
The Economist
Let it rise
A special report on corporate IT
October 25th 2008

Click the window to play video
Click here to go to the AMD site

Let it rise

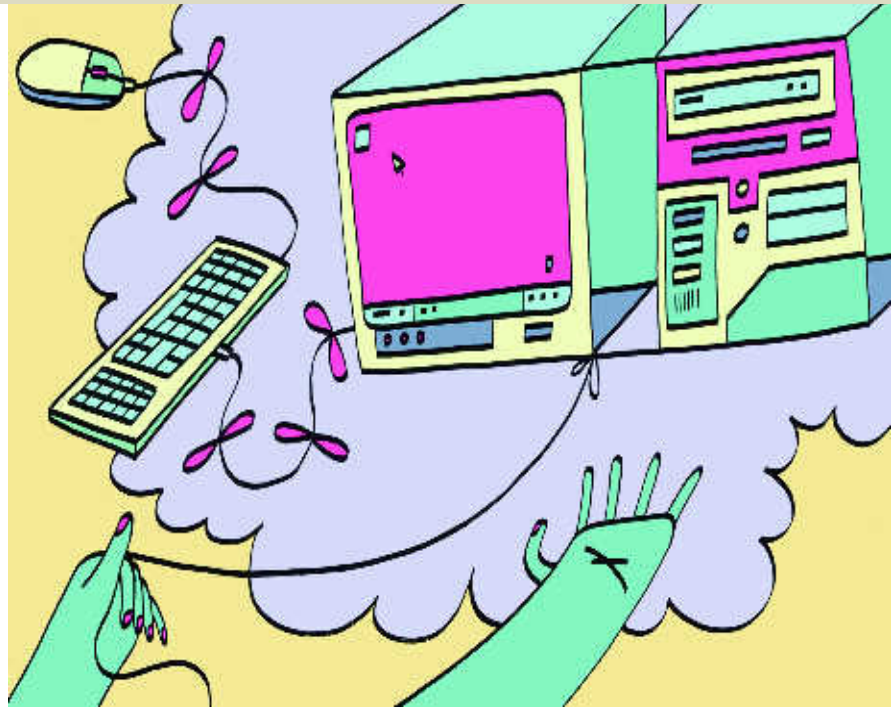

Information technology is turning into a global "cloud" accessible from anywhere, says Ludwig Siegele. What does that mean for the way people conduct business?

IN THE beginning computers were human. Then they took the shape of metal boxes, filling entire rooms before becoming ever smaller and more widespread. Now they are evaporating altogether and becoming accessible from anywhere.

That is about as brief a history of computers as anyone can make it. The point is that they are much more than devices in a box or in a data centre. Computing has constantly changed shape and location—mainly as a result of new technology, but often also because of shifts in demand.

The first "computers" were indeed people. The word originally meant an individual who solved equations, often using a mechanical calculator. Hundreds of them were employed by big companies that needed to do a lot of number-crunching, such as aeroplane manufacturers. It was only around 1945 that the word came to describe machinery.

But even after that, computing kept undergoing mutations—or, in the jargon, platform shifts. The mainframe, the original computing platform, was dethroned by minicomputers, which in turn gave way to personal computers, which are now being pushed aside by hand-held devices and smartphones. With each step the architecture—the underlying structure of computing—became more distributed.

Now, this special report will argue, computing is taking on yet another new shape. It is becoming more centralised again as some of the activity moves into data centres. But more importantly, it is turning into what has come to be called a "cloud", or collections of clouds. Computing power will become more and more disembodied and will be consumed where and when it is needed.

The rise of the cloud is more than just another platform shift that gets geeks excited. It will undoubtedly transform the information technology (IT) industry, but it will also profoundly change the way people work and companies operate. It will allow digital technology to penetrate every nook and cranny of the economy and of society, creating some tricky political problems along the way.

Promise of heaven

Here we go again, you may think. In order to generate new demand, the maturing IT industry keeps creating new buzzwords, often with celestial connotations ("cyberspace", "blogosphere"), which suggest some kind of technological nirvana. The reality is much more down to earth.

Hype is indeed rampant in "cloud com- ▸▸

Also in this section

Acknowledgments

The author would like to thank many people besides those cited in the text for their time, ideas and kindness during the preparation of this report. They include Marc Andreessen of Ning, Brian Behlendorf of CollabNet, Chris Capossella and Tim O'Brien of Microsoft, Simon Crosby of Citrix, Rob Enderle of Enderle Group, Robert Hodges of Continuent, Dirk Holzwarth of Alstom, Greg Papadopoulos of Sun Microsystems, Dennis Quan of IBM, Simeon Simeonov of Polaris Venture Partners, Rahul Sood of SAP and Kishore Swaminathan of Accenture.

A list of sources is at

www.economist.com/specialreports

An audio interview with the author is at

www.economist.com/audiovideo

Cumulonimbus 1

Cloud computing, average weekly search volumes
September average 2007=100

Source: Google Trends

puting". The term entered into IT-speak only a year ago and has spread voraciously. Cloud conferences and cloud blogs are multiplying almost as quickly as cloud start-ups. Established IT firms are slapping the new label on old gear.

In fact, the cloud craze may have peaked already, if the number of Google searches is any guide (see chart 1). Cloud computing is bound to go through a "trough of disillusionment", as Gartner, a research firm, calls the phase in the hype cycle when technologies fail to meet expectations and quickly cease to be fashionable. Much still needs to be invented for the computing sky to become truly cloudy.

Yet even if the term is already passé, the cloud itself is here to stay and to grow. It follows naturally from the combination of ever cheaper and more powerful processors with ever faster and more ubiquitous networks. As a result, data centres are becoming factories for computing services on an industrial scale; software is increasingly being delivered as an online service; and wireless networks connect more and more devices to such offerings.

All this allows computing to be disaggregated into components—or "services", in IT parlance. This is why European technologists such as Lutz Heuser, head of research at SAP, a German software giant, like to refer to it as the "internet of services". The cloud metaphor seems more apt. The internet is used mainly by people with personal computers and a physical network connection. Cloud applications, on the other hand, will be used by billions of devices of all kinds, many of them untethered, but will be connected to the "internet of things".

In some ways the cloud is already hanging in the sky, especially for consumers. According to a recent study, 69% of Americans connected to the web use some kind of "cloud service", including web-based e-mail or online data storage (see chart 2). The best example is Google, the biggest online search company by far, which now offers a plethora of web-based applications such as word-processing or online spreadsheets.

Learning to float

Companies, too, have been moving into the cloud, albeit much more cautiously. Financial institutions in particular have for some time been building "computing grids". Firms that provide enterprise software as a service (saas) over the internet, such as Salesforce.com and NetSuite, have been growing steadily.

In the years to come companies are likely to venture much farther. For one, operators of computing clouds such as Amazon and Google have shown that this is a far more efficient way of running IT systems. Secondly, many firms will find they have no choice. The way in which their IT infrastructure has grown is proving unsustainable. Most corporate data centres today are complex warrens of underused hardware that require more and more people, space and power to keep them going. The current economic malaise will increase the pressure on companies to become more efficient. More has to be done with less, which is cloud computing's main promise.

This special report will chronicle the rise of the cloud and try to predict where it is heading. It will start by looking at the technology. Computing clouds are immensely complex, but can be roughly divided into three layers: infrastructure, applications and the periphery where they meet the real world. These will be discussed in turn. The report will go on to consider the impact the cloud will have on the IT industry and the economy as a whole. The conclusion will look at what might stop the cloud from growing ever thicker: regulation and worries about the safety of both personal and corporate data.

A young cloud 2

Cloud-computing activities, by age cohorts
April-May 2008, % of users*

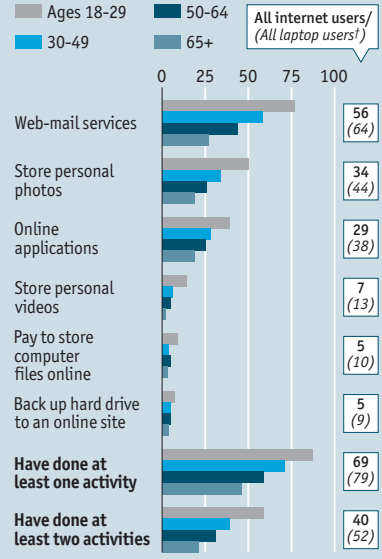

Source: Pew Internet & American Life Project

*Survey based on 1,553 internet users
†35% of internet users, connecting to Wi-Fi away from home or work

Irving Wladawsky-Berger, a technology visionary at IBM, compares cloud computing to the Cambrian explosion some 500m years ago when the rate of evolution speeded up, in part because the cell had been perfected and standardised, allowing evolution to build more complex organisms. Similarly, argues Mr Wladawsky-Berger, the IT industry spent much of its first few decades developing the basic components of computing. Now that these are essentially standardised, bigger and more diverse systems can emerge. "For computing to reach a higher level", he says, "its cells had to be commoditised." ■

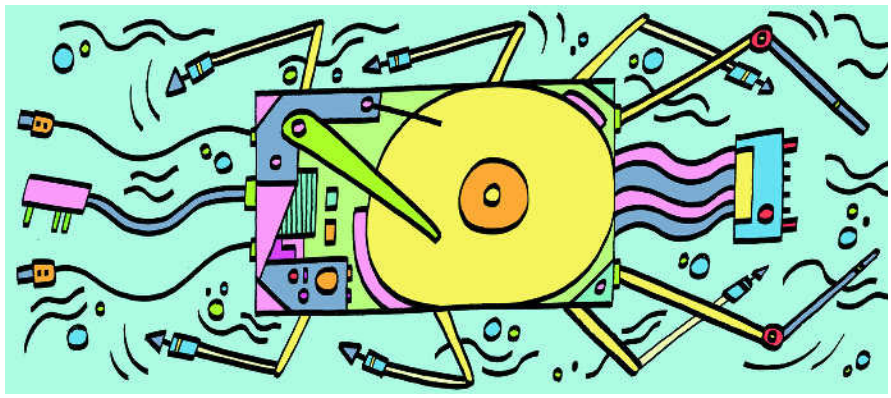

Where the cloud meets the ground

Data centres are quickly evolving into service factories

IT IS almost as easy as plugging in a laser printer. Up to 2,500 servers—in essence, souped-up personal computers—are crammed into a 40-foot (13-metre) shipping container. A truck places the container inside a bare steel-and-concrete building. Workers quickly connect it to the electric grid, the computer network and a water supply for cooling. The necessary software is downloaded automatically. Within four days all the servers are ready to dish up videos, send e-mails or crunch a firm's customer data.

