



The Fog over the Grimpen Mire: Cloud Computing and the Law

Miranda Mowbray

HP Laboratories
HPL-2009-99

Keyword(s):

cloud computing, legal

Abstract:

This paper is about legal questions connected with cloud computing, the business trend in which computation is carried out on behalf of a user on remote machines, using software accessed through the Internet. The user may not know where these machines are; they are "somewhere in the cloud". Some of these legal issues will be resolved by standard agreements between buyers and vendors. I will give some examples from current agreements from prominent cloud service providers. Other issues will probably end up in court. It makes sense to consider these questions now, before they become urgent.



Volume 6, Issue 1, April 2009

The Fog over the Grimpen Mire: Cloud Computing and the Law

*Miranda Mowbray**

Abstract

This paper is about legal questions connected with cloud computing, the business trend in which computation is carried out on behalf of a user on remote machines, using software accessed through the Internet. The user may not know where these machines are; they are “somewhere in the cloud”. Some of these legal issues will be resolved by standard agreements between buyers and vendors. I will give some examples from current agreements from prominent cloud service providers. Other issues will probably end up in court. It makes sense to consider these questions now, before they become urgent.

DOI: 10.2966/scrip.x.x



© Hewlett Packard Development Company, L.P. 2009. This work is licensed under a [Creative Commons Licence](#). Please click on the link to read the terms and conditions.

* Technical Contributor, HP Labs Bristol, UK. Not a lawyer.

1. Introduction

In this paper I will point out areas of potential legal disputes in cloud computing. Some of these will be resolved by technological means, and some by standard agreements between buyers, vendors and subcontractors or by standard industry practices. Others will end up in court. It makes sense to think about the potential legal issues now.

In a 2008 article the journalist Bill Thompson, writing about the fact that cloud computing actually takes place not in an immaterial cyberspace but in physical computers in the real world, said:

*In the real world national borders, commercial rivalries and political imperatives all come into play, turning the cloud into a miasma as heavy with menace as the fog over the Grimpen Mire that concealed the Hound of the Baskervilles in Arthur Conan Doyle's story.*¹

The menaces described in Bill Thompson's article were unreliability of cloud services, including possible inaccessibility of data, and access to data by foreign governments. In the rest of this article I will discuss these and other cloud computing topics for which there are foggy legal issues, including subcontracting, rights to data use, lock-in to a service provider, and security loopholes.

2. Cloud computing: what is it?

Rich Zippel of Sun Microsystems has described cloud computing as "the hottest, and certainly the most abused, buzzword in computing today." Gartner Groups identified it as entering the peak phase of the hype cycle in July 2008.² There is disagreement about the precise definition, but essentially cloud computing means remote computing with software accessed through the internet. This software is usually paid for according to the amount that it is used; in some cases there is also a modest subscription fee, and in others the software is free for use and paid for with advertising. Cloud computing is part of a general architectural trend in the computer industry, moving from users doing computing on their own hardware using copies of software that they own, to users doing computing on other peoples' machines somewhere in the cloud, using software that they rent.

Cloud computing is related to (but not identical to) software as a service, grid computing, Web 2.0, on-demand computing, utility computing, Internet service platforms, ASPs, and everything as a service. These are buzzwords that were previously popular.

¹ B Thompson, "Storm Warning for Cloud Computing" (2008) available at <http://news.bbc.co.uk/2/hi/technology/7421099.stm> (accessed 23 Mar 2009).

² Gartner, Inc, "Gartner Highlights 27 Technologies in the 2008 Hype Cycle for Emerging Technologies" (2008) available at <http://www.gartner.com/it/page.jsp?id=739613> (accessed 23 Mar 2009).

2.1 Examples

Although there is not a consensus definition of cloud computing, there is general agreement that certain examples are indeed examples of the cloud computing ecosystem. These include Amazon Web Services™, which is a set of services that enable customers to use Amazon's computing infrastructure, including a computing power service, a storage service, and e-commerce software provided as a service; Google Apps, which is a platform on which companies they can run their own applications in the cloud, and also includes some useful functionality (such as MapReduce) that can be used by the applications; and Salesforce.com, which sells customer relationship management software as a service.

A theoretical example for which the use of cloud computing would be advantageous is an online Easter egg business. A business that sells Easter eggs over the Internet can expect a very large number of orders shortly before Easter, and hardly any orders the rest of the year. Without cloud computing, the business would have to buy enough servers to meet its needs in the Easter period, and these servers would be sitting idle for the most of the year. With cloud computing, the business can pay to use a cloud service provider's servers during the Easter period, and does not have to pay for the use of these servers outside this period.

