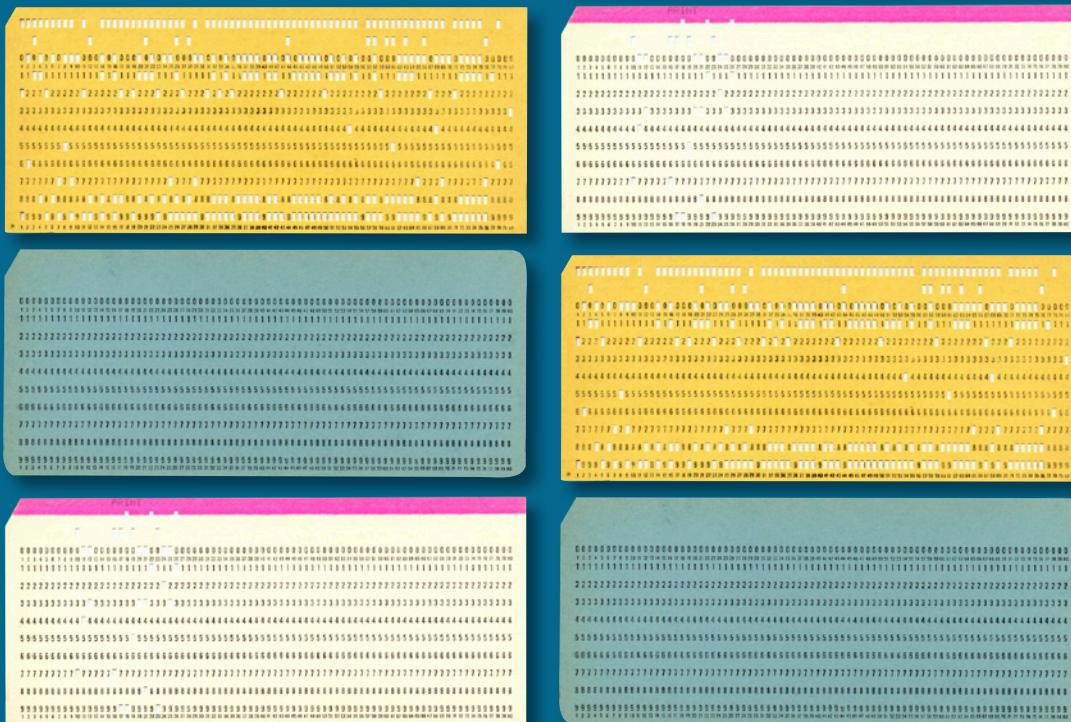


CELEBRATING 50 YEARS OF BRITISH TECHNOLOGY INNOVATION

FROM PUNCHED CARDS TO FLASH AND THE CLOUD

Since the launch of Computer Weekly in 1966, we have moved from a world of punched cards and paper tape to one where flash and the cloud have revolutionised data storage. Anthony Adshead reports



CLAUDIO DIVIZIA/ISTOCK

HOME

CW@50

1966-2016

The relationship of [storage](#) to the architecture of computing is all about capacity, [latency](#) and throughput. In other words, how much data can be kept, how quickly it can be accessed and at what rate.

Since the [launch of Computer Weekly in 1966](#), the world of storage has gone through transformations as remarkable as aviation's progression from the Wright brothers to supersonic flight.

And just as the pioneers of flight would recognise the fundamentals of today's aircraft in basic design, from the viewpoint of 1966, the speeds and magnitudes of storage today would seem utterly alien.

In the 1960s, the key methods of data storage centred on two media: Paper and cardboard; and magnetic media.

Magnetic tape and even the spinning hard drive had already been invented for data, but punched cards and paper tape were used to run programs and store data in most of the nation's datacentres.

PUNCHED CARDS

Dating back to textile and fairground organ applications from the 19th century and beyond, [punched cards](#) were usually the IBM-derived standard 73/8in x 31/4in with 80 columns and 12 rows (0-9 and 11 and 12), although there were variants of card size and column width from other computer makers, such as the UK's [ICL](#).

[Home](#)[The early days](#)[British IT innovation](#)[Enterprise IT](#)[Software](#)[Storage](#)[Datacentres](#)[Internet evolution](#)[Security](#)[Mobile communications](#)[Outsourcing](#)[Government IT](#)[Personal computing](#)[Processors](#)[Women in IT](#)[The CIO role](#)[Retail IT](#)[The next 50 years](#)

Data was represented by punched holes in each column that were read by shining a light on the card. Initially, combinations of punched holes had represented analogue forms of information, but as the 20th century progressed, they came to represent binary data. Information about the dataset – metadata – was represented in rows 11 and 12, and sometimes in unused columns.