This is Microsoft's new data centre in Northlake, a suburb of Chicago, one of the world's most modern, biggest and most expensive, covering 500,000 square feet (46,000 square metres) and costing $500m. One day it will hold 400,000 servers. The entire first floor will be filled with 200 containers like this one. Michael Manos, the head of Microsoft's data centres, is really excited about these containers. They solve many of the problems that tend to crop up when putting up huge data centres: how to package and transport servers cheaply, how to limit their appetite for energy and how to install them only when they are needed to avoid leaving expensive assets idle.

But containers are not the only innovation of which Mr Manos is proud. Microsoft's data centres in Chicago and across the world are equipped with software that tells him exactly how much power each application consumes and how much carbon it emits. "We're building a global information utility," he says.

Engineers must have spoken with similar passion when the first moving assembly lines were installed in car factories almost a century ago, and Microsoft's data centre in Northlake, just like Henry Ford's first large factory in Highland Park, Michigan, may one day be seen as a symbol of a new industrial era.

Before Ford revolutionised carmaking, automobiles were put together by teams of highly skilled craftsmen in custom-built workshops. Similarly, most corporate data centres today house armies of "systems administrators", the craftsmen of the information age. There are an estimated 7,000 such data centres in America alone, most of them one-off designs that have grown over the years, reflecting the history of both technology and the particular use to which it is being put. It is no surprise that they are egregiously inefficient. On average only 6% of server capacity is used, according to a study by McKinsey, a consultancy, and the Uptime Institute, a think-tank. Nearly 30% are no longer in use at all, but no one has bothered to remove them. Often nobody knows which application is running on which server. A widely used method to find out is: "Let's pull the plug and see who calls."

Limited technology and misplaced incentives are to blame. Windows, the most pervasive operating system used in data centres, allows only one application to run on any one server because otherwise it might crash. So IT departments just kept adding machines when new applications were needed, leading to a condition known as "server sprawl" (see chart 3). This made sense at the time: servers were cheap, and ever-rising electricity bills were generally charged to a company's facilities budget rather than to IT.

To understand the technology needed to industrialise data centres, it helps to look at the history of electricity. It was only after the widespread deployment of the "rotary converter", a device that transforms one kind of current into another, that different power plants and generators could be assembled into a universal grid. Similarly, a technology called "virtualisation" now allows physically separate computer systems to act as one.

Virtually new

The origins of virtualisation go back to the 1960s, when IBM developed the technology so that its customers could make better use of their mainframes. Yet it lingered in obscurity until VMware, now one of the world's biggest software firms, applied it to the commodity computers in today's data centres. It did that by developing a small program called hypervisor, a sort of electronic traffic cop that controls access to a computer's processor and memory. It allows servers to be split into several "virtual machines", each of which can run its own operating system and application.

Voracious 3

Worldwide spending, $bn, on:

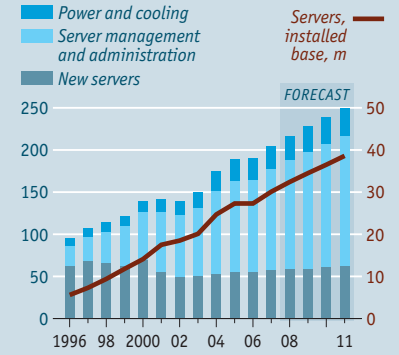

Source: IDC

"In a way, we're cleaning up Microsoft's sins," says Paul Maritz, VMware's boss and a Microsoft veteran, "and in doing so we're separating the computing workload from the hardware." Once computers have become more or less disembodied, all sorts of possibilities open up. Virtual machines can be fired up in minutes. They can be moved around while running, perhaps to concentrate them on one server to save energy. They can have an identical twin which takes over should the original fail. And they can be sold prepackaged as "virtual appliances".

VMware and its competitors, which now include Microsoft, hope eventually to turn a data centre—or even several of them—into a single pool of computing, storage and networking resources that can be allocated as needed. Such a "real-time infrastructure", as Thomas Bittman of Gartner calls it, is still years off. But the necessary software is starting to become available. In September, for instance, VMware launched a new "virtual data-centre operating system".

Perhaps surprisingly, it is Amazon, a big online retailer, that shows where things are heading. In 2006 it started offering a computing utility called Amazon Web Services (AWS). Anybody with a credit card can start, say, a virtual machine on Amazon's vast computer system to run an application, such as a web-based service. Developers can quickly add extra machines

when needed and shut them down if there is no demand (which is why the utility is called Elastic Computing Cloud, or EC2). And the service is cheap: a virtual machine, for instance, starts at 10 cents per hour.

If Amazon has become a cloud-computing pioneer, it is because it sees itself as a technology company. As it branched out into more and more retail categories, it had to develop a sophisticated computing platform which it is now offering as a service for a fee. "Of course this has nothing to do with selling books," says Adam Selipsky, in charge of product management at AWS, "but it has a lot to do with the same technology we are using to sell books."

Yet Amazon is not the only big online company to offer the use of industrial-scale data centres. Google is said to be operating a global network of about three dozen data centres loaded with more than 2m servers (although it will not confirm this). Microsoft is investing billions and adding up to 35,000 servers a month. Other internet giants, such as Yahoo!, are also busy building huge server farms.

In some places this has led to a veritable data-centre construction boom. Half a dozen are being built in Quincy, a hamlet in the middle of America's Washington state, close to the Columbia River. The attraction is that its dams produce plenty of low-cost power, which apart from IT gear is the main input for these computing farms. On average, cooling takes as much power as computing. Microsoft's new data centre near Chicago, for instance, has three substations with a total capacity of 198MW, as much as a small aluminium smelter.

But cheap electricity is only one, albeit important, criterion for choosing the site of a data centre. Microsoft currently feeds 35 sets of data into an electronic map of the world, including internet connectivity, the availability of IT workers, even the air quality (dry air makes a good coolant), to see where conditions are favourable and which places should be avoided. Apparently Siberia comes out well.

Google, for its part, seems to be thinking of moving offshore. In August it applied for a patent for water-based data centres. "Computing centres are located on a ship or ships, anchored in a water body from which energy from natural motion of the water may be captured, and turned into electricity and/or pumping power for cooling pumps to carry heat away," says the patent application.

Many chief information officers would love to take their IT infrastructure out to sea and perhaps drown it there. Even as demand for corporate computing continues to increase, IT budgets are being cut. At the same time many firms' existing IT infrastructure is bursting at the seams. According to IDC, a market-research firm, a quarter of corporate data centres in America have run out of space for more servers. For others cooling has become a big constraint. And often utilities cannot provide the extra power needed for an expansion.

Fewer, bigger, better

So IDC thinks that many data centres will be consolidated and given a big makeover. The industry itself is taking the lead. For example, Hewlett-Packard (HP) used to have 85 data centres with 19,000 IT workers worldwide, but expects to cut this down to six facilities in America with just 8,000 employees by the end of this year, reducing its IT budget from 4% to 2% of revenue.

Other large organisations are following suit. Using VMware's software, BT, a telecoms firm, has cut the number of servers in its 57 data centres across the world from 16,000 to 10,000 yet increased their workload. The US Marine Corps is reducing the number of its IT sites from 175 to about 100. Both organisations are also starting to build internal clouds so they can move applications around. Ever more firms are expected to start building similar in-house, or "private", clouds. The current economic malaise may speed up this trend as companies strive to become more efficient.