A real example of a use of cloud computing is the use of storage and computing power from Amazon Web Services to convert 11 million public domain articles in the New York Times archives from scanned images into PDF format, to make them easy to read over the Internet. It took just under 24 hours.³ One of the articles, dating from 1902,⁴ is headlined

*THE HOUND OF THE BASKERVILLES.
DID IT SLAY SIR CHARLES?
Sherlock Holmes Called Upon to Solve the Tragic Mystery.*

Although it is formatted to look like a news article, is in fact an advertisement in the book section for Arthur Conan Doyle's *The Hound of the Baskervilles*, which had just been published in the US. The foot of the article says that the story is continued on page 1 of the book.

³ D Gottfrid, "Self-service, Prorated Super Computing Fun!" (1 Nov 2007) available at <http://open.blogs.nytimes.com/2007/11/01/self-service-prorated-super-computing-fun/> (accessed 23 Mar 2009).

⁴ The New York Times, "The Hound of the Baskervilles. Did it Slay Sir Charles?" (29 Mar 1902) available at <http://query.nytimes.com/mem/archive-free/pdf?res=9502E0D9103BE733A2575AC2A9659C946397D6CF> (accessed 23 Mar 2009).

3. *Breakfast in Paris, dinner in London, data in Washington*

*"My word, it does not seem a very cheerful place," said the detective with a shiver, glancing round him at the gloomy slopes of the hill and at the huge lake of fog which lay over the Grimpen Mire.*⁵

Some organisations are not very cheerful about the possibility that cloud computing services that process or store of their data outside its country of origin may enable foreign governments to access their data.

French government officials have been forbidden to use Blackberry email devices,⁶ because Blackberries send and receive email using a small number of servers in the US and UK, and the French security service feared that this might cause a threat to national security because of a risk of data interception. Research in Motion, the maker of Blackberries, has denied that there is any risk. Ironically, Research in Motion is a Canadian company, and the Canadian provincial governments of British Columbia and Nova Scotia require public bodies and their Internet service providers to ensure that personal information under their control is stored and accessed only in Canada, unless specified exceptions apply.⁷

Several major cloud computing services are based in the US, where their electronic records may be subpoena'd under the *USA PATRIOT Act 2001* without notification of the data owners.⁸ The UK *Regulation of Investigatory Powers Act 2000* allows a wide range of UK public servants to obtain a warrant to access data stored on computers in the UK if this is necessary, for example, for "the interests of the economic well-being of the United Kingdom" or to prevent or investigate a crime.⁹ It has however been argued¹⁰ that concentrating on laws like these is missing the point, because in many (perhaps most) countries, if the government thinks that it is important for the national interest to look at some data stored on computers within its territory, it will to do so whether or not there is not a national law explicitly authorizing it to do this. Individuals and organisations who want to keep a set of data confidential from some country's government would be wise to avoid using cloud computing services that process or store data in that country.

One way of protecting the privacy of data that is stored remotely is to encrypt it. Governments may have enough processing power to break some forms of encryption, or legal powers to demand the decryption of data, but non-government actors usually do not. A particular issue with cloud computing is that if data undergoes nontrivial

⁵ A Conan Doyle, *The Hound of the Baskervilles* (London: Penguin Books, Red Classics edition, 2007), at 187.

⁶ J Follorou, "La plainte du Blackberry dans les ministères" (2007), *Le Monde*, 20 Jun 2007.

⁷ D Fraser, "The Canadian Response to the USA Patriot Act" (2007) 5 *IEEE Security and Privacy* no.5, 66-68, available at <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=04336282> (accessed 23 Mar 2009).

⁸ *Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism (USA PATRIOT ACT) Act 2001*, Title V, s 505.

⁹ *Regulation of Investigatory Powers Act 2000*, Part II, s 28.

¹⁰ J Madelin, "Transformational Change – large converged networks come to life" (2008), Adaptive and Resilient Computing workshop, London, 30 Sept 2008.

processing on a remote machine, rather than just being transferred and stored, then (with a few currently rare exceptions) the data has to be unencrypted at the point of processing. So data that is processed using cloud services will usually be present unencrypted in a machine somewhere in the cloud. This limits the types of processing in the cloud that are legally permissible for types of data that are subject to certain laws and regulatory regimes, for example personal medical data that is subject to the USA *Health Insurance Portability and Accountability Act 1996* — although some US citizens opt to share their own medical data via the cloud using the Google Health cloud service.

