

# The last data center

Long-term data storage enters a new epoch



Sebastian Moss  
Editor-in-Chief



Images by Sebastian Moss

**B**uried deep in a mountain at the edge of civilization there lies what may end up being humanity's last message.

To get there we traveled over permafrost and up a steep passage, past signs warning of polar bears ahead. Then we descended into the dark mines, our headlamps illuminating falling ice crystals disturbed by our presence, bathing us in glittering and ephemeral showers.

It's not clear how long we traveled down the shaft, time travels differently underground. Here in this Svalbard mountain, it is measured in eons, not hours.

We go past the Global Seed Vault, a backup facility of the world's seeds in case of disaster, and journey further downwards. At last, we come to a door, glowing in the pitch-black. Emblazoned on it are the words "Arctic World Archive: Protecting World Memory."

Before we talk about what's behind that door, we should understand the two great data challenges it hopes to solve. It is joined by dozens of startups, researchers, and even a trillion-dollar corporation in competing to figure out the future of data.

One challenge is philosophical, that of our death and destruction, and what we leave behind. The other is more immediate, that of growing hordes of data, which threaten to overflow our current systems and leave us unable to keep critical data in an economic way.

## The beginning of recorded knowledge

The story of data is an ancient one. Some 73,000 years ago in what is now South Africa, an early human picked up a piece of ocher and scratched a symbol into a shard of stone, in what is our earliest recorded piece of human artwork.

It took the majority of our species' history to get to written recordings, with the Kish tablet in 3500–3200 BC, where humans etched pictographic inscriptions into limestone. Even then, it took thousands of years to advance to clay tablets, and still further to get to papyrus, parchment, and finally paper.

Most of what happened in the world was not recorded. Of what was, the majority

has been lost in wars, fires, and through institutional decay, never to be recovered. Our understanding of ourselves and our past is told through what little survived, providing a murky glimpse that is deeply flawed and relies on the skewed records of kings and emperors.

Now, things are different. We are flooded with data, from individuals, corporations, and machines themselves. But we keep that data primarily on hard drives and solid state drives, which last mere decades if kept unused in ideal conditions, and just a handful of years if actively run.

Other common storage platforms are magnetic tape and optical discs, which themselves come in multiple formats of varying density and lifespan, but are often used for 'cold' longer-term storage.

All have their uses and individual benefits and drawbacks, but the simple fact is that if we stopped transferring data to new equipment, nearly all of it would be gone before the century is out.

"The only written records of our time would be the embossing on stainless steel cooking pots saying 'made in China' and probably the company logos on ceramics," Memory of Mankind (MoM) founder Martin Kunze explained.

Kunze is one of a select few hoping to prevent such a tragic loss for our future. To do so, he is looking to our past.

## A collection of memories

"I studied art with a focus on silicate technologies and ceramics," he said. "The idea of using ceramics as a data carrier is not new, it's 5,000 years old."

Like the archive in Svalbard, and the Deep Sea Scrolls of the past, he has also turned to depositing data underground as a method of long-term storage.

2km deep, inside the world's oldest salt mine beneath the Plassen mountain in

Hallstatt, Austria, rows of neatly organized ceramic tiles attempt to provide a snapshot of our world.

The most immediately discernable tiles are readable to the human eye - with words and images printed onto them at 300 dpi resolution, similar to a normal color printer.

Less visually exciting at a distance, but perhaps far more important, are ceramic microfilm plates. Kunze turned to physical vapor deposition, a method of vacuum deposition that produces thin films and coatings on substrates such as metals, and then laser etches data at five lines per millimeter. This gives around 500 times the density of the original plates - which will be used to store 1,000 of the world's most important books.

This, in turn, sparked a government-funded project to build a femtosecond laser, which could write onto even thinner materials such as glass-ceramic. "It's very early days, but we have proved that it's possible to write and read 10 gigabytes per second" at much higher density, he said.

As for its lifetime, it should "far exceed the existence of our Solar System, so you could say it's eternal," he said. The technology is being developed by a company Kunze cofounded, Cerabyte, currently still in stealth mode.