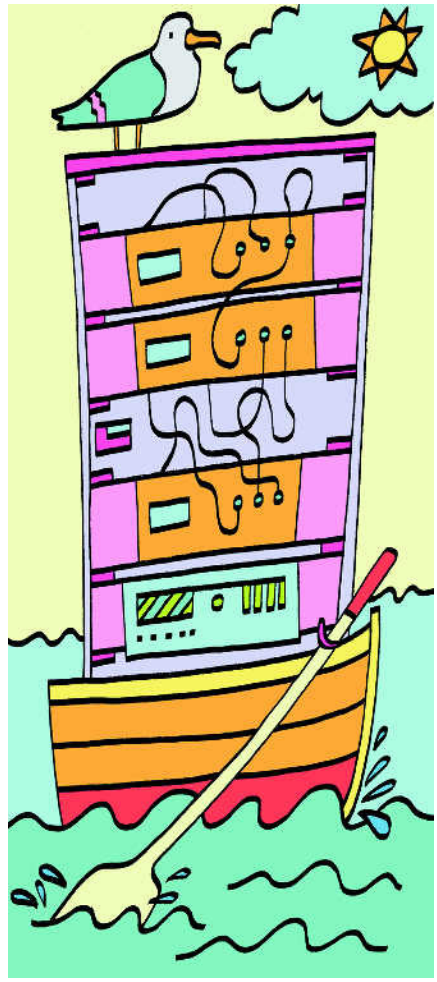

But to what extent will companies outsource their computing to "public" clouds, such as Amazon's? James Staten of Forrester Research, a market-research firm, says the economics are compelling, particularly for smaller firms. Cloud providers, he says, have more expertise in running data centres and benefit from a larger infrastructure. Yet many firms will not let company data float around in a public cloud where they could end up in the wrong hands. The conclusion of this report will consider the question of security in more detail.

It does not help that Amazon and Google recently made headlines with service interruptions. Few cloud providers today offer any assurances on things like continuity of service or security (called "service-level agreements", or SLAs) or take on liability to back them up.

As a result, says Mr Staten, cloud computing has not yet moved much beyond the early-adopter phase, meaning that only a few of the bigger companies are using it, and then only for projects that do not critically affect their business. The *Washington Post*, for instance, used Amazon's AWS to turn Hillary Clinton's White House schedule during her husband's time in office, with more than 17,000 pages, into a searchable database within 24 hours. NASDAQ uses it to power its service providing historical stockmarket information, called Market Replay.

Stefan van Overtveldt, the man in charge of transforming BT's IT infrastructure, thinks that to attract more customers, service providers will have to offer "virtual private clouds", fenced off within a public cloud. BT plans to offer these as a service for firms that quickly need extra capacity.

So there will be not just one cloud but a number of different sorts: private ones and public ones, which themselves will divide into general-purpose and specialised ones. Cisco, a leading maker of networking gear, is already talking of an "intercloud", a federation of all kinds of clouds, in the same way that the internet is a network of networks. And all of those clouds will be full of applications and services. ■

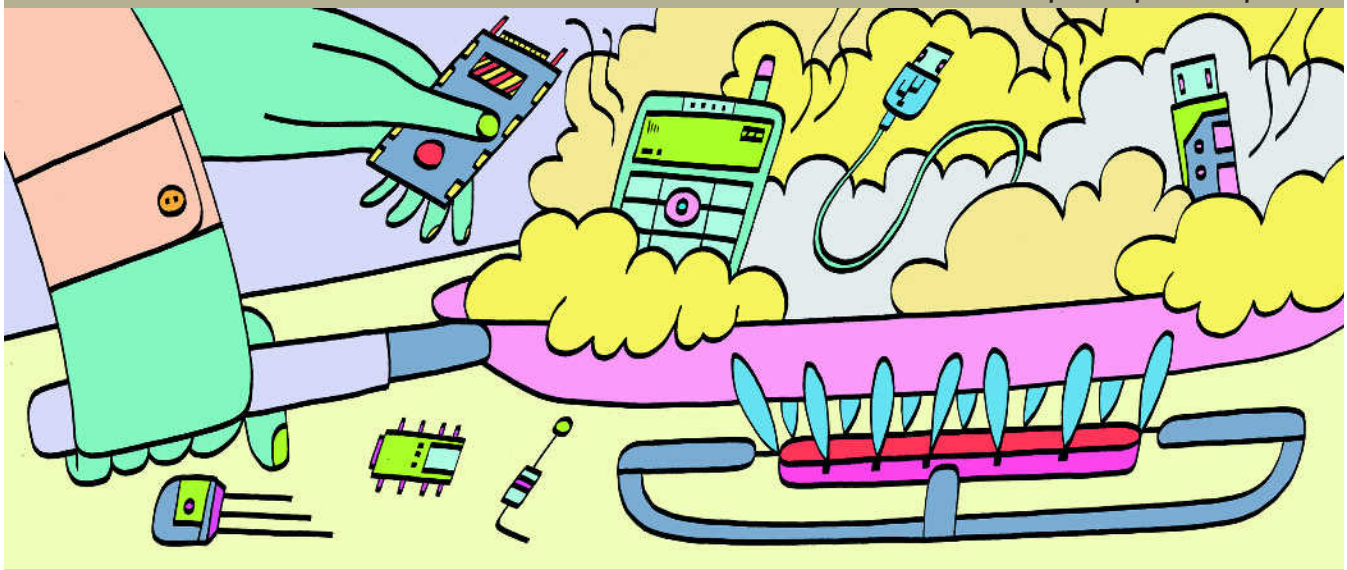

Creating the cumulus

Software will be transformed into a combination of services

WANT to become a programmer in 20 minutes? With the program from Iceberg, a start-up, you can. Just open the "getting started wizard" and pick the type of application, say "project management". Type in its "business objects"—things like "client", "team" and "members"—and tell the tool how they relate to each other. Then design some input forms and define the process for getting a project done. Another click and you are ready to go.

In reality, of course, things are not that simple. And many professional programmers will scoff at the development tool as a mere toy. Yet Iceberg and similar outfits demonstrate that geeks are losing their monopoly on programming. Now, with a bit of patience, anybody can create a simple application, for instance, to collaborate with colleagues or to draw an online map.

This democratisation of programming, however, is only a small part of something much deeper: a fundamental change in the nature of software. It is not just that more and more software will become a service delivered online. More importantly, applications, web-based or not, will no longer come as a big chunk of software, but will be made up of a combination of electronic services—a shift that has picked up a lot of speed since computing began moving into the cloud.

To understand this new way of building applications, known as "service-oriented architecture" (SOA), think of a culinary analogy. Whereas the old chunk of software resembles a precooked meal that just has to be popped into the oven, the new architecture is more like a restaurant. It is a service in itself but also a combination of sub-services. There is the waiter who takes the order and conveys it to the kitchen. There is the cook who prepares the food. And there are the cleaners who keep the place tidy. Together they create the "application": a restaurant.

An attack of the vapours

The importance of this shift from a monolithic product to services is hard to overstate. In a sense, it has seeded the cloud, allowing the droplets—the services that make up the electronic vapour—to form. It will allow computing to expand in all directions and serve ever more users. The new architecture also helps the less technically minded to shape their own clouds, using such tools as Iceberg's.

Just as for the industrialisation of data centres, there is a historic precedent for this shift in architecture: the invention of movable type in the 15th century. At the time, printing itself was not a new idea. But it was Gutenberg and his collaborators who thought up the technologies needed to make printing available on a mass scale, creating letters made of metal that could be quickly assembled and re-used.

Similarly, the concept of modularity has been around since the early days of computing. "Everything in computer science is to just write less code. What is the technique for writing less code? It's called subroutines," said Bill Gates, Microsoft's founder, in a recent interview. A subroutine is a part of a program that can be reused, just like movable type. The idea, says Mr Gates, has always been to apply this principle of a subroutine more broadly.

Yet this did not happen, mainly because the cost of computing fell much faster than that of communications. Ever cheaper and more powerful chips made it possible to move from mainframes to minicomputers to personal computers (PCs) and now to hand-held devices. But connecting all these pieces remained difficult and expensive, which meant that such devices all had to come with their own data and chunky programs. Now, thanks to plenty of cheap bandwidth and more and more wireless connectivity, computing is able to regroup into specialised services, or Mr Gates's subroutines: "We now live in a world where…[a] subroutine can exist on another computer across the internet."

Part of Gutenberg's genius was to recognise the need for all the letters to be identical in height so they could be easily combined. Similarly, for computing services to work there had to be robust technical standards. Only a few years ago this seemed far beyond the IT industry's reach. Most firms insisted on their proprietary technology, mostly to lock in their customers. Again, cheaper communications helped to bring about change. The success of the internet demonstrated the huge benefits of open standards and forced vendors to agree on common ways for their wares to work together. One result is a stack of something called "web-services" standards.

Service-oriented architecture first showed up in open-source software but was quickly adopted by big enterprise-software vendors because they had a pressing need for it, says Jim Shepherd of AMR Research, a consultancy. Big software vendors, for instance, had to find a way to untangle the hairball of code that their products had become, or else they themselves would choke on it. Customers wanted more flexible and extensible programs. ▸▸

Think back to the gastronomic example. A precooked meal is hard to change, and so are traditional software applications. By contrast, a restaurant can easily change its menu and its style of operation. Similarly, SOA-based software allows companies to alter their business processes, such as the way they handle orders to collect cash.

SAP, a German software firm, was one of the first companies to put this service-oriented architecture front and centre. Starting in 2003, it developed, among other things, a new corporate-software package that did away with monolithic applications, such as programs to keep track of a company's finances or manage its relationship with customers. Instead, it introduced a collection of re-usable components that could be strung together at will.

IBM, too, is a fan of SOA and web services. But its approach is somewhat different, given that it does not sell business applications but makes most of its money from IT services and software to manage the underlying computing infrastructure. IBM uses SOA mainly to help firms integrate their increasingly complex and disparate IT systems. Its software turns them into a collection of services that can be woven into business processes.

