CHAPTER 3

The Cloud Threat Landscape

INFORMATION IN THIS CHAPTER

- The cloud threat landscape
- Notorious nine
- Additional threats

Outsourcing any workload to a third party will introduce risk, although whether the overall risk increases is entirely dependent on the provider and the security deployed within the internally provisioned service. Regardless of whether the service is provisioned internally or externally, there will always be an element of risk; this is because while the workload can be outsourced, the risk rarely is. As discussed in Chapter 1, security is seen as a major barrier in the adoption of cloud computing. Many of these security concerns would also apply to internally provisioned services and traditional outsourcing; there are, however, some threats that are specific to cloud computing.

In this chapter, we will review the security threats for the cloud based on the research conducted by the Cloud Security Alliance (CSA) Top Threats Working Group. It is worth noting, however, that the security considerations for cloud computing extend beyond those presented within the findings of the working group and are published as "The Notorious Nine":

1. Data breaches
2. Data loss
3. Account or service traffic hijacking
4. Insecure interfaces and application programming interfaces (APIs)
5. Denial of service
6. Malicious insiders
7. Abuse of cloud services
8. Insufficient due diligence
9. Shared technology vulnerabilities

CSA Guide to Cloud Computing. http://dx.doi.org/10.1016/B978-0-12-420125-5.00003-0

The end of the chapter will include references to additional sources that define the threats to cloud computing. Many of the security considerations are "traditional threats," in other words, those that would exist both in internally provisioned services and within a cloud implementation. For example, the requirement to introduce appropriate physical security controls would exist in both internally provisioned services and cloud implementations. Within a cloud deployment, however, the customer will not have the same level of transparency regarding the level of security deployed by the provider. Subsequently, the controls deployed will likely be articulated as part of the security certification(s) and reviewed by the customer as part of the process of selecting a provider (see Chapter 2). Furthermore, the level of flexibility afforded to end customers regarding the implementation of security controls is generally lower with cloud computing (as was discussed in Chapter 2).

THE CLOUD THREAT LANDSCAPE

Utilizing the cloud provides organizations with many business benefits, but with these benefits come a number of threats. Some of these threats are the traditional threats that we are accustomed to while others are unique to the cloud. By better understanding the various threats that can face our data and services in the cloud we are better prepared to determine how best to secure them.

Before examining the various threats, it is important that we first understand what a threat is. There are many different interpretations and definitions for threats in the context of computer security. The Oxford English Dictionary defines a threat as

> (noun) (1) a stated intention to inflict injury, damage, or other hostile action on someone. (2) a person or thing likely to cause damage or danger. (3) the possibility of trouble or danger

In security fields we tend to focus on the second definition "a person or thing likely to cause damage or danger." However, we need to focus further into what exactly a threat is, particularly in relation to information security.

According to the International Organization for Standardization 27001 Information Security Standard, a threat is defined as

> a potential cause of an unwanted incident, which may result in harm to a system or organization.

Under the Payment Council Industry's Data Security Standard a threat is described as

> Condition or activity that may cause information or information processing resources to be intentionally or accidentally lost, modified, exposed, made inaccessible, or otherwise affected to the detriment of the organization.

The National Institute of Standards and Technology definition of a threat is given in SP 800-301 and defines a threat as

> the potential for a threat-source to exercise (accidentally trigger or intentionally exploit) a specific vulnerability

While the above definitions seem to be more relevant to the information security, the definition supplied by the European Network and Information Security Agency (ENISA) probably provides the most apt definition, in particular when taking cloud computing into account.

According to the ENISA,[1] a threat is

> Any circumstance or event with the potential to adversely impact an asset through unauthorized access, destruction, disclosure, modification of data, and/or denial of service.

Having understood what a threat is, it is important to appreciate how threats against computer systems have evolved over the years. This is not just so we can better understand today's threats but also so we can appreciate that as computing technology evolves and our business and personal use of it also evolves, so too will the threats.

Evolution of Cyber Threats

Since we first started using computers they have been under threat. Those threats come from various sources whether they are from those with malicious intent, from well-intentioned people making mistakes, man-made failures such as power outages, or indeed natural disasters. As our use of computers and the Internet has grown over time so too has the number and the sophistication of the threats facing those systems.