One IBM card held 72 x 10 bits. State-of-the-art punched-card hardware in 1966 was the IBM 2540 (a peripheral to the System/360 [mainframe](#)), which could read 1,000 cards per minute (giving a throughput of 720Kb per minute; write speed/cards punched was 300 per minute) and an input hopper that held 3,100 cards. That is just over 2MB of capacity, but that was theoretically infinitely scalable as long as there was a human available to unload and reload the hoppers.

PAPER TAPE

Meanwhile, [paper tape](#) was also popular. It came in a width of 1in with room for five holes across initially, and up to eight possible later – 5-bit or 8-bit characters. IBM paper tape machines available in 1966, such as the 1621 and 1624, could read tape at a rate of 150 bits per second (9Kb per minute) and write at 15 bits per second.

But talk of latency and throughput is rather misleading when discussing storage in that era. Although [storage hardware](#) may have theoretically offered the levels of performance outlined above, a key limiting factor was simply having access to computing time.

› *Computer Weekly keeps its finger on the pulse in the world of data storage, with regular news, features and practical articles. Catch up [here](#).*

In other words, [batch](#) processing ruled. Jobs were booked into time slots in the [datacentre](#) with a notional departmental spend on computing accrued and were punched to card or tape by specialist data inputters (usually women).

While paper tape and punched cards were commonplace throughout the 1960s, magnetic media were also present. Magnetic [tape](#), such as cards

and paper tape, offered sequential access to data but much greater read/write speeds and better durability. Having said that, there were adverts for paper tape in early [1970s copies of Computer Weekly](#) that felt confident enough to trumpet a claimed superiority over magnetic tape on grounds of capacity, throughput and cost.

Magnetic tape had originated as a means of recording music in the 1920s, but by the 1950s it was available as a means of data storage. It was initially standardised as ½in tape on open 10.5in reels that were present until the 1980s. Tape cartridges emerged in the 1970s and superseded reels over the next decade.

Mid-1960s IBM tapes were nine-track (eight data and one parity) and came in a maximum length of 2,400ft. IBM's 2400 series tape drives could handle reads and writes in 1968 of up to 320Kbps with capacities in the several tens of megabytes.

MAGNETIC DISK DRIVES

At that time, a few UK datacentres would have had access to magnetic disk drives. These had originated with the IBM [Ramac 305](#) in 1957, but by the 1960s had evolved to office desk-sized disk drive units with removable disk packs.

[Home](#)[The early days](#)[British IT innovation](#)[Enterprise IT](#)[Software](#)[Storage](#)[Datacentres](#)[Internet evolution](#)[Security](#)[Mobile communications](#)[Outsourcing](#)[Government IT](#)[Personal computing](#)[Processors](#)[Women in IT](#)[The CIO role](#)[Retail IT](#)[The next 50 years](#)

These disk packs came in diameters of over 1ft and had capacities of a few megabytes each. They were essentially similar to the standardised [hard disk drives](#) (HDDs) of today, comprising multiple platters and read/write heads.

Mid- to late-1960s IBM disk drives, such as the 2300 series, had access times in the tens of milliseconds and capacities of hundreds of megabytes across multiple disk packs.

The step-change with the disk drive was that suddenly access to data was transformed from sequential to random. Sure, read/write heads had to get to the data required, but it was certainly a lot faster than the need to sort through an entire card deck or paper or magnetic tape to find a particular piece of information.

By the end of the 1970s, the paper/card-board-based media had had their day. This was the decade when the disk drive rose to dominance alongside its magnetic cousin, the tape, now increasingly in cartridge format.

The era of having to physically walk to the datacentre was over, with [workstation](#) terminals increasingly prevalent and mainframes running from spinning disk and tape.

And as the 1980s progressed, a new age dawned. The roots of the server-based architecture had existed as a concept since

the 1960s in “remote job entry” on IBM OS/360 systems, where jobs could be input and output using a remote zterminal.

CLIENT-SERVER MODEL

But it was in the 1980s that the [client-server model](#) was enabled by developments in [networking](#) and [operating systems](#) that allowed file shares between multiple users.

