


A Brief Guide to  
**CLOUD**  
Computing

An essential guide to  
the next computing  
revolution



**Christopher Barnatt**

**Christopher Barnatt** is Associate Professor of Computing & Future Studies in Nottingham University Business School, and the author of [ExplainingComputers.com](http://ExplainingComputers.com), [ExplainingTheFuture.com](http://ExplainingTheFuture.com) and their associated YouTube channels. He has written five previous books on computing and future studies, and lectures and consults widely on cloud computing and Web 2.0. Christopher also appears regularly on TV and the radio. You can follow him in the cloud at [twitter.com/ChrisBarnatt](https://twitter.com/ChrisBarnatt).

A BRIEF GUIDE TO

# Cloud Computing

Christopher Barnatt



ROBINSON

Constable & Robinson Ltd  
3 The Lanchesters  
162 Fulham Palace Road  
London W6 9ER  
www.constablerobinson.com

First published in the UK by Robinson,  
an imprint of Constable & Robinson Ltd, 2010

Copyright © Christopher Barnatt, 2010

The right of Christopher Barnatt to be identified as the author of this  
work has been asserted by him in accordance with the  
Copyright, Designs and Patents Act 1988

All rights reserved. This book is sold subject to the condition  
that it shall not, by way of trade or otherwise, be lent, re-sold,  
hired out or otherwise circulated in any form of binding or cover  
other than that in which it is published and without a similar condition  
including this condition being imposed on the subsequent purchaser.

A copy of the British Library Cataloguing in  
Publication data is available from the British Library

ISBN 978-1-84901-406-9

Printed and bound in the European Union

#### **Disclaimer**

While every effort has been made to ensure that the content in this book is as  
accurate as possible, no warranty or fitness is implied. The information is  
provided on an 'as is' basis, and the author and the publisher take no  
responsibility for any loss or damages arising from its use.

1 3 5 7 9 10 8 6 4 2

To Helen

# CONTENTS

Acknowledgements	ix
Preface	xi
<b>Part I: Cloud Computing Basics</b>	
Chapter 1 The Rise of Fluffy Computing	3
Chapter 2 Cloud Computing and Web 2.0	26
Chapter 3 Software in the Cloud	43
Chapter 4 Hardware in the Cloud	76
<b>Part II: Cloud Computing Implications</b>	
Chapter 5 Security, Privacy and Reliability	107
Chapter 6 The Second Digital Revolution	127
Chapter 7 The Battle for the Cloud	150
Chapter 8 Trouble in the Boardroom	179
Chapter 9 The Personal Cloud	200
Chapter 10 Cloud Computing in the Future	215
Cloud Computing Directory	236
Glossary	249
Index	265

## ACKNOWLEDGEMENTS

I would like to thank Leo Hollis for asking me to write this book and then editing it; Kate Pollard for managing the publication process; Richard Rosenfeld for his copy-editing; Mark Daintree for checking parts of the manuscript and helping me introduce others to the cloud; Thomas Chesney, Duncan Shaw and George Kuk for allowing me to bounce ideas around; my parents for their support as I embarked on the journey of yet another book; and Tabitha Browne for being Tabitha Browne.

## PREFACE

Computing is facing another revolution. This time it is called 'cloud computing' and involves accessing software applications, data storage and processing power over the Internet.

Gaining access to computing resources online may not initially seem that radical a proposition. However, cloud computing is already starting to turn the software industry upside down. After all, once people start running programs over the Internet they will have no need to purchase and install them on their own computers. Companies will also not be required to purchase and maintain so much hardware and software if it can simply be rented online. The growth of cloud computing therefore threatens the survival of many software vendors and corporate data centres.

For most people cloud computing will be extremely liberating. This is because it will transform computing into an on-demand utility much like water or electricity. What this means is that, in the near future, computer power will simply be 'on tap' for us to consume as we please. The Internet already provides the infrastructure to allow this to happen. Pioneering cloud computing suppliers are also already peddling their highly cost-effective wares. All that is required is for more people to appreciate the benefits of



not having to invest in computer hardware and software that is rarely fully utilized.

This guide to cloud computing comes in two parts. Part I introduces the basics and explains the advantages of cloud computing, its association with Web 2.0, and the range of online software applications and hardware resources already available. In short, Part I covers what anybody who wants an understanding of cloud computing really has to know about.

Part II provides a broader coverage of the implications of cloud computing, and the different chapters might interest some more than others. The topics covered include security, privacy and reliability, next generation cloud hardware, battles ahead in both the computing industry and many boardrooms, personal clouds, and a glimpse into the future. Come to think of it, every chapter ought to interest everybody!