Cerabyte does not expect to produce the tech on its own, and has turned to Sony - which has an optical disc factory in Austria that is slowly declining with the death of the format - to potentially develop "a minimum viable product that we aim to have in one and a half years," he said.

He's far from alone in trying to reinvent the data storage landscape. While Kunze turned to the past for inspiration, others have gone for an approach that appears ripped from science fiction.

## The data soup

"I guess you could say we're in the business

*"We're manipulating DNA for the purpose of both storage and compute and making some real progress here"*

of trying to build a bio-computer," Dave Turek, chief technology officer at Catalog, said. "We're manipulating DNA for the purpose of both storage and compute and making some real progress here."

Turek knows all about the intricacies of traditional computing - the last time we spoke was after the launch of the Summit

Catalog believes that those researchers, as well as some rival DNA storage companies, have made a crucial error. "You then have to solve a fundamental problem, which is that every time you add another base to your synthetic piece of DNA, it takes 30 seconds. So if the cost is 30 seconds for every two bits of information, that's not going to work very well."

decoding."

This is because the genome industry has prioritized accuracy at the base level over speed, at a fidelity that Catalog doesn't actually need. "We've got to resolve the issue of rapid decoding in parallel to the velocity with which we can write," he said.

"We have some partnerships established



supercomputer, then the world's most powerful, a project he spearheaded during his nearly 23 years at IBM.

Wetware computing is a new avenue. "I'm not a molecular biologist, so we're on even ground here," he joked before launching into a dense explanation of DNA.

The classic double helix construct of DNA is made of ladders that are formed of just four bases: Adenine, Guanine, Cytosine, and Thymine. With that simple starting point, all of the sequences of bases that dictate everything that makes every living thing unique are found. "It's biology's way of encoding information," Turek said.

We have DNA of about 3.5 billion bases long, which contains all that we are, hinting at the storage potential that could be harnessed. It's not a new idea: Early attempts date back to 1988, while American geneticist George Church encoded a book into DNA in 2012.

Those approaches took each of the letters A, G, C, and T, and assigned bit values to them, potentially allowing incredible data storage density. "Immediately, people started seeing DNA as the remedy to the overflow of information created in modern society," Turek said. "And they're completely wrong."

The company has decided to scale back and not work at the base pair level, instead constructing an alphabet composed of small snippets of DNA. Catalog now has 100 of these oligonucleotides, as they are known, with which it can create data by connecting them together in what is called ligation.

"And the machine that we invented and developed automates that process, which is a big deal," Turek said. "It uses inkjet printheads that contain these oligonucleotides. And each reservoir is unique from every other one. So that each of the nozzles can be instructed to fire a particular snippet of DNA out of my alphabet."

The 'Shannon' machine has three thousand nozzles that deposit DNA ink drops at the picoliter level, 500,000 times a second. These mix together to create a long string of data holding DNA. "The way it's currently configured, I can create a trillion unique molecules in a day," he said.

This can then be read with a DNA sequencer, with the company currently using Oxford Nanopore machines. These are much slower than the writing machines - "if I take 10 minutes to write a whole bunch of data, it might take me a week to do all the

to begin to try to do some real innovation on the read side as well. We have ideas today that we think could easily generate two orders of magnitude improvement in speed."

Even at the slightly lower storage ability of oligonucleotides, the amount that DNA can theoretically record is mindblowing. "You could put all the information in the world in this," Turek said, holding up a Pepsi can. Such a feat is still a way off, he admitted, but the company was able to store all of Wikipedia into a few droplets.

As for the long-term storage abilities, the death of a mammoth some 1.2 million years ago gives an insight into its longevity. Last year, researchers successfully sequenced the previously unknown Krestovka mammoth, and its body was not exactly kept in ideal conditions.

Perhaps more enticing still is another idea being cooked up at Catalog - using DNA for compute. The concept is in its early stages, and only works for very specific types of compute, but would still be a profound advancement in computing.

"We're doing one case that is inspired by a real potential user," Turek said. "This is not a theoretical abstract academic exercise that originated in a textbook. And we're using

*"We have turned to warehouse-style robotics, and we have these little robots that operate independently, they can move up and down and along the structure in this crabbing motion"*



that customer to guide us in terms of the nature of the algorithms, and the other kinds of things that need to manifest themselves in DNA."