In the early years of computing, the main source of threats against computer systems were mainly from internal threats such as disgruntled or unhappy employees, or from the well-meaning user who makes a mistake. The other threats faced by these systems were from natural sources or man-made sources such as hardware failures or software bugs. This low level of threats was due to many such computer systems being isolated from other systems outside their own organization's offices and buildings. As a result, the threats against these systems were mostly limited to those with physical access to those systems or from disasters in the locale.

Over time, access to these systems became more and more frequent with companies employing modems and wide area networks to allow remote offices and users to connect to them. While enabling remote users to gain from the benefits of these systems, it also opened up these systems to threats from external parties.

At this stage in the evolution of computing, the external threats posed to organizations' systems were restricted to mainly individuals who broke in and explored these systems out of curiosity to determine how computers, networks, and systems worked. In the main, there was no malicious intent in this type of activity with the primary motive being curiosity.

In the 1980s, we witnessed the introduction of personal computers and their subsequent growth not just in home use but also within corporate environments. Over time, and as a result of these developments, companies and organizations saw their staff becoming more and more productive as they moved from a centralized computing model to a distributed one. The growth in use of personal computers saw data being moved from being stored and managed on a central location onto individual computers located throughout organizations.

In parallel to this growth in the use of Personal Computers, there was also the growth in the use of the Internet. With the growth of the Internet, many organizations took advantage of its openness and global spread to enable them to promote their services, products, and their brands to existing and potential customers. Other Internet-based technologies also enabled workers to share information with others and to be more productive and effective.

All these new technologies brought many advantages to organizations and indeed to society and the economy in general. However, legitimate businesses and organizations were not the only ones taking advantage of these new technologies. Those with malicious intent also saw the opportunities in this brave new world.

In the early stages, the number of attackers looking for financial gain from stealing information from systems also started to increase. While the majority of online attacks still came from those with curiosity as their main motive, many others saw the Internet as a way to promote their political cause or other activism by attacking and disrupting systems to raise awareness of their cause, or by defacing an organization's Web site and posting their messages online.[2]

The threat posed by those with looking to gain financially also increased as they looked to extort money from organizations by defacing their Web sites and extorting payment from them to stop their Web site from being defaced again, or by stealing information from their systems.

With the dawn of the twenty-first century, we saw an explosion in organizations rushing to store and transmit more and more data on their computer systems, we also saw a surge in the use of the Internet by organizations to promote and sell their products and services. As companies rushed to benefit from computers and the Internet so too did those with malicious intent. As the value of information grew and the ability to steal that information through insecure

systems equally grew, we witnessed a change in the online criminals. No longer a niche arena for individuals, or small numbers of like-minded people, cybercrime now attracted traditional organized criminal gangs as they saw many new opportunities to make vast sums of money by exploiting weak computer security with relative low risk of being prosecuted.

This evolution in online threats was also mirrored by the growth in sophistication of computer viruses of the same timeline. The early computer viruses were not very sophisticated[3] and were primarily designed to disrupt the operation of the systems they infected, often in amusing ways, such as the cascade[4] and ping-pong[5] viruses. As these viruses were easily detected due to their disruptive nature, they could be eliminated with the appropriate security tools or by rebuilding the system. Today, however, most viruses are specifically designed to go undetected as their raison d'être is no longer to cause disruption. Instead, criminals create these viruses to go undetected on infected systems so they can be used to steal valuable data such as sensitive financial data, logon credentials to financial systems, or valuable information such as an organizations' intellectual property.

The modern computer virus is also designed not to just steal information but also to enable online criminals use infected computers in other criminal enterprises such as spending spam e-mails, infecting other computers, and extorting money from companies by using the infected computers under their control to take part in a distributed denial of service (DDoS).

Computer viruses are also being developed as advanced weapons to silently attack targets. The Stuxnet[6] virus is a prime example of how a computer virus can be used to silently disrupt the operations of critical target. We will no doubt see further advances in the complexity and capabilities of computer viruses in the future.