The issue of potential access to data by governments is part of a wider issue, which is that the use of services based in other countries may result in customers being affected by laws of those countries. An experience of Steve Marshall, an English travel agent living in Spain, is an example of this. He sells holidays to European customers. Some of these holidays are trips to Cuba. One day about eighty of his company's Cuba-related web sites stopped working. What had happened was that he used a US-based domain name registration service, and the registrar had disabled the sites, without notifying him, as a result of the rules forbidding US companies to do indirect business with Cuba.¹¹

A company using cloud computing may well find itself using hardware and software that are in different countries from its own physical location and the physical location of its customers. We are likely to see legal disputes arising from geopolitical and jurisdictional issues to do with these cases. There is also a potential market for geographically-restricted cloud computing services, where part of the service offering is an assurance that for instance the service will only process data in Europe, so as to conform with European privacy laws, or will only store data in Switzerland, so as to conform with Swiss data protection laws. Indeed, Amazon's computing and storage services have an option for processing and storing in Europe rather than the US. Amazon added this option partly to reduce latency for European customers, but also because of data protection issues.¹²

4. (Re)liability

*I have said that over the great Grimpen Mire there hung a dense, white fog. It was drifting slowly in our direction and banked itself up like a wall on that side of us, low, but thick and well defined ... Holmes's face was turned towards it, and he muttered impatiently as he watched its sluggish drift.*¹³

What guarantee, if any, is there that a cloud computing service will not be too sluggish?

When the Twitter service is temporarily unavailable, a cute cartoon whale called the Fail Whale appears on users' screens. This has happened often enough that the Fail

¹¹ A Liptak, "A Wave of the Watch List, and Speech Disappears" *New York Times*, 4 Mar 2008 available at <http://www.nytimes.com/2008/03/04/us/04bar.html> (accessed 23 Mar 2009).

¹² J Barr, Cloud Computing session, BarCamp Brighton 2, Falmouth, 13-15 Mar 2008.

¹³ A Conan Doyle, note 5, at 189.

Whale now has his own fan club¹⁴ and Facebook group.¹⁵ Yiying Lu, the artist who drew the Fail Whale, later created a girlfriend for him in response to fans' requests. Melik Yuksel responded "If they have kids and spread over the net, I'm blaming you. =)"¹⁶

If you use cloud computing to provide the computing power for your Easter egg business, you would probably like some sort of guarantee that the service will be accurate and available and will not lose your data – or else that if the Fail Whale's kids spread to the cloud computing service, you will get some money back in compensation for the service downtime. The current user agreements for the most prominent cloud computing services give no such assurance.

4.1 Some current user agreements

According to the terms of service for Google Apps¹⁷, the services might be interrupted, untimely, insecure, full of errors, give inaccurate or untimely results, and have low quality, but Google and partners would have no liability to you.

GOOGLE AND PARTNERS DO NOT WARRANT THAT (i) GOOGLE SERVICES WILL MEET YOUR REQUIREMENTS, (ii) GOOGLE SERVICES WILL BE UNINTERRUPTED, TIMELY, SECURE, OR ERROR-FREE, (iii) THE RESULTS THAT MAY BE OBTAINED FROM THE USE OF GOOGLE SERVICES WILL BE ACCURATE OR RELIABLE, (iv) THE QUALITY OF ANY PRODUCTS, SERVICES, INFORMATION, OR OTHER MATERIAL PURCHASED OR OBTAINED BY YOU THROUGH GOOGLE SERVICES WILL MEET YOUR EXPECTATIONS

The terms of service go on to say that you expressly agree that Google and partners shall not be liable to you for "any direct, indirect, incidental, special, consequential or exemplary damages" resulting from any matter relating to Google Services.

Amazon Web Services state in their general customer agreement that "We will have no liability to you for any unauthorized access or use, corruption, deletion, destruction or loss of Your Content or Applications."¹⁸

Even with disclaimers as sweeping as these, service providers may still have some liabilities, although there is not a simple rule to determine whether or not they do in a

¹⁴ Fail Whale Fan Club, "The Fail Whale Fan Club: celebrating Twitter and our favourite error page cetacean" (2008-2009) available at <http://failwhale.com/> (accessed 23 Mar 2009).

¹⁵ Facebook, "FailWhale" (2009) available at <http://www.facebook.com/pages/FailWhale/64467830480> (accessed 23 Mar 2009).