It'll be a long time before such efforts will bear fruit, but Catalog is hopeful that the opportunity is vast. "If you want to build a parallel computer, which I did for 25 years, and you want to add an incremental unit of computing, it typically has a pretty hefty cost to it, and consumes a lot of energy and space," Turek said.

In DNA, he argues, "I can make it parallel cheaply to an extraordinarily large degree. If you say 'I want to run this instruction 100,000 times in parallel,' I would come back to you and say, '100,000? Why not 10 million? Why not 10 trillion?'"

The idea is different, to say the least. We're used to thinking about the computing world in terms of electricity and physics, in bits and bytes. "But look at how people are moving on from von Neumann architectures, and beginning to create quantum computers," Turek countered.

"We think that we're in the right place at the right time, because there is a de facto acceptance of alternative architectures in the computing world. However strange you think this might be, the guys in the quantum

world are stranger - and they can still sell their computers."

Nonetheless, Catalog still has a long way to go before it can convince companies to put their data lakes into data soups, and embrace unconventional storage solutions. Here, rival Microsoft has an advantage - it can deploy its long-term storage concept in its data centers, and rely on a robust sales network to convince users that its approach is the way to go.

### Halls of glass

In Cambridge, the hyperscaler's researchers have been experimenting with fused silica, where a femtosecond laser encodes data in glass by creating layers of three-dimensional nanoscale gratings and deformations at various depths and angles.

"To read it, we image it with a microscope," Microsoft Research's distinguished engineer and deputy lab director Dr. Ant Rowstron explained. "It's got these layers in it, and you focus on a layer. And then we take several images concurrently. We don't spin it, we read an entire sector at a time."

Project Silica was set up by Rowstron and Microsoft after realizing that conventional storage was set to hit a bottleneck. "We began from the ground

up, asking what storage should look like," he recalled, with Microsoft deciding to build upon silica data storage research from the University of Southampton.

It has gone through multiple iterations over the past few years, slowly getting denser as well as easier to read and write. "If we were to fill our 12-centimeter by 12-centimeter reference platter entirely with data, we'd be at around five terabytes," Rowstron said.



Writing is still slow and expensive, requiring high-powered lasers to accurately etch the glass in just the right part. Critics, including some working on the other projects in this piece, worry that there are fundamental limits to how fast you can pump energy into the glass without causing issues.

"It's been a lot of work," Rowstron admitted. "When we first saw the technology at Southampton, it was taking hundreds of pulses to write data into glass. But we've been working on how to form these structures with a very, very small number of pulses."

He declined to disclose exactly how far the project had come, but added "if you compare it to the state of the art, we were significantly below [that number of pulses]. You can think of it as dollars per megabyte writing, and I would say the technology is now in a good spot."

Microsoft is also thinking about how it would be read in a data center. "We have turned to warehouse-style robotics," Rowstron said. "And we have these little robots that operate independently, they can move up and down and along the structure in this crabbing motion, which is pretty cool."

On one end there would be a machine



Credit: Microsoft

writing the data onto the blocks of glass, on the other a reader ready to uncover what is within each block. Other than the robots and the writer, it wouldn't need any power, and it won't need any cooling.

That means that the first data centers Silica will inhabit will be massively over-engineered. "I guess one day we'll end up with buildings or data centers dedicated to just storing that preserved glass," Rowstron said.

That glass can last around 10,000 years. "No one's asked us to go further," Rowstron said. "There are things you could trade-off - you could trade density for lifetime and things like that. But you've got to remember our goal is to get a technology that will allow us to use this in our data centers. No data center is going to exist for 10,000 years."

Others could go further, should there be demand. The original Southampton work found that its much-slower-to-write silica could last 13.8 billion years at temperatures of up to 190°C (375°F). The researchers there

stored the Universal Declaration of Human Rights, Newton's Opticks, the Magna Carta, and the King James Bible on small discs.

"My hope is that in 200 years' time there will be a new storage technology that is even more efficient, even denser, even longer life, and people are going to say 'we don't need to use glass anymore,'" Rowstron said. "But they'll move formats because they want to, not because they have to."