As our use of computer systems has evolved so too have the threats facing those systems; moving to the cloud is just one more evolution in our use of computers, networks, and applications and while the traditional threats facing those systems still remain, there will be other threats that will evolve specifically against cloud computing.

Knowing and understanding what these threats are will make it easier to develop strategies, solutions, and systems to counter and manage those threats.

NOTORIOUS NINE

Data Breaches

Cited as the number one security threat for cloud computing, data breaches refer to the loss of confidentiality for data stored within a particular cloud instance. It is of course worth noting that such a threat is likely to exist even within an on-premise solution, or traditional outsourced solution.

The concern over the loss of confidentiality is entirely understandable, as the potential financial and reputational cost can be significant. This will be entirely dependent on the data that have been stolen; organizations will have many types of data ranging from intellectual property and sensitive business information to personal data (e.g., customer data). For personal data, according to the "2013 Cost of Data Breach Study"[7] conducted by the Ponemon Institute, a data breach (referred to as the theft of protected personal data) can cost up to $200 per record. This cost is entirely dependent on the country in which the surveyed company resides, and as depicted in Figure 3.1.

In terms of deriving the cost per record, costs were divided into two categories, direct and indirect. Direct costs are those that refer to "the expense outlay to accomplish a given activity such as engaging forensic experts, hiring a law firm or offering victim's identity protection services. Indirect costs include the time, effort and other organizational resources spent during the data breach resolution." Dependent on the country in which the surveyed company resided, the costs varied in terms of direct versus indirect. For example, companies surveyed in the United States experienced 32% direct costs compared with those in Brazil where direct costs rose to 59%. According to insurance company Beazley[8] in their small business spotlight, the greatest direct cost associated with responding to a data breach is the notification required. This of course is more relevant to those businesses that have a requirement to notify affected customers. In the United States, for example, and as of the time of writing, and according to Bloomberg Law[9] there are only four states without a data breach notification law; these are Alabama, Kentucky, New Mexico, and South Dakota. However, the data notification requirements across the various states do differ, with varying requirements such as notification triggers and method of notification.

Now of course, the United States is not the only country where data breach notification laws exist; under the European Union's Regulation on the notification

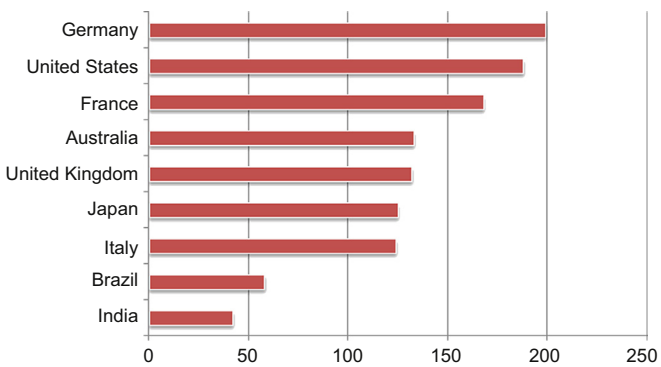

FIGURE 3.1
Estimated cost of breach per record (in USD).

of personal data breaches,[10] providers of publicly available electronic communications services are obligated to notify customers about data breaches. This notification must be done within 24 h to the national competent authority. Moreover, impending legislation, in particular in the European Union, is likely to increase the notification requirements for organizations that experience a data breach.

Notification is one cost associated with data breaches; however, as recent public data breaches have demonstrated, those affected companies have many other costs to contend with, and these may be either direct or indirect. Additional costs can include direct technical costs to identify the cause of the breach, and any remediation work to close vulnerabilities and prevent the issue from reoccurring. In addition, there are likely to be costs associated with the breach itself, such as the potential loss of business. Following the 2006 data breach experienced at the TJX Corporation in which $45 million credit and debit cards were stolen,[11] it was reported that the retailer had faced costs of over $256 million (these figures do vary greatly dependent on source; therefore, the more conservative figure is quoted here), despite initial estimates attributing the costs at a "mere" $25 million. While this level of data breach is certainly at the higher level of examples, it does provide an illustration of the impact an organization faces when experiencing a data breach, and subsequently validates the reason why it is rated as the number one concern when migrating to cloud computing. A large proportion of the costs from the TJX breach was related to the offer of services to its customers; this included credit monitoring services as well as identity theft protection. A breakdown of the estimated costs, and associated activities were presented in an article published by Wired[12] in 2007; while the actual figures in Table 3.1 may be disputed, it does provide an insight into the associated costs related to a data breach.