This book is not a supplier showcase. However, with any technology development as new as cloud computing, it is impossible to explain what is going on without significant reference to actual product offerings. Throughout this book you will therefore find reviews of specific cloud computing services. Many of these are even free to use straight away just by visiting the web addresses included in the text.

Like it or loathe it, cloud computing is already starting to have a significant impact on the personal and corporate computing landscape. Within a decade it is also likely to be the only computing show in town. None of us can therefore ignore cloud computing. Our only real choice is whether we want to be part of the steamroller or part of the road.

Christopher Barnatt  
Lecturer and Futurist  
April 2010

---

### **A Special Note on Web Addresses**

A great many web addresses are included in this book. Most appear in the text without a ‘www’ or ‘http://’ prefix – e.g. docs.google.com – and were checked at the time of publication to work in this format. However, a few web addresses are included with a ‘www’ prefix, for example www.itfarm.co.uk. When this is the case the ‘www’ part does need to be typed into your web browser to access the site. Note that when a web address is included at the end of a sentence, the full stop does not form part of the address. Also note that links to the majority of web resources included in this book are featured in the Cloud Computing Directory on pages 236–48. This directory is also available online from [explainingcomputers.com/cloud](http://explainingcomputers.com/cloud).

Part I

# CLOUD COMPUTING BASICS

# I

## THE RISE OF FLUFFY COMPUTING

This book is being written in the cloud. This means that my word processor is not installed on my computer. The files for each chapter are also not being saved on one of my own hard disks or USB keys. Instead, the program I am using is running somewhere ‘out there’ on the Internet. My files are then being saved to remote online storage. This service is also being provided to me for free.

Exactly where my word processor is running and where my files are I really have no idea. More importantly, I simply don’t care. This is also the whole point. Cloud computing is ‘fluffy’ because the resources being used are irrelevant to the vast majority of users. This is why cloud computing is so scary and resisted in many corporate data centres. It is also why cloud computing is so liberating and powerful for the rest of us.

Just before you think that I have no idea what I am doing, I can reveal that this book is being written in Google Docs. As illustrated in screenshot 1.1, this is an online word processing, spreadsheet, drawing and presentations package now being used by tens of millions of people and an increasing number of companies. In the UK, the *Daily Telegraph* and *Sunday Telegraph* national newspapers are

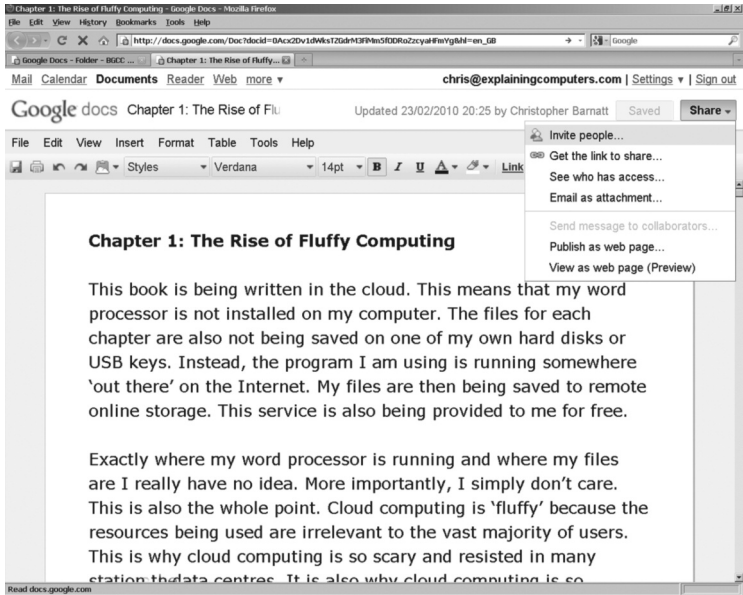
now partially written in Google Docs. This has been occurring since July 2008 when the Telegraph Media Group began a transition from local Microsoft software to the online Google Apps software suite.

Jaguar Land Rover, Rentokil Initial and the District of Columbia government are just three of the far larger organizations that have also adopted Google Apps. Meanwhile universities from Arizona to Delhi, Washington to Notre Dame, and Dublin to Leeds, are in a stampede to sign up to Google Apps Education Edition. Strange and scary as it may initially seem, make no mistake that cloud computing is already entering the mainstream.

This book is your guide to the dawning age of cloud computing in which all manner of computing resources will be accessed over the Internet. It is therefore a book about a technological change as radical as the personal computer revolution. This is, however, not a technical tome. Rather, *A Brief Guide to Cloud Computing* is your route map to the future practices and philosophy of computing and how they will affect us all.