Rowstron believes Silica will prove useful for both major challenges of data storage. "You want to make sure that whatever else happens, data from the world is not lost," he said. But the tech is naturally focused on the more immediately pressing concerns faced by businesses.

"Hard drives are languishing, we've had so little capacity increases in the last five years; tape is suffering," he said. "There is more and more data being produced, and trying to store that sustainably is a challenge for humanity."

Cold storage that doesn't require constant power would be a huge boost for the environment, as would moving away from rare earth metals such as those found in HDDs (of which there may not be enough to meet future demand). It is also much

## The guardians of knowledge

In popular culture, the Library of Alexandria - one of the greatest repositories of knowledge in antiquity - was lost to the flames, its endless rows of papyrus consumed by fire. Scientific theories, fantastical tales, and bureaucratic minutia became embers in a matter of hours.

But this is likely not true, Richard Ovenden, the head of the Bodleian Library, said. Its death was slow, as papyrus became less relevant with new technology - parchment - and no transition was made, like with the Library of Pergamum.

More profound was an institutional decay, as a lack of management, investment, and care led to its tragic demise.

"I think we are in a moment in time where that risk is high again," Ovenden, author of *Burning the Books: A History of Knowledge Under Attack*, said. "You certainly see that with the failure of public libraries in many Western countries."

Our story on how we store data for years to come has focused on the technology, but it's important not to forget the humans working to protect and share knowledge.

Librarians and archivists have passed the torch of knowledge from one generation to the next, ensuring the light of humanity does not go out.

"Throughout history, there have been moments of destruction of knowledge, but they've often been

matched with moments of renewal or of rescue," Ovenden said.

"That certainly is the case today, both in terms of communities rallying around their libraries to protect them, and in library's ability to adapt to the various challenges posed by the big tech companies today."

We live in a strange time where we have more access to information than ever, but it is through the lens of corporations who put profit over fighting misinformation or loss. "Preservation is not in Google's mission," Ovenden said.

"We need places on the Internet where information can be relied on, and I think many of those are libraries and archives that are not profit-making companies."

Preserving the creations of corporations, with their platforms and proprietary information, is also a difficulty.

"As those preserved worlds become much more complex and exist much more in real-time, that task of preservation becomes much more complex and philosophically challenging, as well as financially challenging for organizations," he said.

"There should be a link with regulation and taxation on the profits of the major technology companies in order to allow third-party preservation to take place by not-for-profit preservation organizations."



cheaper over the longer term, when all you have to worry about is where to put the glass.

The technology caught the eye of Guy Holmes, CEO of Tape Ark, an Australian company focused on moving aging tape libraries to the cloud. "I've actually got it just sitting here in my office," he said, grabbing what at first looked like a clear square of glass, but revealed rows of microscopic etchings when a light was shined into it.

"It's very early days, but it seems to have legs internally there, and it appears to be picking up its density," he said. "I spent a lot of time with the guys at the program just talking through density, access times, access frequency, etc. Customers tend to do one restore per week at the moment based on backups. As the tapes age, the number of access requests reduce significantly.

"We're finding there are these archives of infinite retention or tapes that have never been accessed since they were created."

His company remains excited about the technology, but is currently focused on the more immediate challenge of simply getting data off of antiquated tape systems and onto the cloud, including Microsoft Azure. There, it might even go back on tape, but first Holmes recommends trying to see if there is hidden value in the data.

"The number of people that don't even know what's on their tapes is pretty startling," he said. "On there could be a cure for a disease, but nobody knows. So we get pretty excited by some of the projects we do."

Other efforts have immediate clear value. "We got to work with Steven Spielberg on six petabytes of Holocaust survivor videos, which needed long-term preservation, it needed immutability so that nobody could go and do deep fakes on a holocaust video and say that it didn't occur," he said. The work required moving critical data off of both tape and film, in an effort to give it a longer life.

With film, we come back to the door in Svalbard. The company behind the project has its roots in the cinema industry, selling film projectors and digital light modulation technologies to Hollywood and Bollywood.

"And then that movie *Avatar* came out, and there was a super fast transition from a film-based world to digital projector-based world," Piql CEO Rune Bjerkestrand said.