With the era of the network and server, another key distinction in the world of storage began to arise. Until the era of mainframe and [minicomputer](#), all storage was essentially [direct-attached](#) – that is, the storage was dedicated to one computing device only.

Now, with the client-server architecture, the development of file systems that allowed data to be shared (NFS, Novell’s Netware) and the increasing availability of small form factor hard disk drives, the possibility of shared storage emerged.

The 1980s was the heyday of the disk drive, a decade when hundreds of disk drive makers emerged. There were more than 200 – there are now three – and it was a period in which clients/workstations became “fat”, in that they had their own on-board HDDs.

Form factors settled on the 3.5in drive and alongside it [Raid](#) was developed to provide

Mid- to late-1960s IBM disk drives, such as the 2300 series (2311 memory unit with six platters shown here), had access times in the tens of milliseconds and capacities of hundreds of megabytes across multiple disk packs



I, DEEP SILENCE/WIKIMEDIA

[Home](#)[The early days](#)[British IT innovation](#)[Enterprise IT](#)[Software](#)[Storage](#)[Datacentres](#)[Internet evolution](#)[Security](#)[Mobile communications](#)[Outsourcing](#)[Government IT](#)[Personal computing](#)[Processors](#)[Women in IT](#)[The CIO role](#)[Retail IT](#)[The next 50 years](#)

data protection and performance gains when multiple HDDs were clustered together.

[Network-attached storage](#) (NAS) built on those innovations, with NetApp's predecessor, [Auspex Systems](#), providing the first NAS filer in the early 1990s and ushering in the era of shared storage. NAS is essentially a storage server, and it revolutionised access to data, allowing multiple users, running apps from multiple servers, to share capacity across the local area network (LAN).

NAS FILERS

Capacities for the first NAS filers, such as those from NetApp, were in the gigabytes. For example, the NetApp FASServer 400 offered 14GB of capacity in 1993. By the late 1990s, the company's products ranged up to hundreds of gigabytes and broke the terabyte barrier in 1998.

Access to NAS filers was file access, whereby files were organised on a discrete [file system](#) on the filer, from where remote users requested them. This is adequate for many types of data, but file-based access and the Ethernet LAN lacked the performance required for multiple client access to databases and online transaction processing.

And so, the [storage area network](#) (SAN) was born. Instead of access to entire files, clients that shared access to, for example, the same database could work on that file simultaneously – within limits – and read and write data within it.

TALK OF LATENCY AND THROUGHPUT IS RATHER MISLEADING WHEN DISCUSSING STORAGE IN THE 1960s

That was because the file system and its access controls were shifted from the storage to the server. The storage merely held the component blocks, hence block-access storage, and the key defining point of SAN.

Among the first SAN products was Sun's [SPARCStorage Array Model 100](#) of 1994, which brought together multiple HDDs under various configurable Raid levels – 1 (mirroring), 0+1 (mirrored and striped) and 5 (striped with parity). The product claimed four nines availability (99.99%) and base capacity of 6GB expandable to 31GB with 2,000 [input/output operations per second](#) (IOPS).

With the then-new Fibre Channel networking, it offered the possibility of placing a second unit up to 2km away and bandwidth of up to 100Mbps. It claimed to offer all of this for \$1.62 per megabyte.

And so the die was cast for the best part of the next couple of decades. With NAS and SAN, the world had gained shared storage for the new standard client-server architecture. At its heart was the spinning disk hard drive.

INCREASES IN HDD CAPACITY

Key changes from the early/mid-1990s to the dawn of the current decade centred on increases in HDD capacity and access time (largely via increases in operating RPM, but also via connectivity and storage controller/central processing unit performance), as well as the addition of advanced storage services,

[Home](#)[The early days](#)[British IT innovation](#)[Enterprise IT](#)[Software](#)[Storage](#)[Datacentres](#)[Internet evolution](#)[Security](#)[Mobile communications](#)[Outsourcing](#)[Government IT](#)[Personal computing](#)[Processors](#)[Women in IT](#)[The CIO role](#)[Retail IT](#)[The next 50 years](#)

such as [data protection](#) and storage tiering. The world of SAN and NAS looks like a relatively long, stable period in retrospect, but it was by no means the dominant form of storing data in the enterprise. Most organisations relied on hard disk storage direct-attached to servers.