### **So What Exactly is the Cloud?**

You may be wondering what any of this has to do with clouds! The answer is that, for many years, the Internet has been represented on network diagrams by a cloud symbol. When, around 2008, a variety of new services started to emerge that permitted computing resources to be accessed over the Internet, the label 'cloud computing' started to be used as an umbrella term. So does this mean that we ought to be talking about 'Internet computing'? Well, perhaps. However, in the strictest sense, the 'cloud' is a label for online computing resources rather than the entire Internet. The term 'cloud computing' is also useful because it distinguishes the kinds of things we have been doing online



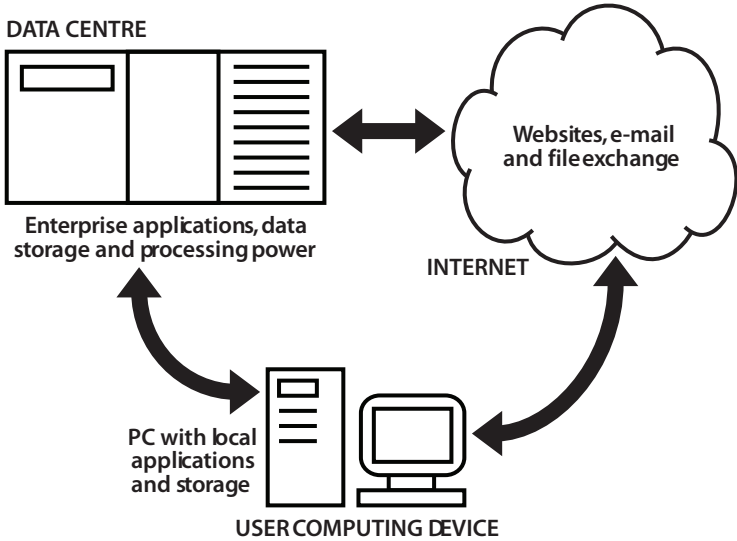
**Screenshot 1.1: Word processing in Google Docs**

for a couple of decades from a totally new age of online software and processing power.

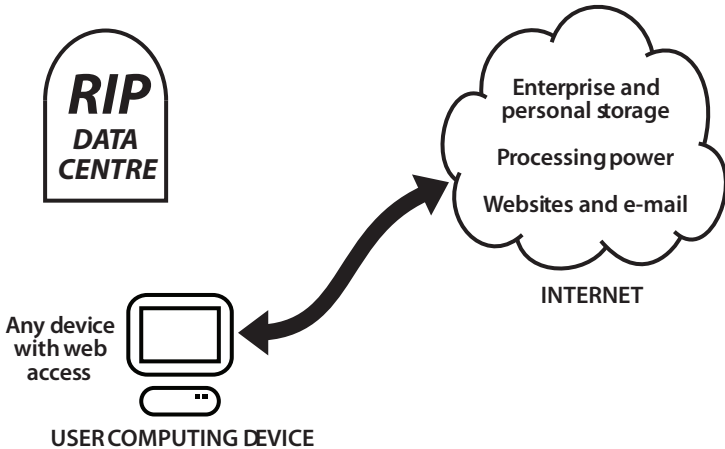
Figure 1.1 illustrates the key differences between traditional and cloud computing. As shown in the top half of the figure, at present local software is installed and data is stored on most personal computers. Most computer users in organizations also access enterprise applications, data storage and processing power from a corporate data centre. The Internet may additionally be used and often relied upon. However, until now, for most people Internet usage has been confined to accessing information from websites and exchanging e-mails and file attachments.

The lower half of figure 1.1 shows the brave new world of cloud computing. Here the corporate data centre has been decommissioned. Also, software applications and data are

## TRADITIONAL COMPUTING MODEL



## CLOUD COMPUTING MODEL



**Figure 1.1: Traditional and cloud computing**

no longer installed and stored on a user's computing device. Rather, enterprise applications, personal applications, data storage and remote processing power are all accessed from the cloud.

The scenarios shown in figure 1.1 indicate the two most extreme positions, with a hybrid model somewhere in between being most likely in the medium-term. Even so, the implications of ceasing to install all applications and store all data on personal computers or in a data centre will be very significant. Not least, as explored in chapter seven, there will be major ramifications for traditional software companies such as Microsoft. As discussed in chapter eight, the implications for those who currently work in company IT departments will also be just as great.