The approaches may be different but the vision is the same: to create IT systems that adapt to the business needs of companies and allow them to connect. "When I want to do something new", explains Steve Mills, the boss of IBM's software group, "I do not need to build a new application but can use the pieces I already have." To Peter Zencke, who led the development of SAP's new package, its most exciting feature is that "any of the process components can now become a service provided by some other firm."

Despite millions of dollars spent on marketing SOA, it has not really taken off yet. But many web-based applications for consumers rely on this concept. The prime example is Google Maps. When the online giant launched the service, programmers quickly figured out how to mix the maps with other sources of information. This is how, for example, Housingmaps.com was created, a combination of a Google map with the rental and sales listings from Craigslist, a website for classified ads. It was one of the first "mash-ups", as such combinations have come to be called.

Since then the number of such mash-ups has exploded, thanks mainly to services like Microsoft's Popfly and Yahoo! Pipes. In essence these are graphical programming tools. Users drag and drop "modules"—data feeds providing such information as pictures, headlines and search results—and weave them together.

Most of these mash-ups are still toys, but firms offering software as a service have started offering similar combinations. In April Salesforce.com and Google announced that they would integrate their online services. Users of Salesforce, which helps firms manage their customer relationships, can now quickly switch to Google's web-based applications.

Smaller firms have already started to weave a network of services. OpSource, a Silicon Valley start-up, for instance, provides basic services for other saas firms and web companies. TriCipher, another Californian newcomer, authenticates users of web applications. Ribbit, for its part, allows these services to add voice communications to their offerings.

Yet it is unlikely that the software cloud will end up as a vast nebula of thousands of specialised services. Even creating a service-oriented architecture is "no silver bullet" against complexity, in a famous phrase by Frederick Brooks, an elder of computer science. Although web services allow online offerings to connect, for instance, it is costly to synchronise their data. And it would not make sense for any firm to bet its business on simple mash-ups.

As software markets mature, they tend to form two kinds of clumps: integrated suites of applications, and platforms on top of which others can build programs. Both forms are already emerging. On the applications side there is Google Apps and Zoho, which is even more comprehensive. It encompasses a total of 18 applications, including word processing, project management and customer-relationship management (CRM).

As for platforms, there are already plenty, in different shapes and sizes. For enterprise applications, SAP has built one called Netweaver. Oracle offers something similar called Fusion. Last year, Salesforce launched a "platform as a service", allowing other firms to use the plumbing that supports its own CRM offering.

More recently platforms for consumer services have been proliferating. Facebook, a social network, was the first to become one in 2007. Other big online firms have followed suit or will do so soon: Google with App Engine, Yahoo! with Y!OS and Microsoft with a "cloud operating system" thought to be called Windows Strata. Some predict a platform war to rival the epic fights between Microsoft's Windows and Apple's Macintosh.

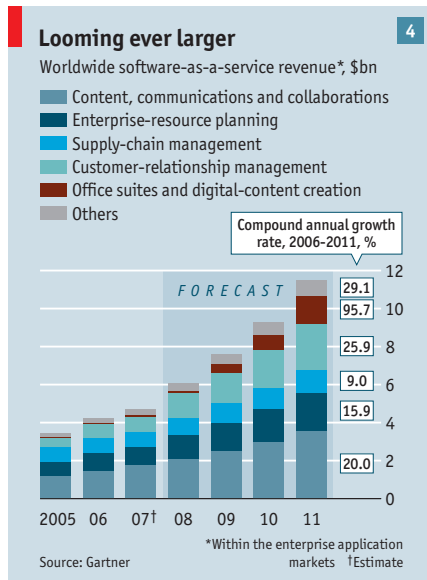

Never say die

What shape will the software cloud take, other than being a vast collection of services? In one way it will look much like the old software world. There will be a few big platforms, akin to today's operating systems, and most applications will be written to one of these platforms.

What is less clear is just how much of business and consumer software will migrate into the cloud, and how fast. The answer depends on whom you ask. Unsurprisingly, Marc Benioff, Salesforce's founder and chief executive, argues that web applications will spell the "death of software". But people are not about to throw out their powerful PCs or other "client" devices, if only because many of them still work offline at times. Similarly, companies will always want to keep some applications in-house, for reasons of security, regulation or simply to maintain control. Ray Ozzie, Microsoft's chief software architect, promotes something called "software plus services", meaning that customers will settle on "the right mix of old and new stuff".

If history is any guide, Mr Ozzie is more on the mark. Even the biggest changes in IT have never spelt the death of anything, notes Josh Greenbaum of Enterprise Applications Consulting. IBM, for instance, is still making money with mainframes.

So the software cloud, just like its hardware underpinnings, will be very diverse. But how will people make use of this kind of computing? ■

On the periphery

The cloud's communications with its clients will become ever more intelligent and interactive

IT WILL take something with a lot more bang to replace a medium that is thousands of years old. That was the prevailing reaction when Amazon last November announced the launch of Kindle, an electronic book reader the size of a paperback that can store more than 200 volumes. Yet by the end of this year Amazon will have sold nearly 380,000 Kindles, says Mark Mahaney, an analyst with Citigroup, a bank. "Turns out the Kindle is becoming the iPod of the book world," he recently wrote in a note to clients, in a reference to Apple's iconic music player.

It is certainly not the Kindle's looks that explain its success. Compared with the iPod, its design looks very last century. Software and battery life, too, leave a lot to be desired. The chief attraction of the device is the ease with which it can be used to buy books and other content. Equipped with a mobile-phone modem, the Kindle can simply pull new reading material out of the air. Users do not even have to have a wireless service contract. "Our vision is to have every book that has ever been in print available in less than 60 seconds," explains Jeff Bezos, Amazon's boss.

It remains to be seen whether the Kindle will become a cultural phenomenon like the iPod, of which around 160m have been sold so far. Amazon, for its part, is downplaying the Kindle's success and will not confirm any sales estimates. But it is safe to say that, once the next generation of wireless networks is up and running, hundreds of millions of devices will come, like the Kindle, with built-in radio connectivity (see chart 5). Digital cameras will automatically upload pictures. Smart meters will send readings of how much electricity a house consumes. All kinds of sensors will be able to send messages, even things like dipsticks when tanks of liquid are low.

The relationship of these devices to cloud computing may not be obvious. But if huge data centres and applications make up the cloud itself, then all the hardware and software through which it connects and communicates with the real world are its periphery. In IT speak, this is known as the "front end" or "client side".

As the Kindle and other examples show, this layer does not have much to do with the user interface or client device of old. It will do a lot of computing itself. It will come in all shapes and sizes, depending on what the user wants to do. And it will not just distribute information, as the web does, but collect it as well. The analogy that springs to mind here is a theatre performance with audience participation: the electronic cloud will adapt to whatever it engulfs.

As you like it

Just like computing itself, the dominant user interface has evolved continually. In the days of the mainframe, when computers and their peripherals filled entire rooms, people communicated with these machines first via punch cards and then via green-glowing monitors, which were simply dumb terminals. Only with the rise of personal computers did the user interface become more intelligent, responsive and graphical.

The first version of the web was thus a brief step backward. To be sure, browsers brought colour and graphics to the hitherto text-based internet, but they were as dumb as the mainframe terminal. This has changed only in recent years. A bundle of web-development techniques dubbed AJAX and multimedia software such as Adobe's Flash and Microsoft's Silverlight now allow programmers to write what are called "rich internet applications" (RIA).

Chatterboxes 5

Worldwide shipments of electronic products with built-in wireless capability, bn

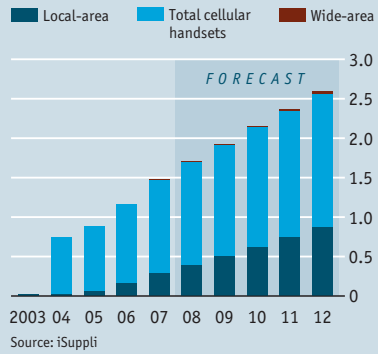

Source: iSuppli

Whatever the buzzword, the principle is much the same. Servers no longer dish up simple hypertext markup language (HTML), the web's early lingua franca. Increasingly, web pages are bona fide pieces of software that are executed in the browser. Users of Web 2.0 sites who venture into menu items such as "view source" in their browsers can sometimes see thousands of lines of code.

In recent months the browser has become even more of a platform for other programs, akin to an operating system such as Windows. The main driver of this trend is Google, with its huge strength in distribution that can only gain from more and more software being offered as a service. In May 2007 the Silicon Valley firm launched Gears, a program that allows web applications to be used offline, and in September this year it released a new browser called Chrome. Its most important feature is that it can execute several sophisticated web applications at once.

Although for now the internet browser will remain the main vehicle for people to interact with the cloud, other forms are coming to the fore. One is the "widget", a snippet of code that often lives on a PC's desktop and allows the user to get a quick personalised view of a set of data. The idea is that a salesperson, for instance, should not have to fire up an entire application for customer-relationship management to find out which leads to follow up.

More importantly, there is now a greater variety of hardware through which to access the cloud. Already, desktop and laptop computers are starting to lose their monopoly for surfing the web as smaller devices such as smart mobile phones and various forms of portable computers start to compete with them.

Asus, a Taiwanese computer-maker, started the trend when it launched a small, cheap laptop called "Eee" a year ago. Now there are dozens of these devices. Gartner reckons that 5.2m of these "mini-notebooks" will be sold this year, 8m next and as many as 50m in 2012.