What these figures, or rather what these activities, clearly demonstrate are that the costs associated with a data breach can be significant, and any potential breach is quite rightly seen as a major concern. In addition, it is worth noting that some of these figures seem low and therefore it is assumed they are per record (e.g., cost per call is $25, but is likely per customer). From a cloud perspective, it is worth noting that as the risk is not outsourced, the remediation costs will be borne by the customer and not the provider. As discussed in Chapter 7, the data controller will almost always be the end customer and therefore they will be responsible for ensuring that not only is the appropriate due diligence undertaken but their own customers (data subjects) will look to them to remedy the situation. It may be possible to point the finger at a provider, but the truth is that the data subjects (whose records have been stolen) are not direct customers of the cloud provider and their decision to no longer work with the company they trusted to look after their data will affect the bottom line of the data controller. This is referred to as the abnormal churn rate, which can be as high as 4.4% dependent on geography and likely sector.

Table 3.1 Assumed Costs Related to TJX data Breach

Number of client records breached	45,600,000
Cost of detection and determination of response	$319,200,000
Internal investigation ($3.00 per record)	$136,800,000
Legal and external advice ($2.50 per record)	$114,000,000
Public relations and Investor relations ($1.50 per record)	$68,400,000
Cost of customer remediation	$1,140,000,000
Letters, e-mails, and phone calls	$25
Call center to address response ($5.00 per record)	$228,000,000
Cost per call	$25
Percentage of clients that call	20
Credit watch per year	$50
Years of credit watch	2
Percentage of clients that request credit watch	20
Cost of corporate remediation	$630,000
Fines	$150,000
Increased cost of audit/assessment oversight	$200,000
Legal defense and investigation	$100,000
Systems remediation	$180,000
Cost of down time	$100,000
Value per day of down time	$100,000
Number of days	1
Cost of brand impact	$700,000
Lost existing customers	500,000
Lost new customers	200,000
Cost of fraudulent use of data	$228,000,000
Average cost per compromised record	$50
Percentage of client records that result in fraudulent use	10
Total cost of breach	$1,688,630,000
Average cost per client record	$37
Probability of a breach	33%
Expected value based on probability of a breach	$557,247,900

Small caveat to the above statement: the provider could also experience a loss of trust if the breach is significant and public enough to negatively impact the trust of other customers, both potential and/or existing.

Other types of data can also have a significant financial impact. Research conducted by the Centre for Strategic and International Studies identifies the following categories in its report entitled "Economic Impact of Cybercrime"[13]:

- Intellectual property: "The cost to companies varies from among sector and by the ability to monetize stolen data (whether it is IP or business confidential information). Although all companies face the risk of

loss of intellectual property and confidential business information, some sectors—finance, chemicals, aerospace, energy, defense, and IT—are more likely to be targeted and face attacks that persist until they succeed." From a cloud perspective, while personal data will demand due diligence, the hosting of data classed as intellectual property should be commensurate to its value. This should include not only the cost of the research, but also the opportunity costs such research represents to the business.

- Financial crime: "Financial crime usually involves fraud, but this can take many forms to exploit consumers, banks, and government agencies. The most damaging financial crimes seek to penetrate bank networks, with cybercriminals gaining access to accounts and siphoning money." The migration of cloud services, particularly for financial services will witness greater focus from nefarious actors looking to commit fraud by targeting systems hosted by external providers. This renewed focus was reported by CNBC when "cybercriminals acting in late 2013 installed a malicious computer program on the servers of a large hedge fund, crippling its high-speed trading strategy and sending information about its trades to unknown offsite computers." Admittedly, these types of attacks are not solely targeted at cloud computing, but demonstrate the threat landscape for financial fraud involves malicious actors that are very technically adept and well resourced.
- Confidential business information: "The theft of confidential business information is the third largest cost from cybercrime and cyberespionage. Business confidential information can be turned into immediate gain. The loss of investment information, exploration data, and sensitive commercial negotiation data can be used immediately. The damage to individual companies runs into the millions of dollars."