In practical terms, the cloud is made up loads of giant data centres – also known as ‘server farms’ – run by Google, Amazon, Microsoft, IBM, Apple and a host of other traditional and emerging computing giants. While the sorts of online services that can be offered from the cloud are quite varied, they can largely be classified under the two broad headings of ‘online software’ and ‘online hardware’. Lots of information and examples to demonstrate what this means in practice are provided in chapters three and four respectively.

### **Killer Benefits**

So why, you ask, would anybody want to cloud compute? Well, as figure 1.1 indicates, over time cloud computing will remove the need to install and maintain many local or corporate computing resources. In the shorter term, cloud computing also offers a couple of immediate killer benefits.

The first immediate benefit of cloud computing is that data and applications are accessible from any computer on the Internet. This book, for example, is being written in my home office, my university office and on an Eee PC netbook



in all sorts of other interesting locations. By writing the book in a cloud word processor, I never have to worry whether the next computer on which I work will contain the most recent version of each chapter.

The second immediate advantage is that cloud computing is collaborative. I happen to be writing this book all by myself. However, if I had a co-author my cloud word processor would really come into its own, with each of us always being able to work on the latest version of any chapter, and even at the same time. Just in case this is getting you excited, I should point out that lots of information on collaborative working with specific cloud computing applications is provided in chapter three.

Text documents are also far from the only forms of data starting to take up residence in the cloud. Photographs, music, e-mail, personnel files, accounts, videos, books and more are already being manipulated over the Internet as cloud computing advances. As will be revealed throughout this book, whatever you do or want to do with a computer it is likely that there is already a suitable cloud computing application.

Alongside any-device access and collaborative working, cloud computing also has a number of wider benefits. As we will explore shortly, these include cloud computing being cheaper and more environmentally friendly than traditional computing. Online processing power and cloud data storage are also essential for the development of new computing applications like augmented reality.

### **Security, Privacy and Reliability**

Before we get too evangelical, it has to be said that cloud computing does have its fair share of potential drawbacks. Most obviously, there is the dependence of any cloud computing application on a reliable Internet connection.

While for some this may present a problem, it nevertheless remains the case that a great many homes and most organizations do now have a reliable and high-speed Internet connection. Over the next few years, those that do not have one can probably expect connectivity to improve. Which brings us to most people's far bigger worry – security and privacy.

Many cloud computing concerns about security are likely to be perceptual rather than real. Anybody who sends an e-mail is already trusting the confidentiality of their material to their Internet service provider, not to mention all those companies who run the Internet infrastructure over which their message will travel. This means that any document written in, say, Microsoft Word and sent as an e-mail attachment is already no more secure and confidential than a document written in Google Docs or any other online word processor. To a large extent, when it comes to security and privacy, cloud computing developments simply highlight how trusting of the Internet we have already perhaps blindly become.

Security issues may also at least in part be accounted for contractually. For example, when in October 2009 Los Angeles City Council decided to move its 30,000 employees to Google's Government Cloud services it got a powerful security-breach penalty clause added to the contract.

For individuals and those working in smaller organizations, cloud computing is also likely to reduce their risk of security disasters. As YouTube user AnswerFortyTwoX posted in response to a security-conscious viewer of one of my early cloud computing videos:

From the perspective of a hacker . . . it is infinitely easier for me to break through the meagre security on a personal

computer than it is for me to take on a Google server. In a way, your documents are safer.

As further explored in chapter five, the benefits of cloud computing already outweigh the potential security, privacy and reliability drawbacks in the vast majority of cases. Stand-alone desktop and data centre computing may continue to prove essential for certain specialist types of users and applications. However, for the majority of us most of the time, cloud computing is the future.

### **Key Cloud Computing Characteristics**

Given that it is fluffy by nature, it is quite difficult to provide a totally precise and widely accepted definition of cloud computing. This is partly because different technical computing specialists will opt for a different emphasis in their definitions than most end-users. After all, while for most people cloud computing is about simplifying things and masking complexity, for computing specialists it already involves new technologies and new career structures. You should therefore not expect a data centre manager, a programmer, a typical business computer user or a private individual to define cloud computing in precisely the same manner. This said, definitions do remain important.

I have already explained that cloud computing involves software applications, processing power and data storage being accessed online. Building on this, from most people's perspective we can also state that cloud computing is where dynamically scalable, device-independent and task-centric computing resources are obtained over the Internet, with any charges (where payable) being on a per-usage basis.

While the above is a somewhat long and involved definition, it does at least bring together in one sentence

cloud computing's four key characteristics. I will also now explain these in more depth.