Perhaps the best indicator of things to come is Intel, a huge chipmaker. It made a fortune selling processors for servers, personal computers and laptops. In June the ▸▸

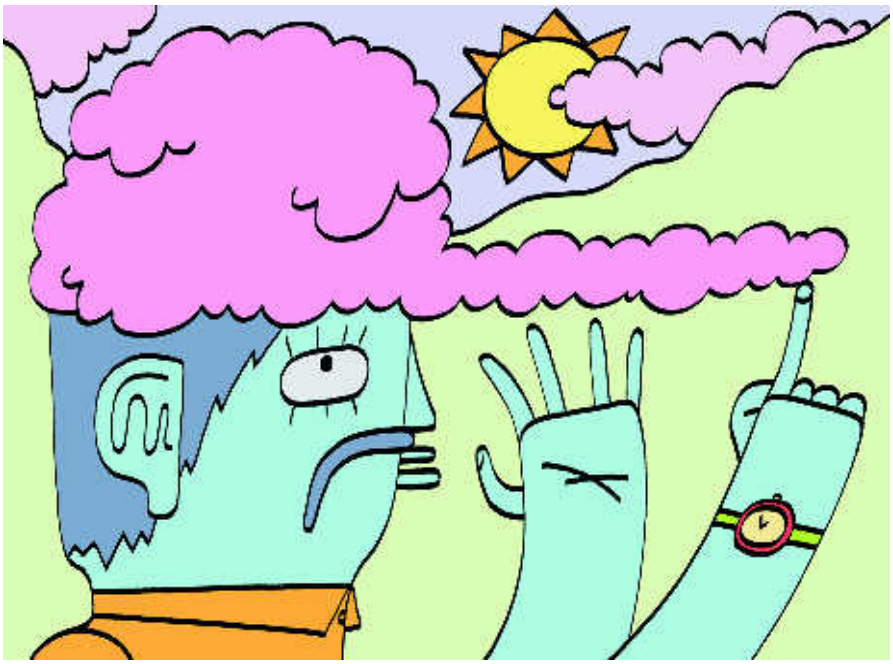

firm launched a new line of chips called Atom, designed to power what it calls "netbooks" and "mobile internet devices" (MIDs), mainly intended for surfing the web. Intel is also the driving force behind WiMAX, a technology for wireless broadband access to the internet. It wants to put a WiMAX radio chip into as many devices as possible, from portable computers to specialised gadgets such as the Kindle.

Apple's iPhone and its App Store, which allows iPhone and iPod owners to download applications, also provide a foretaste of how important wireless devices will be for the cloud. Apple launched App Store only in July. Two months later it had already tallied 100m downloads, meaning that it took off much faster than Apple's highly successful iTunes music store. Many of the programs on offer connect to the cloud, including news feeds, multi-player games and a service that keeps track of the latest polls for America's presidential election.

You can take it with you

The plethora of devices wirelessly connected to the internet will speed up a shift that is already under way: from a "device-centric" to an "information-centric" world, in the words of VMware's Paul Maritz. Up in the cloud there will be a body of data for each individual that will accompany them through life, he explains, and it will not be tied to any particular device, as it is today.

Again, what will make this possible is virtualisation—this time of client devices, not servers. With the help of software from VMware and others, some firms have already virtualised their employees' desktop computers, which allows them to be managed centrally. Operating systems and applications will no longer run only on the employee's PC but on a virtual machine in a data centre that can be accessed remotely, theoretically from any PC in the world. Sooner or later mobile devices will also become virtualised. Users will be able to use their applications and data on whichever gadget they have at hand.

Yet the cloud's interface is designed not merely to provide information but to gather it as well. The future belongs to services that respond in real time to information provided either by their users or by non-human sensors, predicts Tim O'Reilly, the founder of O'Reilly Media, a publisher of technology books who coined the term "Web 2.0". Such "live applications", he says, will get better the more data they are able to collect—and there will be plenty as the cloud expands.

One of the first examples of such a service was Google. What originally put the search service ahead of the competition when it was launched a decade ago was its way of harvesting the information provided by web users in linking to other sites: the more links point to a page, the more useful it must be. These days most links are generated by computers, so the original form of this "page rank" algorithm has long since been scrapped. But Google's approach is still the same: mining information provided by web users, such as their search histories, to provide more relevant search results and more effective and targeted advertising.

The direct link to users also allows firms such as Google continuously to improve their interface, something traditional software-makers were not able to do. At any given time Google is running dozens of tests to optimise the look and feel of its offerings. This makes web applications far less technology-driven and much more user-oriented, says IBM's Mr Wladawsky-Berger. "They are much more inspired by what goes on in the real world."

A raft of start-ups is also trying to build a business by observing its users, in effect turning them into human sensors. One is Wesabe (in which Mr O'Reilly has invested). At first sight it looks much like any personal-finance site that allows users to see their bank account and credit-card information in one place. But behind the scenes the service is also sifting through its members' anonymised data to find patterns and to offer recommendations for future transactions based, for instance, on how much a particular customer regularly spends in a supermarket. Wireless devices, too, will increasingly become sensors that feed into the cloud and adapt to new information.

Nokia, for its part, is planning to build all kinds of sensors into mobile phones to monitor things like movement, barometric pressure or even the owner's health, which many experts expect to become a big new trend. Sensors could also be used to record people's activities, creating what some already term a "lifelog"—raising all kinds of privacy concerns.

As wireless technology gets better and cheaper, more and more different kinds of objects will connect directly to the cloud. SAP, the German software-maker, has launched a research project called "The Internet of Things" to see what can be done with the resulting information. As part of that project, an initiative called the "Future Factory" is now under way to investigate how intelligent tags can make manufacturing more adaptive and efficient.

More and more data get you only so far, however. In the end, Google's search results and its text-based online advertisements are relevant to users only because the firm has devised clever ways to sift through them, says Hal Varian, the firm's chief economist. The big challenge of the cloud will be to connect the myriad data in it and make them profitable. ■

Highs and lows

As IT gets cloudier, the economics of the business will change

EVEN elephants can die. In 1993 extinction came close for IBM, then the world's largest computer-maker (it has since been overtaken by HP). Its mainframe business was collapsing and profits were plummeting. At that point Louis Gerstner took over as chief executive and managed to turn the company around. "Only a handful of people understand how precariously close IBM came to running out of cash in 1993," he writes in his memoir, "Who Says Elephants Can't Dance?" "Whether we would have had to file for bankruptcy, I can't say."

There are many reasons why IBM nearly went belly-up, not least the fact that Big Blue had become a bureaucratic monster. But most critically, it had failed to adapt to the industry's first big platform shift, which only really made itself felt in the early 1990s: the move from mainframes to smaller machines, first so-called minicomputers, then personal computers. "IBM was slow, very slow, in delivering distributed computing, and many small companies moved in to fill the gap," Mr Gerstner writes.

It is unlikely that the move into the cloud will produce a similar near-death experience or even a real casualty—if only because IBM still stands as a warning. But that does not mean that the structure of the IT industry will remain unchanged, nor that the economics will stay the same. Once the IT sky really clouds over, individual firms' share of the global IT budget (see chart 6) will shift.

The move to distributed computing, which started in the mid-1980s, led to a big change in the IT industry. In the era of the mainframe computing came in a vertically integrated package, mainly from IBM. With distributed computing the industry became a stack of horizontal layers. In corporate IT these were mainly hardware, the network, infrastructure software (such as operating systems and databases), enterprise applications and IT services.

Not all of these layers were created equal. Computer-makers commanded a thicker one, for instance, but software companies were more profitable. The key program was the operating system, both on servers and on personal computers ("clients"). It was the standard to which other components of IT systems had to conform. Usually this was a version of Windows, which made Microsoft the IBM of this new era of computing.

Cloud computing is unlikely to bring about quite such a dramatic shift. In essence, what it does is take the idea of distributed computing a step farther. Still, it will add a couple of layers to the IT stack. One is made up of the cloud providers, such as Amazon and Google. The other is software that helps firms to turn their IT infrastructure into their own cloud, known as a "virtual operating system for data centres".

Drawing a neat diagram of the IT stack will also become increasingly difficult because the layers are becoming less distinct. In a world of services it often does not make sense to think of hardware and software separately, argues Padmasree Warrior, the chief technology officer of Cisco. Both need to be blended to offer new services, she says.