¹⁶ Yahoo!, "flickr: Eve Whale" (2009) available at http://www.flickr.com/photos/crispy_chips/2622841707/#comment72157606109949329 (accessed 23 Mar 2009).

¹⁷ Google, "Welcome to Google Apps: Google terms of service" (2009) available at http://www.google.com/apps/intl/en/terms/user_terms.html (accessed 23 Mar 2009).

¹⁸ Amazon Web Services LLC, "AWS Customer Agreement" (2009) available at <http://aws.amazon.com/agreement/> (accessed 23 Mar 2009).

particular case. For instance for UK customers, UK consumer protection laws may still apply.¹⁹ Under the *United Nations Convention on Contracts for the International Sale of Goods 1980*, sellers have a duty to deliver goods that are “fit for the purposes for which goods of the same description would ordinarily be used”, unless the parties have agreed otherwise.²⁰ A very unreliable computing service might be unfit for purpose. The Amazon Web Services customer agreement states however that “The parties expressly exclude application of the United Nations Convention for the International Sale of Goods to this Agreement.”

The master subscription agreement for Salesforce.com used to have a general disclaimer of warranties similar to the one in the Google Apps agreement. The current version of the Salesforce.com agreement²¹ does include some warranties.

We warrant that (i) the Services shall perform materially in accordance with the User Guide, and (ii) subject to Section 5.3 (Google Services), the functionality of the Services will not be materially decreased during a subscription term. For any breach of either such warranty, Your exclusive remedy shall be as provided in Section 12.3 (Termination for Cause) and Section 12.4 (Refund or Payment upon Termination) below.

Sections 12.3 and 12.4 say that a customer can terminate the agreement “for cause” if she has given 30 days’ written notice to Salesforce.com of a material breach of the agreement, and the breach remains uncured at the end of the 30 days; and that in that case, Salesforce.com will refund any prepaid fees covering the remainder of the term of all Order Forms after the effective date of termination.

In no event shall any termination relieve You of the obligation to pay any fees payable to Us for the period prior to the effective date of termination.

Thus if Salesforce.com breaches either of these warranties, all that the exclusive remedy offers to the customer is a refund of fees that she has paid in advance for future services that she will not use. Salesforce.com also warrants that it has the legal power to enter into the agreement, and that it will not transmit any malicious code to the customer unless the customer transmitted the code to Salesforce.com first, but disclaims all other warranties.

4.2 Software sales vs. network outsourcing

The customer agreement for Google Apps quoted in the previous subsection takes a familiar form: it gives a blanket warranty disclaimer and a blanket liability disclaimer, just like the shrink-wrapped licences with which software is sold on physical media. However, it is not clear that physical software sales give the best analogy for cloud

¹⁹ P Massey, “Privacy, Regulation, Security and Cloud Computing”, *Powered By Cloud*, London, 2-3 Feb 2009.

²⁰ *United Nations Convention on Contracts for the International Sale of Goods* (1980), art 35.

²¹ Salesforce.com, “Master Subscription Agreement” (2000-2009) available at <http://www.salesforce.com/company/msa.jsp> (accessed 23 Mar 2009)

computing. When software is sold on a physical medium, the purchaser owns the software and only pays once to use it as many times as she wishes.

An alternative analogy is the network outsourcing business, in which the network owner makes a regular payment to the outsourcer (often related to the amount of services that are actually used) to manage the network. As in cloud computing, outsourced computations may be carried out on hardware that is owned and maintained by the service provider. Network outsourcing contracts can be a hundred pages long. They typically include a detailed service level agreement, which may specify a percentage network uptime, data backup intervals, response times, an audit standard and a data security standard which must be provided by the service provider, with penalties (usually in the form of credits for future services) if the specified service level is not met. Drafting, negotiating and disputing such contracts are lawyer-intensive activities. Jennifer Jones' advice for companies purchasing this type of service is "First, hire all the lawyers".²²

It is likely that in the future, cloud computing will go both ways. There will be one market in which services are cheap or free and advertising-supported, and in which the customer takes nearly all the risk. This is the typical model for cloud computing agreements at the moment. But there will also be a market for cloud computing in which the service provider takes more of the risk, in return for more payment.