*Cloud computing is dynamically scalable*

Cloud computing is dynamically scalable because users only ever have to consume the amount that they actually want. Just as we are used to drawing as much or as little electricity as we need from the power grid, so any user can draw as many or as few computing resources from the cloud as they require at any particular moment. This means that individuals and organizations will no longer have to invest in computing resources that often sit idle. Nor will they have to wait in frustration for complex computing tasks to be completed due to a lack of available processing power.

One of the first major suppliers of dynamically scalable cloud computing resources was Amazon. Yes, the company that started out selling books online – though which in reality has always been a logistics business – is now selling computer-processing power by the hour. And I mean that quite literally. Via a service called Elastic Compute Cloud, or 'EC2', Amazon sells cloud computing processing capacity in what are termed 'instances'. As Amazon explain, EC2:

... allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change.

At the time of writing, the smallest standard Amazon EC2 instance is a 1.2GHz 32-bit virtual processor core with 1.7Gb of memory and 160Gb of storage. This can be

provided running either Windows or Linux for around 5p an hour.

If the last two paragraphs sound like a load of technical gibberish, then please don't panic and just read on! The key point right now is that cloud computing developments like EC2 allow anybody to purchase any capacity of computer power they want both cheaply and on an hourly basis. Amazon EC2, its competitors, and their business implications, are discussed in chapter four.

### *Cloud computing is device-independent*

I have already noted that cloud computing resources can be accessed from any computer on the Internet. It is, however, worth stressing that this does not just mean any computer, but any *kind* of computer. Provided that it has an Internet connection and a web browser, it really does not matter if the computer is a traditional desktop or laptop PC, or even a netbook, tablet, smartphone, e-book reader, surface computer, ambient device or any of the other new computing appliances discussed in chapter six.

Again, this is a radically new development. While over the past decade it has become easy to exchange data between different computers, this still requires the right software to be installed. For example, if a presentation is created in Microsoft PowerPoint and sent to somebody as an e-mail attachment, the recipient needs to have Microsoft PowerPoint on their computer to open and edit the file in a completely compatible manner. However, if the presentation is created in Google Docs then it can be opened and edited on any kind of computer with an Internet connection and a web browser. This even includes viewing and editing the presentation on some Internet-enabled mobile phones.

Device independency will be of increasing importance as more of us access the Internet using smartphones, tablets and

other pocket and handheld devices. For many years, manufacturers have been building cut-down versions of Word, Excel and PowerPoint into mobile devices to allow us to access and work on documents on the move. However, this has very much been a second-best solution, with a lack of total compatibility between the desktop and mobile software versions. For anybody who wants to work on documents on a phone using exactly the same software that they use on a laptop or desktop, cloud computing is therefore already a Godsend.

### *Cloud computing is task-centric*

Cloud computing is task-centric because the usage model is based entirely around what users want to achieve, rather than any particular software, hardware or network infrastructure. Users do not have to purchase or install anything before using a cloud computing resource. Nor do they have to maintain or pay for anything during periods in which no resources are being used.

As David Malcolm Surgient of ZDNet explains, cloud computing ‘abstracts away’ the traditional, infrastructure heavy view of pre-defined computing environments still maintained by most organizations. As David nicely puts it:

In most cases, users of the cloud generally want to run some business service or application for a specific, timely purpose; they don’t want to get bogged down in the system and network administration of the environment. They would prefer to quickly and easily access a dedicated instance of an application or service.

For many years Microsoft marketed its wares under the slogan, ‘Where do you want to go today?’ The answer was presumably that we wanted to go to a computer and install

some Microsoft software that would, in turn, allow us to get on with what we actually wanted to do. Cloud computing cuts out the ‘going to a computer and installing something’ step. Rather, if you want somebody to see or edit a document you just send them a link.

The above means that cloud computing will allow us to simply get on with those many activities that involve a computer. Nobody today settles down to use a pencil. In contrast, lots of people do still consciously sit down to use a computer. Cloud developments may, however, start to catalyze a mentality shift from tool-in-hand to task-at-hand computer application.

### *Cloud computing has no fixed costs*

In business, a fixed cost is something that has to be paid regardless of the number of people who use a certain facility or a company’s level of production. A variable cost is different because it changes according to the number of people involved and output levels. For example, the annual cost of renting a factory is likely to be fixed. In contrast, the cost of staffing a factory, and of the raw materials it requires, will vary according to how much is produced.

Traditionally, computing has involved substantial fixed costs. Most significantly, these have included the cost of building, equipping and maintaining data centres. However, because cloud computing is dynamically scalable and task-centric, for most users it has no fixed costs. Rather, all costs are on a per-usage or variable basis. As demonstrated by the example of Amazon EC2, processing power can already be purchased from the cloud by the hour.