The loss of confidentiality for an organization can have a significant impact regardless of whether the data are hosted externally or are an internally provisioned service. Using cloud computing can have enormous efficiency gains, but as the example of Code Spaces (more detail under Data Loss) demonstrates, the need for security remains and indeed one can argue that with the volume and complexity of threats increasing the need for security has never been more important. Ultimately, the loss of confidentiality will impact the cloud customers significantly, and also be to the detriment of the provider.

Data Loss

Unlike data breaches, loss of data refers to the unavailability of data stored within the cloud for the end customer. We touched on the subject briefly in the first chapter using MegaUpload as the example; however, the legal status of the provider is only one example that may potentially impact the service.

Provider Viability

What do you do when our cloud service provider (CSP) goes bankrupt? This was a question that customers of Nirvanix faced[14] when they were notified they had 2 weeks to migrate their data. Posted on their Web site on September 30, 2013, customers were advised they had until the 15th of October to ensure their data had been removed.

Two weeks. It is hardly a sufficient time frame to analyze alternate providers, conduct due diligence, and then implement a migration plan despite the company providing a list of recommendations. Indeed, reports[15] suggest that the provider had many customers with over a petabyte of data and while official notice was provided it was a full 10 days after the reports began to appear in the mainstream press.

Recognizing the impact of a provider going bankrupt has led to the introduction of legislation that allows the end customer a legal right to claim back data from a bankrupt provider. Introduced in July 2013, the European country Luxembourg introduced Article 567 p2, of the Code of Commerce.[16] This allowed the end customer the opportunity to recover "intangible and nonfungible movable assets" under the following conditions:

- "The bankrupt company must not be the legal owner of the data but only hold it;
- The claimant must have entrusted the data to the bankrupt company or be the legal owner of the data;
- The data must be separable from the other intangible and nonfungible movable assets of the company at the time of the opening of bankruptcy proceedings."[9]

The associated costs of the recovery of the data will be the responsibility of the claimant; therefore, while the law provides the means to recover data the cost of recovery will need to be factored in. Although of course the law does include cloud computing providers, its scope is considerably wider and includes any third parties entrusted with customer data. Although a significant legal document, its global scope is limited to Luxembourg; however, it serves as an indicator that the legal framework is focusing attention on the viability of providers.

Insufficient Disaster Recover (DR)/Business Continuity Planning (BCP) Practices

The benefit of migrating to the cloud is that defining the level of availability is as simple as a line entry in the contract. This really sounds simple does it not? By stating availability as 99.999% and then sitting back to use the service safe in the knowledge that the likelihood of the service going down is so unlikely (because there is the safety of a sentence in the contract). Sadly the reality is very far from this perfect world.

What happens when the service level agreement regarding the availability of service is not met? Invariably, the response as defined within the contract results in credit being issued to the end customer. Depending on the provider this is likely to be a tiered model, with greater compensation/credit being provided depending on the amount of downtime experienced. While receiving credit may be an appropriate level of compensation for many provisioned services, for many customers getting 10% credit for an hour's downtime may not compensate the loss of service. This loss of service itself is most likely to be a result of a power outage, according to recent research.[17] Of the 27 publicly reported outages in 2012, the main cause for the outage was power loss, as depicted in Figure 3.2.

What was particularly interesting within the research was that the average time to recover the services from the outage was 7.5 h, and the examples used within the research used some of the biggest names in cloud computing.

Errors

While the examples of malicious actors involve a conscious decision to affect the availability of services, not all actions are a direct result of someone malicious intentionally looking to impact the availability of a paid service. There could be something as simple as a human error, for example, an operator inadvertently deleting something or powering down an important asset. While the action may be an accident, the result is likely to be the same, namely, the unavailability of data to the end customer.