Service level agreements are already being offered for some cloud computing services, but they are at present rudimentary compared to those offered for network outsourcing. Amazon has had a service level agreement for its storage service Amazon S3 since October 2007,²³ and for its computing service Amazon EC2 since October 2008,²⁴ which promise 99.9% availability measured over a month for S3 and 99.95% availability over a year for EC2, (excluding *force majeure* downtime) or a refund of 10%-25% of customer's payment for the last billing period, paid in service credits. In order to receive the service credits, customers have to document details of outages, and make a claim to Amazon. The Amazon S3 service had outages lasting over an hour in February and July 2008.²⁵ In contrast, some contracts for outsourced management of onsite networks specify 99.999% network availability. However there are many applications for which the availability given by Amazon S3 and EC2 is perfectly adequate.

There is an intriguing clause in the Amazon Web Services customer agreement which forbids customers from disclosing, for three years after the end of the term of the agreement, "the nature, content and existence of any discussions or negotiations

²² J Jones, "Data Diligence" (2005) *ComputerWorld*, 14 November 2005 available at <http://www.computerworld.com/managementtopics/management/story/0,10801,106127,00.html> (accessed 23 Mar 2009).

²³ Amazon Web Services LLC, "Amazon S3 Service Level Agreement" (2009) available at <http://aws.amazon.com/s3-sla/> (accessed 23 Mar 2009).

²⁴ Amazon Web Services LLC, "Amazon EC2 Service Level Agreement" (2008) available at <http://aws.amazon.com/ec2-sla/> (accessed 23 Mar 2009).

²⁵ J Brodtkin, "More outages hit Amazon's S3 storage service" *NetworkWorld*, 21 Jul 2008 available at <http://www.networkworld.com/news/2008/072108-amazon-outages.html> (accessed 23 Mar 2009).

between you and us.” It is not indicated what these negotiations might be about, but one possibility is that Amazon may be prepared to offer selected customers a more reliable service in return for higher fees.

Legal requirements for the handling of particular types of data (for example health data and financial data) are one of the forces creating a market for cloud services with more stringent service level requirements. Some laws and regulatory regimes place requirements for auditing or data security that may not be provided by current cloud services. Pharmaceutical companies and financial organisations have best-practice requirements that some kinds of data have to be stored on an identifiable server. It should not be impossible to build a cloud storage service which can identify for the customer the precise server or servers on which their data is currently stored, although such a service might be less efficient. Similarly, although UK companies storing personal data with some cloud computing services might find themselves in breach of the seventh principle of the *Data Protection Act 1998* if the standard subscription agreement for the services does not give sufficient (or indeed any) guarantees that the computers that the data will be stored on are appropriately secure,²⁶ it is possible to create cloud services that meet industry security standards; for example, Google Apps has SAS70 Type II certification.²⁷

On the other hand, other legal requirements may present more difficult obstacles to the use of cloud computing for some applications. In particular, some licensing models for software and for copyright data fit poorly with cloud computing. Jimmy Lin of the University of Maryland gives an example of data whose license restrictions forbid it to be copied onto computers outside Maryland.²⁸ Per-seat pricing for software can have the effect of disallowing use of the software in cloud computing.

5. Subcontracting

*Every minute that white woolly plain which covered one-half of the moor was drifting closer and closer to the house. Already the first thin wisps of it were curling across the golden square of the lighted window. The farther wall of the orchard was already invisible, and the trees were standing out of a swirl of white vapour.*²⁹

A cloud service provider may subcontract parts of the services, and this subcontracting will usually be invisible to the buyer. This raises issues and potential disputes that are also raised by complex subcontracting agreements in other industries. These issues include contracting and auditing requirements, and questions of the distribution of payment if all goes well, and of liability if it does not – for instance, if

²⁶ J Salmon, “Clouded in uncertainty – the legal pitfalls of cloud computing” (2008) *Computing*, 24 Sept 2008 available at <http://www.computing.co.uk/computing/features/2226701/clouded-uncertainty-4229153> (accessed 23 Mar 2009).

²⁷ Google, “What does a Google Apps SAS70 Type II audit mean to me?” (27 Jan 2009) available at <http://www.google.com/support/a/bin/answer.py?hl=en&answer=138340> (accessed 23 Mar 2009).

²⁸ K Greene, “Google’s Cloud Looms Large” (2007) *MIT Technology Review*, 3 Dec 2007 available at <http://www.technologyreview.com/business/19785/page2/> (accessed 23 Mar 2009).

²⁹ A Conan Doyle, note 5, at 191.

two cloud computing subcontractors provide micro-services which are of good quality in themselves but do not integrate properly.