Software applications purchased from the cloud similarly incur only variable costs. Where provision is not free, charges are typically based on the number of people using an application each month, or the number of records or projects

being worked on. So, for example, the online database software Zoho Creator is free for two users using up to three database applications. It then costs \$15 per month for three users, \$25 a month for up to five users, \$45 a month for up to ten users, and so on – with the number of databases and records permitted also rising as the charge increases. As with Amazon EC2, the cost structure for Zoho Creator is also entirely elastic. This means that users are able to upgrade or downgrade their requirements at any time. A company with seasonal business can therefore license the software for a great many users only at their busy times of year.

The fact that cloud computing has only variable costs is extremely important for small companies. This is because small companies have not, until now, been able to afford the kinds of sophisticated business applications available to large corporations. However, because cloud computing suppliers such as Clarizen, Employease, Netsuite, Salesforce and Zoho do not charge an initial fixed-cost outlay, they are now levelling the software-access playing field. Indeed, as explained in chapter three, the latest types of human resource, project management, customer relationship management (CRM) and other applications can now be accessed from the cloud by any business, large or small.

### **The Only Show in Town**

All of the basics of cloud computing have now been outlined. You may well think that everything expected of a first chapter has therefore been covered and that it is high time for chapter two. After all, you now know what cloud computing is, how it differs from current computing practice, what its key benefits and drawbacks may be, and hence why individuals and organizations may opt to cloud compute.

The above noted, what I have not familiarized you with yet are the broader reasons why we may all soon have to



cloud compute. In other words, it is now time to explain why cloud computing will soon be the only mainstream computing show in town.

In October 2009, technology analysts Gartner cited cloud computing as the most strategic technology development for 2010, and hence a trend that no business can ignore. Their logic for this was complex and in part linked to everything discussed thus far. However, I would propose that cloud computing will become inevitable for the following three reasons:

- Cloud computing will be essential to remain competitive
- Cloud computing will be essential to be green, and
- Cloud computing will be essential for next-generation applications

I will now justify each of these rather bold claims.

### **The Competitive Cloud**

We have already seen how cloud computing is dynamically scalable, task-centric and charged on a per-usage basis. As a consequence, cloud computing is likely to be more cost effective than current computing arrangements for most organizations. Indeed, the cost savings already speak for themselves.

The Telegraph Media Group, for example, expects to reduce its software costs by 80 per cent over three years following its switch from local Microsoft software to Google Apps. Some suppliers even claim cost savings of as much as 90 per cent. While such figures do need to be treated with some scepticism, there are nevertheless a great many companies – and especially small companies – with very positive stories to tell. US biotechnology pioneer Genentech,

for example, claims savings of ‘millions of dollars’ from using Google and a range of other cloud computing providers. Not least this is because Genentech has avoided the construction of a new \$10 million data centre.

In fact, the cost savings are potentially so great that in an increasing number of instances, cloud computing is likely to become the only competitively acceptable computing option. There is also a strong historical precedent for this. As explained by Nicholas Carr in his excellent book *The Big Switch*, around a century ago using electricity from a national grid – rather than generating it within a business – became a competitive necessity. This was simply because big electricity suppliers could leverage economies of scale that were unavailable to their customers.

Today the cloud is becoming the centralized power plant of the Information Age. In tandem, the mistrust being voiced by many in computing who do not want to use this new central utility service is also nothing new. As Carr notes, in 1900 there were about 3,600 public electricity-generating plants in the US, and yet around 50,000 private plants fuelling the energy needs of individual companies. Firms may have switched from water power to steam power to electricity. However, they still took some convincing to switch to an external electricity supplier. However, once convinced of the competitive necessity to switch they did not hesitate. For as Carr puts it:

By avoiding the purchase of pricey equipment, they reduced their own fixed costs and freed up capital for more productive purposes. They were also able to trim their corporate staffs, temper the risk of technological obsolescence and malfunction, and relieve their managers of a major distraction.

The above are today also the precise reasons that most companies will shift from internal to cloud computing. Cloud vendors are already able to reap very significant economies of scale by optimizing the use, support and upgrading of their infrastructure. For most companies, running computing internally will therefore soon be recognized as too costly a strategy except in very specialist situations. Developing a reliance on central resources is also not something to be feared. It was, after all, the centralization of food and water supplies that created civilization and permitted the rise and survival of the modern city.

As we go about our daily lives, few of us even consider our reliance on a wide variety of utility services and infrastructures beyond our individual control. If we do ever give them a thought, then I would suggest it is usually with thanks and relief that we do not have to individually forage and fight to gather the sustenance necessary to keep us alive. In a decade the same will apply when it comes to cloud computing – with even the idea that a typical organization should run its own data centre being unthinkable.