### The return of film

While film appeared to be on its way out, the company was convinced that there was still something to the medium. Piql studied different types of film from around the world,

hoping to find one that would last a long time, and store a lot of data. Some last up to 500 years, but don't have much density, and are dangerously flammable.

"So we embarked on developing our own film, Piql Film, which is a nano silver halide film on a polyester base," Bjerkestrand said.

The company takes binary code and converts it into grey pixels, for a total capacity of 120GB per roll of film. Under good conditions, that film can last around 1,000 years, Piql claims.

The company has split its data archiving in two. One is more traditional, where it works with companies like Yotta in India to convert data to film and store them in commercial data centers around the world - with Piql highlighting its use against ransomware or data center disasters.

Then there's the Arctic World Archive, pitched as a more humanitarian mission to preserve crucial data for future generations.

Due to its distance from other land masses, the fact that it has been declared demilitarized by 42 nations, and its lack of valuable resources, the hope is that Svalbard is unlikely to be nuked in any future conflicts. However, just like everywhere else on the planet, the Norwegian archipelago cannot claim total safety - it is one of the fastest-warming places in the world due to climate change, and the territory has faced increased aggression from Russia.

Still, the land is remote and its empty mines are of little use to would-be invaders. At the depositing ceremony we attended, works of art from national archives around the world were placed in storage (and images of an Indian couple's wedding, who paid to have "their love recorded for eternity"). Large reels of film are carefully vacuum sealed and then placed in the shipping container 300m underground, nestled deep in the permafrost.



Anyone can pay to store their data at the archive, but it is primarily finding business with governments and public bodies. Microsoft's software collaboration platform GitHub also stored 6,000 software repositories on the site (a backup is also stored at the Bodleian library, see box out).

"It's for people that want to bring their valuable, irreplaceable information into the future to the next generation," Bjerkestrand said, but added that the company was not planning to be a curator in and of itself. "That has been a serious discussion, and we came to the conclusion, that no, why should we have an opinion on what's worth bringing into the future?"

When pressed on how the AWA will last 1,000 years or more when tied to a corporation, Bjerkestrand said that the company planned to spin it off as a non-profit foundation.

Currently, however, people pay for storage for up to 100 years, after which the data is sent back to them if they don't keep paying. In an eventuality where the institution that first contracted Piql collapsed, that could cause problems: "They need to make sure that somebody gets the right to the reel," he said.

A foundation, on the other hand, would have a longer-term focus. "Such a foundation would have interest across the world for organizations to support, it's basically supporting world memory to survive into the future," Bjerkestrand said. "So I think it's a good cause that you could get sponsors, donors, and supporters for."

When thinking about scale in that way, you cannot assume that future generations will be able to understand how data storage formats work. Here, Piql has an advantage: As it's film, you can simply hold it up to the light, and the first few frames are pictures of how to access the information.

That data then needs to be able to be readable with what is available on the roll. "We have a fundamental principle that it should be self-contained," Bjerkestrand said. "We don't compromise that it should be open source license and free to retrieve the data with the tools that are on the data. We had scenarios where we could do more data on the film, but then it would be too complex to read."

Simply recording reality and putting it on one of the above media doesn't work if future civilizations can't understand it, posing a critical challenge for much of the data of our day.

### The death of tools

"My biggest concern is the loss of knowledge

of the software that's needed to quickly interpret digitized content," Vint Cerf, TCP/IP co-creator, said.

"An increasing amount of digital content that we create was made by software, which is needed to correctly understand, render, and interact with data, like with spreadsheets and video games where you actually need a piece of software, plus a bunch of data in order to exercise it. If you don't have that software running anymore on the platforms that are available 100 years from now, then you won't be able to do that."

Cerf, known as one of the fathers of the Internet, told *DCD* that he was "worried about the kinds of software that's needed to interact with databases, for instance, timesheets and other kinds of complex objects, where we may not have as widespread implementations available, some of them may even be proprietary."

Fellow Internet Hall of Fame member Brewster Kahle shares the concern about specific pieces of software and data being proprietary or in the hands of corporations.

"If you look at the history of libraries, they are destroyed or they're strangled such that they're left irrelevant," he said. "And that used to be by king and churches, but these days it is governments and corporations."