Follow the money 6

Worldwide end-user spending, $trn

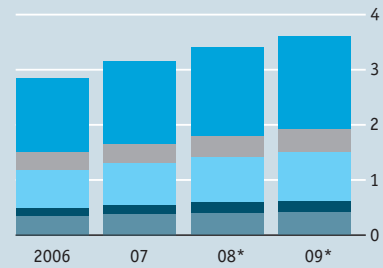

Source: Gartner *Forecast

Mix and match

Even though the IT stack may not change all that much, the perceived value of the different layers will shift, and with it the amount of profit IT firms can make from each of them. Who will lose and who will win depends on how much of computing ▸▸

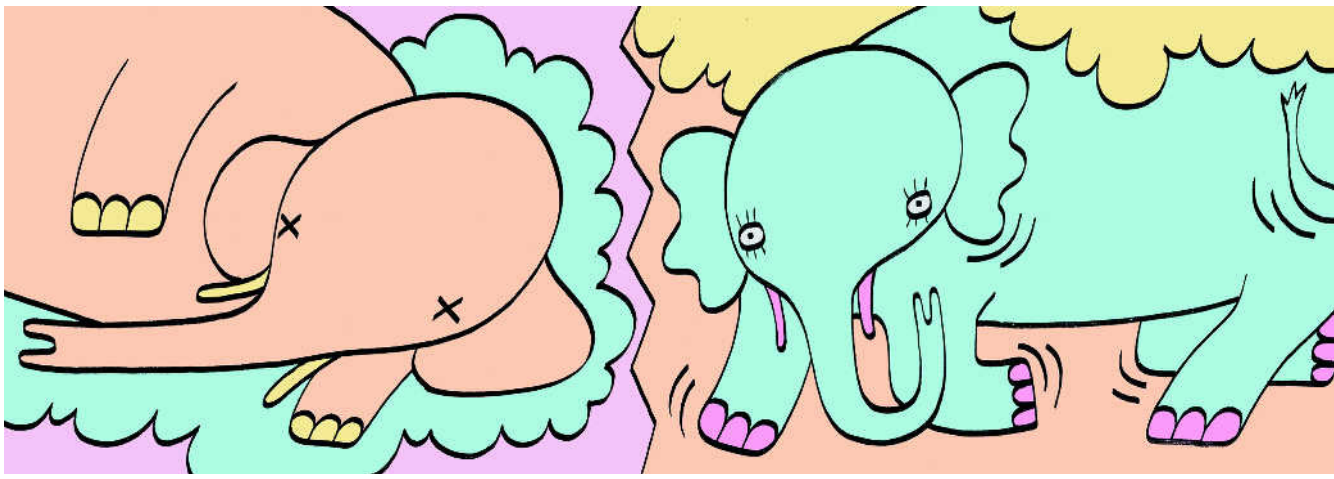

eventually moves into the cloud.

In the first round almost everybody in the IT industry will do well as the clouds are being built. The biggest winners are likely to be hardware-makers, says Mark Stahlman of Gartner: "Hardware always wins when new demand for computing is uncovered. And we haven't had such a sweeping global demand since the 1990s."

But in the longer term there will be relative winners and losers. The hardware business could actually find itself in the losing group. Its margins could get squeezed as cloud computing matures because there will be fewer customers with more buying power, says James Staten of Forrester Research. Large cloud providers can dictate how to build servers and at what price, he notes.

All that may explain why hardware-makers were among the first to jump on the cloud-computing bandwagon. So far, they have done only what Mr Staten calls "cloud-washing": relabelling existing products that help customers build a more flexible IT infrastructure. But they are also preparing for a time when more money can be made building clouds than building computers. IBM and HP, for instance, have teamed up with other firms and universities to design new cloud architectures.

Which side of the fence?

In the long run, says Mr Staten, hardware-makers may be torn between supplying cloud providers or becoming providers themselves. Being both will not be easy, because the firms would be competing with their biggest customers. Dell seems to have decided to be a cloud supplier. Sun Microsystems is a candidate to become a provider; it is offering a cloud-like service called Network.com, albeit not very successfully. HP and IBM, already used to the balancing act of selling hardware and providing IT services, will try to do both.

Makers of traditional software will find the going even tougher. With the advent of open-source software, in particular Linux, selling operating systems had already become less profitable. In a virtual world they will become even more commoditised, which is bad news for Microsoft. Many business applications no longer need a big, general-purpose server operating system but can use a specialised one, which should put pressure on prices. On client computers, more and more applications are written to run in browsers, not on any particular operating system.

Makers of business applications are also on the defensive. Traditionally they have made billions by selling their programs, often demanding hefty sums to install them and then charging an annual maintenance fee for upgrades and technical support. But this highly lucrative business model has come under increasing pressure, says Michael Cusumano, a professor at the Massachusetts Institute of Technology (MIT).

For one, he says, software vendors will have to find new ways to charge for their wares: in the cloud, tying licensing fees to the number of users, for instance, will be difficult, since services will mostly be consumed by other machines. More importantly, the corporate world has become less and less willing to buy software for large sums of money, so software firms listed on America's stockmarkets now make most of their profits from maintenance and other services (see chart 7). SAP will increase its annual maintenance fees to at least 22% of a program's value over the next few years, in line with those of Oracle, its main rival.

Yet the biggest challenge for software firms is to become providers of online services themselves, says Brent Thill of Citi Investment Research. So far they have moved slowly, offering SaaS only on the side, if at all. This was partly because their customers were not that keen. But more importantly, notes Mr Thill, the software houses are still wedded to their old business model. With SaaS they do not get a big upfront payment, only subscription fees.

Once Salesforce and NetSuite had shown that the SaaS model works, the incumbents began to move faster. In September last year, for instance, SAP presented "Business ByDesign", a package of web-based enterprise applications for smaller businesses. But success will not come easily. SAP has slowed down the introduction of the new service because it still needs to work out how to run it cheaply enough to make a reasonable profit.

Pure SaaS providers also have a lot on their minds. Some experts, such as Joshua Greenbaum of Enterprise Applications Consulting, reckon that few will ever be as profitable as traditional software firms. Although it is almost a decade old, Salesforce started making money only in 2006, mainly because it first had to spend heavily on marketing to attract customers. But now that the service has 1m users and revenues of more than $1 billion, these costs will come down, says the firm.

The companies that have the best chance of making money from the cloud are those that get things to connect and work together and help customers move their computing around. This is music to the ears of big IT firms, not least IBM. Nearly 80% of its revenues come from infrastructure software and IT services, which it can offer globally. HP is catching up, having taken over EDS, another big IT-services firm. Both Microsoft and SAP, for their part, believe that firms will want to have a choice in where to do their computing, as well as the flexibility to move things around over time.

Two potentially important contestants are rarely mentioned: Cisco and EMC, the leading makers of networking and storage gear respectively. Having invested a lot in software and services, Cisco has become more than just the source of most of the world's routers, the traffic cops of the internet. It is betting that in the cloud the network layer will become more important, for instance to ensure that computing workloads are able to move around securely. EMC, for its part, has made two dozen cloud-related investments and launched a cloud-infrastructure division.

Whoever manages to own the dominant operating system for the data centre could become a big winner. VMware is bound to have a shot at this. As well as being the market leader in virtualisation, it has the support of EMC, which holds 86% of the firm. But the competition is likely to be intense.

Will this prospective platform war produce a dominant company in the mould of IBM or Microsoft that is able to extract more than its fair share of the profits? Probably not, because it will be relatively easy to switch between vendors, says George Gilbert of Tech Strategy Partners, a consultancy. Nor is it likely that one firm will manage to build a global cloud monopoly. Although there are important economies of scale in building a network of data cen-

Serving up a treat 7

US-listed software-firms' sales, % of total

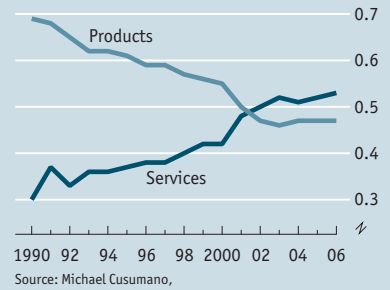

Source: Michael Cusumano, Massachusetts Institute of Technology

tres, the computing needs of companies and consumers vary too widely for one size to fit all.

Even if the cloud is likely to transform the IT industry, some things will stay the same. One is the importance of lock-in. If anything, companies and developers will be even more dependent on cloud platforms and applications than they are on the old kind. SaaS promotes the "hollowing out" of IT: a firm that needs to migrate to another system will no longer have the required expertise. When Facebook, say, makes a change to its platform, developers have no choice but to go along with it. Some are already calling for a "Cloud Computing Consortium", in the mould of the World Wide Web Consortium (W3C), to set standards that allow applications to migrate easily from one platform to another. One standard initiative, called "OpenSocial", already allows the same web-based application to run in several social networks, which are also clouds of sorts.

But standards go only so far. Some fear that one company could try to monopolise other key parts of the cloud; ironically, Microsoft worries that Google is doing exactly that with the online advertising market. To Steve Ballmer, Microsoft's boss, Google's advertising platform is like a flywheel that picks up speed as more websites attract more advertisers, and vice versa.

Eric Schmidt, Google's chief executive, denies any evil intent to achieve world domination. He argues, with some justice, that it would be hard for Google to control the cloud, if only for technical reasons: much of it is already based on open standards, and its loose structure does not lend itself to locking customers in.

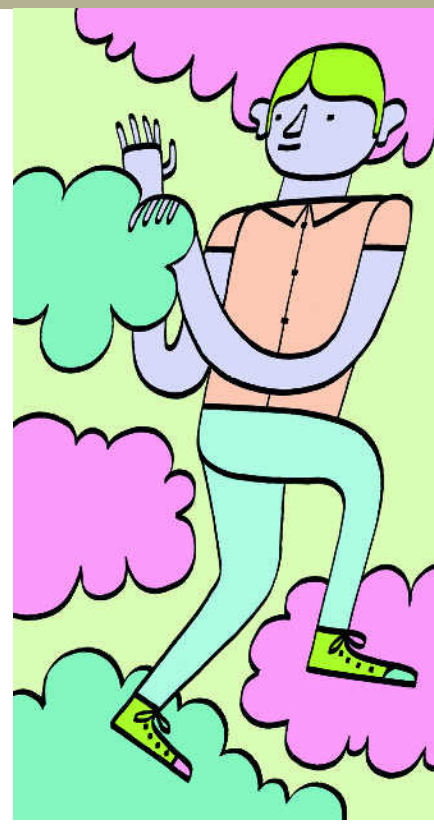

Mr Schmidt promises that Google will not lock its users in either. "Our competitive advantage is not from lock-in", he says, "but from having specialised knowledge of how to build data centres and how to build new software that is not reproducible, such as our search algorithm. This is how we make our money."