### **The Green Cloud**

The second major factor that will force most of us to cloud compute will be the legal and moral requirement to use less energy and waste fewer resources. In other words, cloud computing will be essential to be green.

One reason why cloud computing is more environmentally friendly than in-house or desktop computing is because large external vendors can run their infrastructure highly efficiently. Cloud computing suppliers with hundreds of thousands of servers will rarely, if ever, need to have many machines powered while standing idle. In contrast, in most corporate data centres – let alone in smaller organizations – a great many servers and indeed desktop computers are kept

running at well below full capacity most of the time. This means that electricity is constantly being wasted to power and cool under-utilized hardware.

Cloud computing vendors are already proclaiming their green credentials. Netsuite, for example, advertises that in 2008 its customers saved \$61m in energy bills by using its web-based customer relationship management, accounting and other cloud applications. Whole countries are also waking up to the environmental potential of cloud computing. The Icelandic government, for example, has now recognized that in terms of carbon footprint, Iceland is one of the best places in the world to host large numbers of computer servers. With approximately half of the energy used by a large data centre going into cooling, putting cloud server farms in very cold countries also makes plenty of sense. This is because the cold air required to cool the computers can be drawn in from outside the data centre, rather than being chilled with electrical air-conditioning units. Iceland can also power its data centres with geothermal rather than fossil fuel electricity.

In preparation for its anticipated cloud computing ‘cold rush’, Iceland is laying high-capacity, fibre optic cables to connect the country with North America and Europe. Just outside the Icelandic capital Reykjavik, work is also well under way on one of the first of Iceland’s massive cloud computing centres.

As Iceland is well aware, the world’s data centres already have about the same carbon footprint as the airline industry. Computing is therefore going to come under significant pressure to be more environmentally friendly. The adoption of cloud computing can also help us achieve this in three important ways.

Firstly, as noted above, a move from in-house to cloud computing will allow large vendors to optimize the power

usage of the servers on which we will all rely. Local server capacity will therefore no longer be powered but idle. We just have the significant hurdle of getting IT managers to stop server-hugging (which I address in chapter eight).

Secondly, if cloud computing is adopted then many individuals will be able to work on low-power computing devices that draw most of their computational capacity from the cloud. This means that over-specified personal computers will not be sitting on desks and drawing loads of power while they wait for their users to actually do something. Once again, the savings can be significant. The Canadian vendor ThinDesk, for example, has achieved energy reduction savings of up to 80 per cent for small and medium-size firms who have switched to low-power computers that rely on cloud services delivered from its TELUS data centre.

Thirdly, while a mass migration to the cloud will make computing less of an environmental problem, innovative cloud computing adoption may also make computing part of a broader green solution. Collaborative Internet tools are already enabling more people to reduce their business travel by teleworking from home at least some of the time. The more we cloud compute, the more we will therefore be able to take planes out of the sky and vehicles off the road. These broader implications of cloud computing are explored in more depth in chapter ten.

### **The Next Generation Cloud**

The final factor that will make cloud computing essential is that many next-generation computer applications will only work in the cloud. Local software and data inevitably constrain collaboration and the anytime, anyplace, anywhere use of information resources. This means that we will be prevented from obtaining the benefits of new

‘crowdsourcing’ developments unless we cloud compute.

Crowdsourcing is where the Internet is used to help generate value from the activities of a great many people. Today, crowdsourcing mainly involves lots of people working together collaboratively to tackle a problem that in the past would have been left to just one individual or a small group. As will be discussed in chapter ten, we are therefore starting to see an increasing number of so-termed ‘open source’ developments where all of the involved intellectual property is created and shared online for mutual benefit. Already crowdsourced products and services in use or under development include computer software, robots, 3D printers, prosthetic limbs and electric cars.

An increasing uptake of cloud computing will make it easier for individuals to consciously work together on crowdsourcing projects. In addition, there is also a significant potential to crowdsource data not just from people, but also from things.

As will be discussed in chapter six, more and more everyday items – including fridges and even clothing – are now being given their own Internet connection. Using data from cameras and other sensors, smart cloud computing applications are also starting to recognize and monitor objects that do not have their own connection to the web. As these trends take hold and more objects directly or indirectly go online, we will start to witness the development of some quite innovative cloud computing applications.