Such an example was reported by ZDNet in 2011,[18] whereby a software bug resulted in the deletion of customer data. The status page from Amazon Web Services (AWS) at the time reported the following:

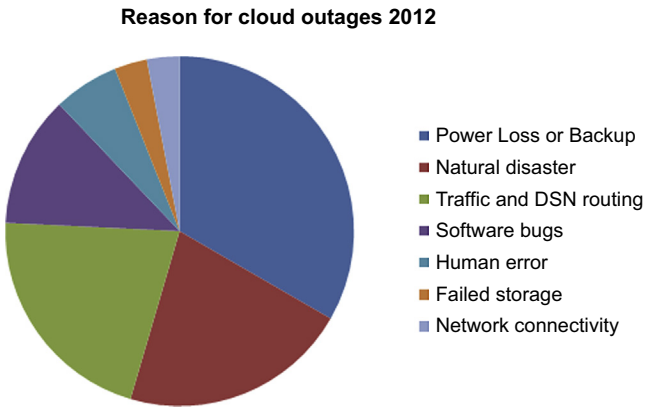

FIGURE 3.2
Research into reasons for cloud outage.

> Independent from the power issue in the affected availability zone, we've discovered an error in the EBS software that cleans up unused [EBS] snapshots…During a recent run of this EBS software in the EU-West Region, one or more blocks in a number of EBS snapshots were incorrectly deleted.[10]

The power issue in the status update refers to lightning that impacted European operations of AWS. Despite the power issue, the software bug resulted in a number of customers being without access to their data for a period of time. While the issue itself was not malicious the net result would appear to be exactly the same.

While these examples of potential threats to a cloud service can be mitigated by employing a secondary service, or with the requisite assurance that the provider employs sufficient business continuity practices, such costs should be factored in. Therefore, the cost presented by the provider is unlikely to be the total cost of ownership for the provision of an outsourced solution. Equally, the aforementioned examples are only a small snapshot of some of the reasons for data loss, one glaring omission are the actions of malicious actors, or "hackers" if we adopt the media definition. The recent case of Code Spaces provides a stark warning to organizations looking to leverage cloud computing without implementing the appropriate level of security. In June 2014, it was reported[19] that the company was "forced to close its doors after its AWS EC2 console was hacked." The company faced a DDoS attack in June 2014, and also an "Amazon Web Services (AWS) Elastic Compute Cloud (EC2) console and left messages instructing the company's management to contact them via email." Although they were able to change their passwords, the intruder leveraged backup accounts created in the intrusion. However, the "hacker removed all Elastic Block Storage (EBS) snapshots, Simple Storage Service bucks, AMIs, and some EBS and machine instances. Most of the company's data, backups, machine configurations and off-site backups were either partially or completely deleted, leaving Code Spaces unable to operate." Of course, this particular example could have applied to internally provisioned, just as easily as to those hosted with a CSP. However, as Nathan McBride, Chief Cloud Architect for AMAG Pharmaceuticals, puts it, "if you're going to put your eggs in the AWS basket, you have to have the mechanisms in place to really solidify that environment."[20] To be fair, this statement could be made about any cloud provider.

Account or Service Hijacking

As the case of Code Spaces demonstrates, knowledge of one's password can have serious repercussions. Consider the implications of someone knowing your social media password, for example, Twitter or LinkedIn. These issues have played themselves out with some very notable instances, for example, in early 2013, the Twitter account of Burger King was compromised and the attacker, in addition to sending tweets under the guise of Burger King, changed

the logo of the account to that of its competitor McDonalds.[21] This of course is one of many examples and is not specific to cloud computing.

In a cloud environment this threat was one that was realized by Wired journalist Mat Honan who wrote of his experience in an article entitled "How Apple and Amazon Security flaws led to my epic hacking."[22] The article demonstrated the ease with which hackers were able to take control of key cloud-related services using simple social engineering techniques leveraging information that the journalist describes that "anyone with an Internet connection and a phone can discover." Indeed this information (at the time in 2012, and as reported by the journalist) consisted of a billing address and the last four digits of a credit card to encourage the vendor to issue a new password, and ultimately allow a malicious individual access to the target's iCloud account. According to an interaction with an individual the journalist believes was party to the hack, they then proceeded to outline the manner in which they accessed these two critical pieces of information to enable a reset of the Apple ID. A billing address in this example was straightforward as the journalist had a personal Web site, and registered the domain to this particular address. Therefore, a simple WHOIS search provided the billing address.

