drive mechanism.

Over the past few years, the areal density scaling of data on hard disks has slowed from its historical average of around 40 percent a year to between 10 and 15 percent. The reason has to do with some fundamental physics: To record more data in a given area, you need to allot a smaller region to each bit. That in turn reduces the signal you can get when you read it. And if you reduce the signal too much, it gets lost in the noise that arises from the granular nature of the magnetic grains coating the disk.

It's possible to reduce that background noise by making those grains smaller. But it's difficult to shrink the magnetic grains beyond a certain size without compromising their ability to maintain a magnetic state in a stable way. The smallest size that's practical to use for magnetic recording is known in this business as the superparamagnetic limit. And disk manufacturers have reached it.


Until recently, this slowdown was not obvious to consumers, because disk-drive manufacturers were able to compensate by adding more heads and platters to each unit, enabling a higher capacity in the same size package. But now both the available space and the cost of adding more heads and platters are limiting the gains that drive manufacturers can make, and the plateau is starting to become apparent.

There are a few technologies under development that could enable hard-drive scaling beyond today's superparamagnetic limit. These include heat-assisted magnetic recording (HAMR) and microwave-assisted magnetic recording (MAMR), techniques that enable the use of smaller grains and hence allow smaller regions of the disk to be magnetized. But these approaches add cost and introduce vexing engineering challenges. And even if they are successful, the scaling they provide is, according to manufacturers, likely to remain limited. Western Digital Corp., for example, which recently announced that it will probably begin shipping MAMR hard drives in 2019, expects that this technology will enable areal density scaling of only about 15 percent per year.

In contrast, tape storage equipment currently operates at areal densities that are well below the superparamagnetic limit. So tape's Moore's Law can go on for a decade or more without running into such roadblocks from fundamental physics.

Still, tape is a tricky technology. Its removable nature, the use of a thin polymer substrate rather than a rigid disk, and the simultaneous recording of up to 32 tracks in parallel create significant hurdles for designers. That's why my research team at the IBM Research–Zurich lab has been working hard to find ways to enable the continued scaling of tape, either by adapting hard-drive technologies or by inventing completely new approaches.

In 2015, we and our collaborators at FujiFilm Corp. showed that by using ultrasmall barium ferrite particles oriented perpendicular to the tape, it's possible to record data at more than 12 times the density achievable with today's commercial technology. And more recently, in collaboration with Sony Storage Media Solutions, we demonstrated the possibility of recording data at an areal density that is about 20 times the current figure for state-of-the-art tape drives. To put this in perspective, if this technology were to be commercialized, a movie studio, which now might need a dozen tape cartridges to archive all the digital components of a big-budget feature, would be able to fit all of them on a single tape.



To enable this degree of scaling, we had to make a bunch of technical advances. For one, we improved the ability of the read and write heads to follow the slender tracks on the tape, which were just 100 or so nanometers wide in our latest demo.

We also had to reduce the width of the data reader—a magnetoresistive sensor used to read back the recorded data tracks—from its current micrometer size to less than 50 nm. As a result, the signal we could pick up with such a tiny reader got very noisy. We compensated by increasing the signal-to-noise ratio inherent to the media, which is a function of the size and orientation of the magnetic particles as well as their composition and the smoothness and slickness of the tape surface. To help further, we improved the signal processing and error-correction schemes our equipment employed.

To ensure that our new prototype media can retain recorded data for decades, we changed the nature of the magnetic particles in the recording layer, making them more stable. But that change made it harder to record the data in the first place, to the extent that a normal tape transducer could not reliably write to the new media. So we used a special write head that produces magnetic fields much stronger than a conventional head could provide.



Photos: David Parker/Science Source; right: IBM

A Data Deluge: Modern tape libraries can hold hundreds of petabytes, whereas the IBM 726 [right], introduced in 1952, could store just a couple of megabytes.

Combining these technologies, we were able to read and write data in our laboratory system at a linear density of 818,000 bits per inch. (For historical reasons, tape engineers around the world measure data density in inches.) In combination with the 246,200 tracks per inch that the new technology can handle, our prototype unit achieved an areal density of 201 gigabits per square inch. Assuming that one cartridge can hold 1,140 meters of tape—a reasonable assumption, based on the reduced thickness of the new tape media we used—this areal density corresponds to a cartridge capacity of a whopping 330 TB. That means that a single tape cartridge could record as much data as a wheelbarrow full of hard drives.

In 2015, the Information Storage Industry Consortium, an organization that includes HP Enterprise, IBM, Oracle, and Quantum, along with a slew of academic research groups, released what it called the “International Magnetic Tape Storage Roadmap.” That forecast predicted that the areal density of tape storage would reach 91 Gb per square inch by 2025. Extrapolating the trend suggests that it will surpass 200 Gb per square inch by 2028.

The authors of that road map each had an interest in the future of tape storage. But you needn’t worry that they were being too optimistic. The laboratory experiments that my colleagues and I have recently carried out demonstrate that 200 Gb per square inch is perfectly possible. So the feasibility of keeping tape on the growth path it’s had for at least another decade is, to my mind, well assured.

Indeed, tape may be one of the last information technologies to follow a Moore’s Law–like scaling, maintaining that for the next decade, if not beyond. And that streak in turn will only increase the cost advantage of tape over hard drives and other storage technologies. So even though you may rarely see it outside of a black-and-white movie, magnetic tape, old as it is, will be here for years to come.

This article appears in the September 2018 print issue as “Tape Storage Mounts a Comeback.”

About the Author

Mark Lantz is manager of the Advanced Tape Technologies at IBM Research Zurich.