For example, within a few years augmented reality will be commonplace. This will allow us to hold up a smartphone or other mobile device and see real-time data overlaid on the camera image shown on its display. When viewed on screen, buildings in the street or any product in a shop will be clickable if additional information is required. To allow this to happen, image recognition cloud applications will need

to identify objects in view and associate them with GPS coordinates from our handsets and massive online databases. Sound unlikely? Well, as discussed in chapter six, first-generation augmented reality browsers are already on the market and running on some smartphones.

Pretty soon there will be so many cameras, microphones and other sensors online – at least in public spaces – that a great many objects will start to cast a constant ‘data shadow’. While this may raise concerns, it will also allow us to reap crowdsourcing benefits similar to those of online social networking. Satellite navigation systems, for example, could advise on routes based not only on internal maps, but also the position and predicted intent of all other vehicles on the road. However, this will only happen if a great deal of data is pooled and shared in the cloud rather than being held and processed on local computing devices.

Future developments in artificial intelligence (AI) will probably also depend on crowdsourced data. Programming a mobile phone or a robot to recognize everything in view is likely to remain very difficult if internal data and processing power have to be relied on. On the other hand, a phone or robot with access to cloud resources including video feeds from other nearby cameras will be in a far better position to make sense of the world around it. This means that for computers to be usefully smart they will require access to information from our immediate environment that can only be crowdsourced from the cloud. In turn, along with augmented reality, the development of artificial intelligence is likely to be a very strong reason for the mass adoption of cloud computing.

## Coming Full Circle?

Whenever I give a lecture or run a workshop on cloud computing, at least one person questions whether it is all basically just a return to computing's early days. After all, as they argue, in the 1960s and 1970s most computer users worked on 'dumb terminals' that were totally reliant on a connection to a mainframe or minicomputer housed in a remote data centre. Having cited this historical fact, they then argue that cloud computing is actually a very old idea and potentially a step backwards.

So is cloud computing bringing us full circle? Well the answer to this is a very definitive 'yes' and 'no'. On the 'yes' side, cloud computing will make most of the computing devices we use dependent on remote resources most of the time. This is, however, already commonplace and not something we should fear.

Mobile phones, for example, are computing devices that are pretty much useless without a connection to a 'cloud' infrastructure and this is readily accepted. Digital televisions are also complex computing devices that are totally dependent on an external data feed and this similarly does not make people worry.

While most people may not think of mobile phones and digital televisions as computers, these devices help us realize why cloud computing is quite distinct from the mainframe-and-dumb-terminal era. The key difference is that yesterday's dumb terminals were dependent on a very specific, small-scale computing infrastructure. In contrast, most cloud computing appliances will not depend on any specific local server or mainframe, or indeed on any dedicated local computing resource at all.

I have already explained how cloud computing is dynamic, device-independent and task-centric. We could also add that cloud computing is resilient. Granted,



cloud computing will make us ever more dependent on remote computing resources. However, we must recognize that such remote resources are likely to be far better managed and less prone to failure than those in most internal IT departments. This is simply because the likes of Google, Amazon and IBM are able to specialize in running a great many data centres – what we can think of as large wisps of the cloud – and that scale plays a major role in the delivery of service continuity. If, for example, you were dependent on a constant electricity supply to stay alive, would you rather be connected to your own generator or an electricity grid fed by scores of power stations?

In late 2009, Google's Chief Executive Eric Schmidt stated that the cloud computing revolution will be bigger than the advent of personal computing. While he may be somewhat biased, I also tend to agree with him. The early years of the personal computing revolution may have briefly empowered individuals and freed many from the shackles of their IT department. Unfortunately, the network and Internet developments of the nineties and noughties then allowed IT staff to claw back into the corporate data centre a great deal of central control.

Within organizations, one of the great promises of the cloud is that it will turn computing into such a utility activity that it will never again be able to be centralized in-house. In the home, the cloud revolution is also destined to turn computers into on-off appliances that simply just work. To the future relief of so many, the days of ludicrously complex, misbehaving operating systems and corrupt or incompatible software will come to an end.

The IT function in many companies is today at a crossroads in the face of cloud computing developments that threaten to give users the kind of flexibility once briefly promised by personal computing and then cruelly snatched

---

away. Battle lines are therefore being drawn not just in the computer industry, but within a great many companies. Small businesses are also being freed to compete on a level computing playing field, while environmental campaigners are becoming wise to online opportunities that can make computing more green.

The remaining chapters of this book will guide you through the aspects of cloud computing that everybody needs to know about and their associated implications. The content included is as current as it can be. However, exactly how the cloud computing revolution will unfold is inevitably at present unclear. This said, I think we can already be fairly certain that the computing landscape that the advancing cloud will leave in its wake will be very different from the one we know today.