There are two factors in cloud computing which make these problems more acute than in other industries. The first is that the choice of subcontractor might be changed hourly according to availability and price. In this case the contracting and auditing involved with a change of subcontractor will need to be done rapidly and frequently, with as much automation of these processes as possible. To complicate matters, the new subcontractor may be in a different country from the old one, with the result that different laws may apply.

The second issue is that rather detailed data, which may be commercially valuable, flows from the customer to the vendor to the subcontractor. In other industries for which subcontracting is common, the subcontractors typically receive only a small amount of data about the customers, if they receive any at all. Suppose a cloud computing vendor goes bankrupt. Can the subcontractors hold onto the customer data as an asset? Can they threaten to publish sensitive data if they are not paid for their services? It is also not clear what rights a purchaser of the bankrupt vendor might have to the data.

In normal operation, the fact that the data flows to subcontractors means that there are additional opportunities for it to be stolen or misused.

6. Who can use which data?

So as the fog-bank flowed onward we fell back before it until we were half a mile from the house, and still that dense white sea, with the moon silvering its upper edge, swept slowly and inexorably on. "We are going too far," said Holmes.³⁰

As mentioned in the previous subsection, data used in cloud computing services may flow onward to several unknown computers within the cloud. It is currently a foggy question as to who within the cloud computing ecosystem should have rights to use data originating from buyers of cloud services, and for what purposes. The model of advertising-supported cloud computing services assumes that providers will obtain some rights to use customers' data for purposes other than those necessary for the technical provision of the service.

A cloud computing service that processes data for a small business will have access to four different types of data. First, there is data about the small business' own customers. For example, the data for an Easter egg business might include a list of past Easter egg purchasers, and information about their purchases. Second there is account data about the small business itself, including its contact and payment details. Third is data generated by the operation of the services (some of which, for example the internal state of applications, may not be accessible to the small business). Last there is activity data, which tracks when and for which applications the business' account with the service provider is used.

The Service Level Agreements for cloud computing services do not always make it clear what rights the providers have to use this data (including the circumstances

³⁰ A Conan Doyle, note 5, at 191.

under which they can sell it), and in particular do not always distinguish clearly between these different types of data. Some clarification may be necessary in contracts for future services. For example, the Amazon S3 storage service lets customers choose to store data objects that they upload to the service in Europe rather than the US, but it is not clear from the S3 terms of service whether or not their account details and activity data will also be stored in Europe if they make this choice.

7. Lock-in

*The fog-bank lay like white wool against the window. Holmes held the lamp towards it. "See," said he. "No one could find his way into the Grimpen Mire tonight." She laughed and clapped her hands. Her eyes and teeth gleamed with fierce merriment. "He may find his way in, but never out," she cried.*³¹

A survey of cloud computing customers by RightScale³² found that their main concern was the possibility of being locked in to a cloud computing provider. Lock-in is possible at several different layers of the computing stack: a company might find it difficult to change their cloud infrastructure provider and also difficult to change their provider of cloud-based software for managing marketing campaigns.

If one company's application programming interface for cloud computing services becomes a *de facto* standard, the company may gain a market position such as that enjoyed by Microsoft in operating systems, with implications from competition law. Even if a company does not dominate at a particular layer of the stack, they can make it difficult for their customers to change provider. For instance, Salesforce.com uses a custom programming language, custom objects, and custom user interface tags.

Competing services may use different assumptions about computing environments or interfaces with other parts of the cloud, so that lock-in at one layer may produce lock-in at another. Indeed, since it is predicted that most money is to be made at the high level of cloud computing, selling specialized software applications as services, the intention of some telecoms companies is to make use of their market position at the infrastructure level to capture consumers for higher-level cloud services.³³

There are some attempts to prevent lock-in by providing open source alternatives, such as the open source database Prophet,³⁴ and RightScale proposes a cloud meta-service that enables customers to switch easily between one service provider and another; this however begs the question of lock-in to the meta-service provider.

³¹ A Conan Doyle, note 5, at 196.

³² M Crandell, "RightScale: the cloud management platform" (2009) Powered by Cloud, London, 2-3 Feb 2009.

³³ D Lupafya, panel session at Powered by Cloud, London Feb 2-3 2008.

³⁴ J Vincent, "Prophet: a path out of the cloud" (2008), Open Source Convention (OSCON), Portland OR, 21-25 July 2008 available at <http://assets.en.oreilly.com/1/event/12/Prophet.%20your%20path%20out%20of%20the%20cloud%20Presentation.pdf> (accessed 23 Mar 2009).