As one of the developers of the World Wide Web precursor the WAIS system, Kahle found himself at an inflection point

of humanity - where data was set to be shared and accessible by the world, but was transient and easily lost. This, he hoped, gave an opportunity for a new kind of library.

"As we're coming to a change in media type, can we go and start a library, right away?" he wondered, launching the non-profit Internet Archive, which tries to create a long-term copy of websites, music, movies, books, and more.

"The goal of the Internet Archive is to try to build the Library of Alexandria for the digital age," he said. And, so far, the effort has been wildly successful: "We're a small organization, we're \$20-25 million a year in operational costs, and yet we're the 300th most popular website in the world."

But its continued success is threatened by a changing world.

"I could only start the Internet Archive after we had gotten the Internet and the World Wide Web to really work, both are open systems. So if we go into a period where the idea of public education or universal access to all knowledge starts to eclipse, we're in trouble," Kahle said.

"We see that now within corporate environments, and we're starting to see it in government environments, whether it's banning books or where you can pressure organizations to do things without going through the rule of law, but you go through the rule of contract."

## Layers of control

As the web becomes more centralized and in the hands of a few cloud providers, their terms of service risk controlling which information lasts long enough to be stored for greater timescales. At the same time, as systems get built on top of cloud providers, or platforms like Facebook, storing them would require also recording all that backend software that is not shared.

Kahle is less worried about the longevity of data storage devices, and more about ensuring immediate access in an open world.

"100 years ago, microfilm was a new technology, and it was greeted with this fanfare that we'd be able to make it available so that people in rural areas could have access to information just like the people in big universities," he said. "Well, that didn't really come about, it ended up just being used to just reinforce the power structures and the publishers that existed at the time."

The question of world memory is less about the medium it is stored on, and more about how it is used when it is on that medium, he argued. "Microfilm may last 500 years... if you don't throw it away. But it turns out, people will just throw it away even if it hasn't been copied forward.

"We need not just formats that will last a long time, we need to keep the material in conversation, in use so that people will continue to love it and keep it going." ■

## Our obituary

We don't know when humanity will end. It could be soon, with climate change, nuclear war, or another pandemic. Perhaps we will survive, eventually spreading to distant stars.

When we do, we may come across a golden disc floating through space, carrying a message from our past. Or, possibly, someone else might come across it.

"There were two audiences we designed for - one was the extraterrestrial audience, and that was the one that I was most concerned with, and the other was the message to ourselves," Jon Lomborg, NASA's design director for the Golden Record, recalled.

"In a sense, it was a message to ourselves, as well as a message to extraterrestrials saying this is a snapshot that one group of people thought would be a good capture of the Earth at this point in time."

While destined for the stars aboard the Voyager spacecraft - which are the furthest man-made objects from our planet - the team had to contend with very terrestrial challenges.

First, there was copyright, with some songs not allowed on the golden vinyl.

Then there were issues with nudity, after the earlier Pioneer plaque which included a line drawing of a naked man and woman drew criticism from the more prudish for showing a penis, and those on the other end of the spectrum for the censorship of female genitalia.

Lomborg's anatomically correct drawings of a man and woman were not included on the Golden Record.

Another image that was not included was that of an atomic bomb or mushroom cloud, despite the pervading threat of nuclear holocaust when the vessels launched in 1977. "Carl [Sagan] didn't want it to seem like a threat," Lomborg said. "He didn't want anything on record that could seem like 'if you mess with us, this is what we could do.' He wanted to greet the cosmos with open arms."

The discs contain music from around the world (which is available online for those curious), as well as different languages, all of which will last an incomprehensibly long time. The side facing outwards is covered by a box that "lasts about a billion years, and then it's about another billion years to erode the outward facing side," Lomborg said. "The inward-facing side lasts a lot longer - they think up into the trillions of years."

Should it be found by another species billions upon billions of years from now, it will be our last message to the cosmos. And, by dint of data storage limitations and intentional curation, it will be an imperfect one.

But Lomborg is happy with that: "In a sense, this is our obituary; let's be remembered for Mozart, not Hitler. We'll likely be done in by our own flaws, our own shortcomings may well destroy us, and that's punishment enough. But that wasn't the whole story of us."