The rise of server virtualisation changed all that. One of server virtualisation's key effects on storage has been to highlight any bottleneck in access to data. Where there was one application per physical server and consequent levels of storage traffic, there are now many virtual machines per box.

With this, the volume and randomness of input/output (I/O) to storage resources increased by orders of magnitude, in the so-called [I/O blender effect](#). With storage direct-attached in relatively few HDDs, performance suffered. So, as virtualisation swept the world's datacentres, we saw the consolidation of shared storage as the norm. Upgrading to shared storage and revamping [backup](#) processes is still a key item on many organisations' agenda.

FLASH STORAGE

But at the same time, and to some extent driven by virtualisation but also by the needs of very high-speed transactional processing, there is the rise of [solid-state flash storage](#).

[Flash storage](#) offers incredibly low latency, measured in milliseconds or even microseconds. Initially offered as an add-in alongside existing spinning disk array products, the flash market has spawned entire product categories of [all-flash](#) and [hybrid flash](#) arrays. And while those types of storage product mirror somewhat the existing architecture of HDD-based arrays, new architectures threaten the long dominance of NAS and SAN.

Here we are talking about [hyper-converged](#) or so-called [hyper-scale](#) architectures. This sees compute, storage and networking capacity all in one unit. It has been driven by internet giants such as Google and Facebook, which require vast amounts of compute and storage to service webscale operations.

With the in-house expertise to do it and a keenness to avoid the cost and lock-in of enterprise storage products, they built grids of what are, essentially, servers. Unlike the norms of enterprise storage, where if, say, a HDD fails, you replace it, with webscale architectures a failed compute/storage/network node is replaced in its entirety.

OPEN SOURCE HARDWARE

While the internet big guns have got together to promote so-called open source hardware, some storage startups have begun to ape the developments of hyper-converged infrastructure, so

THE 1980s WAS THE HEYDAY OF THE DISK DRIVE, WHEN HUNDREDS OF DISK DRIVE MAKERS EMERGED

[Home](#)[The early days](#)[British IT innovation](#)[Enterprise IT](#)[Software](#)[Storage](#)[Datacentres](#)[Internet evolution](#)[Security](#)[Mobile communications](#)[Outsourcing](#)[Government IT](#)[Personal computing](#)[Processors](#)[Women in IT](#)[The CIO role](#)[Retail IT](#)[The next 50 years](#)

we have seen hardware products that combine compute, storage and networking from the likes of Nutanix, Simplivity and VMware with its EVO:Rail.

Increasingly, also we are seeing these products offered as software that can be run on dedicated commodity server hardware or as virtual machines within a virtualisation hypervisor.

This development – so-called [software-defined storage](#) – threatens to shatter the lock the big storage hardware makers have had on the market. It could break the link between storage hardware – essentially commodity products – from the intelligence of storage in the software.

There is no reason, in theory, why enterprise storage cannot be built from commodity hardware and storage software. At the same time, the enterprise storage market as it has existed for 20 years is also threatened by the cloud.

CLOUD FOR STORAGE

It is true that, at present, the cloud is hampered by concerns over availability and security, and it certainly lacks the rapid access times required for the hottest data, but those obstacles are likely to be overcome in time.

In fact, significant numbers of organisations already use the cloud for some form of data storage, notably for backups and

archives that do not require rapid access times, and we already see the use of the cloud in hybrid form as a target for backups, as a nearline tier to in-house storage, for example.

But eventually, concerns over access and security are likely to dissolve. When that happens, the possibility is that [cloud storage](#), with a network of providers able to buy or build storage hardware at economies of scale, will become the norm and organisations need only have minimal local capacity, or even none at all.

SIGNIFICANT NUMBERS OF ORGANISATIONS ALREADY USE THE CLOUD FOR SOME FORM OF DATA STORAGE

MAJOR TRANSFORMATION

So, that's 50 years of data storage. In that time, we have travelled from paper and card as the key medium and a world where you would never see a screen. Instead, you would queue up at the counter and book a computing slot, with the results coming back on reams of paper the

next day – or maybe in an hour or two, if you were lucky.

Now we live in a world where vast amounts of business-critical transactions can be read and written every second. Magnetic tape is still with us, of course, and with its exponential increases in capacity and throughput, it is likely to remain for cost-effective long-term storage for a few years. But apart from that, all that remains the same are the bits and bytes of digital information.

If the next 50 years see transformations of the same magnitude, what will data storage look like then? ■