8. *Hey, you, get off of my cloud*

“Hist!” cried Holmes, and I heard the sharp click of a cocking pistol. “Look out! It’s coming!” There was a thin, crisp, continuous patter from somewhere in the heart of that crawling bank. The cloud was within fifty yards of where we lay, and we glared at it, all three, uncertain what horror was about to break from the heart of it.³⁵

In this section I will describe several horrors – mostly related to security – that might be about to break out from the cloud. Cloud services can provide considerable opportunities for legitimate businesses, but carelessly or maliciously designed cloud services may also offer entrepreneurial opportunities for tax evaders, industrial spies, data thieves and denial-of-service extortionists. These opportunities are a potentially fertile source of future legal cases.

The geographical flexibility offered by international computing networks is already used by offshore banks and casinos to pick jurisdictions that suit them. Cloud computing offers the prospect of further decentralized and globalized companies, which deal with customers using machines in countries with little consumer protection, store customer details in countries with light privacy requirements, and bank the profits in countries with low taxes and high banking secrecy. It may be complicated to determine the jurisdiction in which to sue such a company.

Since data generally has to be unencrypted at the point of processing, this means that if it is processed using cloud computing, it will generally be present in unencrypted form on a machine in the service provider’s or subcontractor’s network. There is therefore a risk of theft or sabotage by a rogue employee of the service provider or subcontractor.

An advantage of cloud computing is that it enables services to be provided within the same network to many different customers, in an environment which gains efficiency by sharing resources. There need therefore to be technological protections to prevent customers (who may be commercial rivals) from spying on each others’ data or interfering with each others’ computations. There is also a possibility that system errors may produce unintentional leaks of information from one customer to another. A flaw in the Google Docs application, now fixed, had the effect that some users inadvertently shared some of their documents.³⁶

Large cloud service providers may have a large amount of commercially valuable data and services in their network. As a result, the network will become a juicier target for data thieves and denial-of-service extortionists. Thefts of data via cloud services are not unknown; in 2007, cybercriminals targeted Salesforce.com’s customers and succeeded in stealing e-mail addresses and phone numbers by using a phishing

³⁵ A Conan Doyle, note 5, at 191.

³⁶ D Raywood, “Google admits that some of its Docs have been accidentally shared” (2009) *S C Magazine*, 10 Mar 2009 available at <http://www.scmagazineuk.com/Google-admits-that-some-of-its-Docs-have-been-accidentally-shared/article/128491> (accessed 23 Mar 2009).

attack.³⁷ Moreover, the use of homogeneous virtual machines by cloud computing services may amplify the power of attacks. Nitesh Dhanjani notes that if a remotely exploitable vulnerability were found in the generic kick-start image that Amazon recommends to its customers, or if a malicious user were able to terminate virtual machine instances outside her role, these vulnerabilities could be used to attack the whole of the Amazon EC2 cloud.³⁸ The counter-argument is that small and medium businesses' own networks almost certainly have less protection against external attackers than Amazon's network does, so that despite these issues, businesses that take up cloud computing may actually increase their security from such attacks.

Some cybercriminals, for example those using botnets, already exploit the possibility of using remote access over the Internet to circumvent security systems designed to prevent or detect their attacks. If a carelessly designed cloud computing service enabled its customers to circumvent some security systems, for example if it allowed them to disguise their location by launching an external communication from a machine within the cloud, many more people might exploit this possibility. For example, the fact that the CoDeeN academic content distribution network allowed client access from outside the hosting organisation was used by spammers and password crackers, but also by people accessing academic journals to which they did not have a subscription, and by people remotely viewing academic web pages on University intranet sites that were not intended to be externally accessible.³⁹

The usual business model for cloud computing is that it is partly subsidized by the sale by providers to third parties of services based on the providers' access to customers' data. This may simply involve the customers' data being used to target third-party advertisements that the provider includes with the service, but there has also been discussion of the sale of customers' data for lead generation and market intelligence. Cloud computing customers need to check what data their providers will release to third parties, and under what conditions. Businesses are generally very reluctant to enable their direct competitors to have access to their customer contact lists or detailed sales data. But even innocuous-seeming data might turn out to be commercially sensitive. For example, someone who obtained information about the load levels for a company's hypervisor might be able to use this information for insider trading if the company used the cloud service for a crucial application. There is a particular issue about the sale of anonymized data. In the last few years there have been technical advances in the art of de-anonymization (strictly speaking, de-pseudonymization). Pseudonymized data is still proof against opportunistic thieves, but an attacker who is out to get a particular customer, already knows a small amount of data about him, and is prepared to do some data processing, may be able to recover

³⁷ A Greenberg, "Cloud Computing's Stormy Side" (19 Feb 2008) available at http://www.forbes.com/2008/02/17/web-application-cloud-tech-intel-cx_ag_0219cloud.html (accessed 22 Mar 2009).