Yet Google is more like Microsoft than it likes to admit, says Nicholas Carr, a technology writer and blogger. Microsoft, he argues, achieved its dominant position in the PC world not least by commoditising products, such as the browser, that are complementary to its cash cows, such as Windows: as their cost came down, demand for Microsoft's products went up.

Similarly, Google's natural instinct is to do its utmost to encourage people to spend more time online, because that will give the company more opportunities to sell advertisements and collect data about them. According to Mr Carr, almost everything the company does—building huge data centres, fighting copyright restrictions, digitising the world's libraries, developing a new browser and, most recently, even helping to launch satellites—is aimed at increasing the use of the internet. "Google wants information to be free", he recently wrote in his blog, "because as the cost of information falls it makes more money."

But Google may never become as powerful as Microsoft because regulators are unlikely to let it. Microsoft was eventually put in the dock for abusing its monopoly because it got too greedy, pushing most of the rest of the industry to complain. Given that the world has already lived through the Microsoft drama and that Google will affect many more industries, the search company is likely to be restrained much earlier. The firm is currently in negotiations with the US Justice Department about a controversial advertising partnership between itself and one of its competitors, Yahoo!, which would further strengthen Google's position in online advertising.

Even if the economics of the cloud are still in flux, though, it is already clear that it will have far-reaching implications for businesses and for society as a whole. ■

The long nimbus

The cloud will make businesses more adaptable, interconnected and specialised—and often smaller

"BUSINESSES, as well as most organisations outside the business world, begin to shift from hierarchical processes to networked ones. Nearly every facet of human activity is transformed in some way by the emergent fabric of interconnection. This reorganisation leads to dramatic improvements in efficiency and productivity." So said *Wired* magazine, the central organ of Web 1.0, in July 1997 in an essay entitled "The Long Boom", arguing that the world was in for "25 years of prosperity, freedom and a better environment". Back then, the article reflected the general optimism that led up to the internet bubble. Now, after two busts, several wars and growing fears of global warming, it makes for somewhat surreal reading.

The cloud lends itself to similar hyperbole. Yet so far there has not been much debate about its economic fallout—probably because the "new economy" ended badly and the newest one is currently doing even worse. There will be many ways in which the cloud will change businesses and the economy, most of them hard to predict, but one theme is already emerging. Businesses are becoming more like the technology itself: more adaptable, more interwoven and more specialised. These developments may not be new, but cloud computing will speed them up.

Corporate IT has always promised to make companies more agile. In the 1990s many companies re-engineered their business processes when they started using a form of software called enterprise-resource planning (ERP), which does things such as managing a firm's finances and employees. But once these massive software packages were in place, it was exceedingly difficult to change them. Implement-

ing SAP, the market leader in ERP, is like pouring concrete into your company, goes an old joke among IT types.

This helps to explain why in many firms IT departments and business units have traditionally been at loggerheads. In recent years tensions have worsened. Companies must grapple with ever-changing markets and regulations, yet IT budgets are being cut. Many firms now have a huge backlog of IT projects.

Hence the interest in cloud computing. It turns capital expenditure into operational expenditure, which makes things much easier and cheaper. Instead of having to shell out a lot of money for, say, a server to test an application and, even with luck, wait a few weeks for it to be up and running, managers just have to whip out a credit card, open an account at Amazon Web Services (AWS) and fire up a virtual machine for a few dollars.

In many firms senior managers probably do not even know that business units are using AWS or similar services. In some organisations, however, it has already become a much-appreciated R&D tool. Pharmaceutical companies, for instance, are regularly tapping into AWS to calculate simulations. Sogeti, a European consultancy, has used a cloud built by IBM to allow it to test new ideas and quickly put together an IT system for a company-wide brainstorming event.

Cloud services have also been hugely successful with start-ups, which can now enjoy infrastructure of the same quality as large companies. In fact, AWS is probably the main reason why there are now so many firms offering all kinds of "Web 2.0" services. Their usefulness may sometimes be hard to gauge, but that is a good thing. It is a sign of lively "combinatorial innovation", made possible because entrepreneurs can cheaply try new combinations of technology, says Google's Mr Varian.

Many start-ups would probably not even exist without the cloud. Take Animoto, a service that lets users turn photos into artsy music videos using artificial intelligence. When it launched on Facebook, a social network, demand was such that it had to increase the number of its virtual machines on AWS from 50 to 3,500 within three days. "You could give me unlimited funding," says Adam Selipsky of AWS, "and I wouldn't know how to deploy that many servers in 72 hours."

Combinatorial innovation should also be made easier by the fact that the cloud will be a huge collection of electronic services based on standards. But this service-oriented architecture will be even more important for existing firms because it should free their inner workings—their "business processes"—from the straitjacket of their ERP systems and allow these processes to be more easily adapted, for instance to launch a new product.

Again, the software industry has been promising this for some years under the banner of service-oriented architecture, discussed in an earlier chapter. Yet the adoption of SOA has been slow and many projects have failed, says Chris Howard of the Burton Group, a consultancy. The reasons are not just technical but cultural; for example, some business units are not used to sharing data. Cloud computing will help resolve some of these problems. Many web-based services are built to be integrated into existing business processes.

Don't do it yourself

What effect will all this have on the nature of the firm? If IT systems really allow companies to become more modular and flexible, this should foster further specialisation. It will become even easier to outsource business processes, or at least those parts of them where firms do not enjoy a competitive advantage. Companies will increasingly focus on their "core" and shed the "context", in the words of Geoffrey Moore, managing director of TCG Advisors, a consultancy.

This also means that companies will rely more on services provided by others. They will increasingly form "process networks", a term for loosely connected groupings of specialised firms coined by John Hagel, a business strategist at Deloitte & Touche, an auditing firm. His prime example is Li & Fung, a company based in Hong Kong that has assembled a global process network of nearly 10,000 business partners in the clothing industry from which it puts together customised supply chains for clothes designers.

Both trends could mean that in future huge clouds—which might be called "industry operating systems"—will provide basic services for a particular sector, for instance finance or logistics. On top of these systems will sit many specialised and interconnected firms, just like applications on a computing platform. Yet this is only half the story. The cloud changes not only the plumbing and structure of firms and industries, known as the "transactional layer", but also their "interactional layer", a term coined by Andy Mulholland, chief technologist of Capgemini, a consultancy. He defines this as the environment where all the interactions between people take place, both within an organisation and with its business partners.

Despite all the technology that has en-

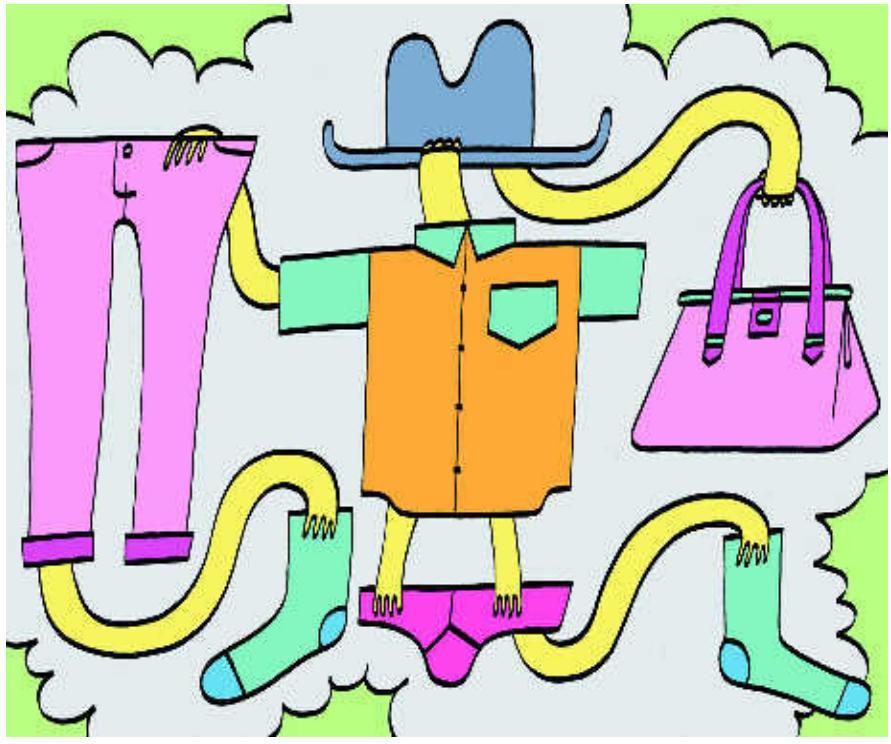

tered the workplace in recent years, so far this layer has not really changed. PCs certainly made people more productive, but most of their programs were not designed for collaboration. The enterprise applications they worked with were still centralised systems. And e-mail has in some ways made things worse as the flood of messages takes up lots of time and attention.