³⁸ N Dhanjani, "Amazon's Elastic Compute Cloud [EC2]: Initial Thoughts on Security Implications" (28 April 2008) available at <http://www.dhanjani.com/blog/2008/04/index.html> (accessed 23 Mar 2009).

³⁹ V Pai et al, "The Dark side of the Web: An Open Proxy's View" (2004) 34 *ACM SIGCOMM Computer Communication Review* issue 1, 55-62 available at <http://portal.acm.org/citation.cfm?doid=972374.972385> (accessed 23 Mar 2009).

all the data about him in a large pseudonymized data set.⁴⁰ Previously-used practices of selling cloud customers' personal or sensitive data in pseudonymized form – and regulatory guidelines approving these practices – may have to be revised.

In addition to sales of data that have been authorized by the customer, there is also a possibility of unauthorized sales of data. Imagine that there is a cloud service provider that provides a niche application which is particularly useful for one particular company, and gives high quality at a low price. What the company does not know is that the reason that the price is so low is that the service provider is secretly selling the data that passes through their system to one of the company's competitors. There is currently research being done in designing services to automatically match businesses to the cloud-based service offerings that give the best fit to their needs at the best price. If the company used such a matching service, it would be automatically matched to the provider practising industrial espionage as a service, rather than to competing providers that would not be able to offer such a low price.

8. Conclusion: Cloud Computing advantages

*And now I come rapidly to the conclusion of this singular narrative,
in which I have tried to make the reader share those dark fears and
vague surmises which clouded our lives so long*⁴¹

This paper has described several dark fears and vague surmises about cloud computing. However, it is likely that a combination of technical solutions, business practices, and standard contracts between service providers and customers will be able to resolve most if not all of them. What appears at the moment to be a demon hound from Hell may turn out to be just a dog.

Moreover, there are plenty of applications for which these issues do not arise. These are applications, such as the format conversion of public-domain New York Times articles, which do not involve private data, are not subject to special data handling regulations, and do not require high quality of service – 99% availability would be fine, and lost or corrupted results can be recalculated. (In fact, the New York Times format conversion was run twice because an error in the output was spotted after the first run was finished). Cloud computing can be used today for such applications. It is worth working to extend its usefulness for other applications, because it potentially offers advantages to everyone involved.

For buyers, one advantage of using cloud computing, as opposed to buying all the hardware and software necessary to meet their computing needs, is that they only need to pay for the computing services that they actually use. Moreover the cost of computing is moved from upfront capital expenditure to an operating expense, a feature that is especially welcome to business users in a period of recession and tight credit: it reduces risk for start-up companies, and enables larger enterprises to avoid the significant capital expense of building a new data centre. Another advantage is

⁴⁰ A Narayanan and V Shmatikov, "Robust De-anonymization of Large Datasets (How to Break Anonymity of the Netflix Prize Dataset)" (2008) IEEE Symposium on Security and Privacy (S&P), Oakland CA, 18-21 May 2008, available at http://arxiv.org/PS_cache/cs/pdf/0610/0610105v2.pdf (accessed 23 Mar 2009).

⁴¹ A Conan Doyle, note 5, at 197.

ease of use. If you use cloud computing services for your Easter egg business, it will be the service provider's responsibility to maintain, patch and upgrade the servers on which these services run. The services will potentially be accessible from anywhere with Web access (although there may be geographical or other restrictions for legal reasons). If more orders for Easter eggs arrive than were expected, the burden of rapidly finding the additional computing resources to process the orders will be taken by the service provider, not you.

For subcontractors, cloud computing offers a whole new market. For vendors of computing services, the advantages include the potential for higher margins and for advertising revenue. They also may see growth in the size of the market for computer services, because cloud computing makes viable some small business models which would have struggled in its absence. Finally, cloud computing may provide a good source of income for lawyers.

9. Acknowledgements

Thanks to my colleagues Peter Toft, John Manley, Anna Fischer, Nigel Edwards and Patrick Goldsack; and also to the organisers of the GikIII and Powered by Cloud workshops, where I presented early versions of this paper.