The dominant model is still that people first labour individually and then merge their respective efforts, says Mr Mulholland. "It's not much different from the age of paper," he writes in his book "Mesh Collaboration". "Collaboration often means pulling up your chair next to your colleague so you can look at the same screen."

Consumers have pulled ahead of companies in using cloud-based services that allow for better collaboration, such as blogs, wikis and social networks. The first generation of people that has grown up using all these tools is now entering the workforce. Being used to a culture of sharing information freely, these "digital natives" will be impatient with the rules of traditional corporate IT.

So it is helpful that firms have at last begun to embrace Web 2.0 technologies in earnest, a trend predictably called Enterprise 2.0. By 2013 companies around the globe will spend $4.6 billion on such tools, according to Forrester Research. What nobody knows is how firms will trade off the advantages of letting employees collaborate with the outside world against the associated risks, for instance that confidential information is leaked. Because of such security concerns (see chart 8), many firms block access to such sites as Facebook on company computers.

Companies may not have much choice but to open up, says Mr Mulholland. Employees will increasingly resist constraints on their use of technology, and they will have a growing need to reach beyond the corporate firewall. Twenty years ago, he argues, 80% of the knowledge that workers required to do their jobs resided within their company. Now it is only 20% because the world is changing ever faster. "We need to be open to new and unknown connections with people and content," he says.

Dangers within 8

How IT organisations feel about risk of employee-driven, unsanctioned use of Web 2.0*
June 2008, %

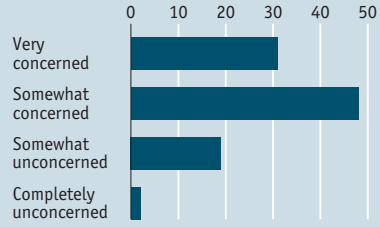

Source: Forrester Research

*Survey of 242 US IT decision-makers at companies with 500+ employees

Exploiting the "mesh", in the words of Mr Mulholland, will also mean that employees may build simple applications of their own that allow them quickly to automate repeat tasks, using internal and external IT services. A new product by Serena Software called Mashup Composer gives a foretaste of things to come. Using a visual interface, even the not so technically minded can quickly put together a service to deal, say, with travel requests or approve documents.

Yet the impact of the cloud will also be felt on a macroeconomic level. Just as it makes small firms more competitive, it will help developing economies to move ahead. "The biggest promise is that it will make computing cheaper and easier to use—and thus allow it to penetrate new markets," says Russ Daniels, chief technology officer for HP's cloud-services strategy.

Hop, skip and jump

The mobile phone has already enabled developing countries to skip fixed-line networks. Cloud computing could prove to be a similar "leapfrog" technology because it dispenses with the need to build a cumbersome IT infrastructure. "Software developers from a developing country can build just as great an application on our platform as somebody who lives in Palo Alto," says Mr Benioff of Salesforce.

Indeed, countries such as India will certainly take a big chunk of the global market for cloud services. Zoho, a popular suite of web-based applications, is operated by an Indian company, AdventNet. Indian hospitals are already offering specialised healthcare services this way. The insurance arm of ICICI, an Indian bank, has used the technology to come up with innovative services such as a personalised insurance for diabetes. Premiums are adjusted depending on how well policyholders stick to a fitness plan. All this suggests that the economic impact of the cloud may be felt not only in the IT industry itself but in other sectors too. The internet disrupted the music business; Google disrupted the media; cloud-based companies could become disrupters in other inefficient industries. Buzzwords such as "Health Care 2.0", "Banking 2.0" or even "Education 2.0" could soon acquire real meaning.

But given that the cloud is global by nature, how might it be regulated? The conclusion of this special report will offer some answers. ■

Computers without borders

The cloud may be the ultimate form of globalisation

READERS may be excused if they have never heard of Kinakuta. It is a tiny island dreamed up by Neal Stephenson in his novel "Cryptonomicon". In the science-fiction classic, libertarian technologists and entrepreneurs try to turn it into a data haven for such things as anonymous online banking and electronic money.

Reality has indeed caught up with fantasy. Iceland is about to become a Kinakuta of sorts. Data Islandia, a local company, is trying to establish the island as a vault for a growing pile of data that firms must retain in order to comply with all kinds of regulations. It has a compelling pitch. With its cool climate, abundant geothermal energy and secure remoteness, Iceland appears to be a prime location for data archives.

As often, however, truth is stranger than fiction. In a way, Data Islandia is erecting borders in the cloud: it intends to store European data according to European regulation and American bits according to American rules. What is more, to keep the data safe during transport, they are picked up with a "data scooter" (in essence a container filled with disk drives) and taken to Iceland by aeroplane—as though fibre-optic links had never been invented.

This illustrates the political tensions

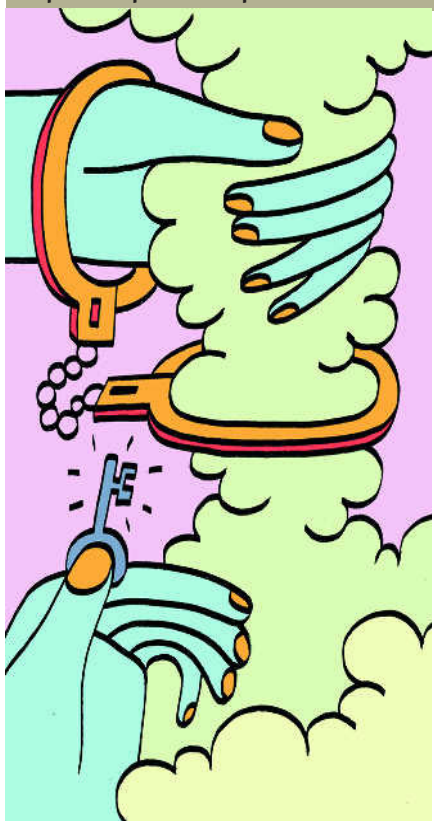

that will arise with the cloud. In one way it is the ultimate form of globalisation: vast virtualised computer systems and electronic services know no borders. Yet governments are likely to go to great lengths to avoid losing even more control.

When the internet went mainstream in the late 1990s, libertarian thinkers argued that cyberspace was a distinct place calling for laws and legal institutions of its own. After all, they said, it was built in such a way "that it interprets censorship as damage and routes around it". But many governments quickly found ways to block content they deemed offensive. Just look at China and its "great firewall".

Controlling where data are stored and how they are treated is harder, though, because information can float freely in the cloud. And it is not just undemocratic governments that want to control their citizens' and companies' data: indeed there are nearly as many sets of data regulation as there are countries. "If we wanted to be on the safe side in terms of regulation, we probably would need 95 individual data centres," says Chuck Hollis, a technologist at EMC, the leading maker of storage gear, which owns Mozy, a cloud service that allows users to back up their data.

There are technological fixes to this problem, too. Customers of Amazon's storage service, for instance, can have their data kept either in an American or a European data centre. In future, cloud providers will offer many more options for where data are kept and how they are protected. It seems possible that data as well as content will eventually travel with security, location and expiry policies attached.

It is when computers become virtual machines that things get really tricky. These days IT systems are at the core of many companies—and just like data, these systems can now live in a variety of places. What happens if they start to migrate to another country where power is cheaper or regulation laxer? Similarly, if services are a combination of elements provided in different jurisdictions, who is liable if something goes wrong?

The cloud's potential political and social effect is only now entering the public debate. IT firms are putting it on the agenda on both sides of the Atlantic. In mid-September, for instance, Google organised a discussion on the subject in Washington, DC. A few days later SAP presented a white paper in Brussels pointing out that policymakers are not aware of the dramatic economic impact of the "Future Internet", as the paper calls the cloud.

IT industry leaders note that officials from many countries have begun to take an interest in the cloud. Some just want data centres to be built in their country to create jobs; others are concerned about issues of law enforcement and jurisdiction. The danger, they say, is that cloud providers might be obliged to build more data centres than are needed and have to comply with many different regulatory regimes. Some of them have been floating the idea of "free-trade zones" for data centres where common rules would apply.

Yet such ideas appear at odds with another big question that could keep the cloud from growing: how to protect privacy. "Consumers expect their information will be treated the same way on the cloud as if it were stored at home on their own computers," says Ari Schwartz of the Center for Democracy and Technology, an advocacy group. Many of the devices that feed into it, such as sensors and cameras, will be intrusive. For example, Google's vans go round taking street-by-street pictures of cities for the company's online map service called "Streetview". The pictures are meant to help people find their way around, but also often show passers-by in embarrassing situations. To protect their identity, Google now blurs their faces and licence plates.

Hands off my data

But it is not only personal information that could get out into the open. Privacy is a worry for companies too—and not just because criminals or spies might intercept their data. Once they are in the cloud, governments can also get their hands on them more easily. SWIFT, the organisation that manages international bank transfers, is planning to build a data centre in neutral Switzerland. That will allow it to keep data about European transfers on the old continent, where it cannot be subpoenaed by the American government. Web-based e-mail is not safe either. Thanks to the Stored Communications Act, American law enforcers can read people's messages—and do not even have to tell the recipient.

Just as too much regulation may keep the cloud from rising high, so could lack of privacy. If consumers and companies cannot be sure that their information is safe, they will err on the side of caution. But despite all the caveats, the precipitation from the cloud will be huge. ■