The last four digits of the credit card was garnered through a social engineering technique to the Amazon call center, and with that the attacker(s) had all of the information necessary to carry out a call to the Apple call center. What this example demonstrated is the ease with which the password mechanism was circumvented through a simple social engineering technique (the journalist confirms they repeated the social engineering tricks the attacker outlined with equal success). Clearly, stronger authentication mechanisms would have posed more of a problem for the attacker.

According to penetration tester Peter Wood,[23] the ease with which cloud credentials can be hacked is exasperated by the ability to log in from anywhere as the content is primarily delivered through a browser. He continued to highlight that "Spear phishing is massively increasing as a primary entry point technique," and noted the increasing use of social engineering techniques (such as those experienced by Honan). "We get social engineering attacks by telephone almost every week," said Wood. Of course, the rise of spear phishing attacks are being made simpler due to the relative ease in being able to research the target. With social media profiles offering attackers a veritable banquet of valuable information, delivering an attractive e-mail to the target to induce a click is really very simple. Indeed, such is the prevalence of spear phishing that according to security firm Trend Micro, "In an analysis of targeted attack data, collected between February and September 2012, Trend Micro found 91% of targeted attacks involved spear phishing."[24]

We could quite easily spend the rest of this chapter, and beyond, to dive into the rise and psychology behind spear phishing attacks. Tempting really does distract from the purpose of focusing on the specific threats impacting cloud

computing. However, it is worth noting that broadly speaking, the majority of spear phishing threats will look to leverage the influencing levers. According to psychologist Robert Caldini, there are six principles of influence[25] that are used to convince others. In today's spear phishing attacks, it is not uncommon to see these principles being used, some of course will be more prevalent than others:

- Reciprocation: By carrying out a favor to someone, they invariably feel obligated to return the favor. Consider social media, when writing a recommendation does that individual feel obligated to return the favor?
- Scarcity: Something is more attractive when it becomes scarce; with many social engineering attacks using e-mail this principle is used. For example, warnings that an account may be closed unless verification is provided (e.g., entering personal data) use scarcity principles.
- Consistency: Once we have committed to something we are more likely to see this through.
- Liking: We are more likely to be influenced by those we like.
- Authority: A sense of duty of obligation exists to those in authority. We will invariably see this principle used through the use of e-mails that purport to be from financial institutes, for example, with attackers making great efforts to appear the communications are coming from the legitimate institute.
- Social validation: There is greater safety in numbers, and social validation looks to leverage this belief.

To meet this particular threat, of course, user education is an important approach. However, organizations leveraging cloud services should not only consider using stronger authentication principles, but also ensure that the service management (e.g., resetting passwords/access) have, and enforce strong validation against callers.

Insecure Interfaces and APIs

APIs within cloud environments are used to offer end customers software interfaces to interact with their provisioned services. There are multitudes of APIs available within a cloud environment; these can include provisioning new hardware and monitoring the cloud services, as just two examples. According to API Management Company, Mashery, there exist three categories of Cloud APIs[26]; these are

- Control APIs: APIs that allow the end customer to configure their cloud provisioned service. Amazon EC2 provides a multitude of APIs that allow customers to configure their services, as defined within the Amazon Elastic Compute Cloud: API Reference.[27] Examples include the allocation of internet protocol (IP) addresses, creating/editing of access control lists, or monitoring of specific instances.

- Data APIs: APIs within which data may flow into or out of the provisioned service. Such data flows can also be into alternate cloud providers, so that data can flow from one provider and into the provisioned service provided by an alternate provider.
- Application functionality APIs: Although the earlier APIs provide the ability to transfer data between alternate providers, or indeed management of the overall solution, the application functionality APIs can provide considerably more functionality that the end customer can interact with, ranging from the simple availability of shopping baskets to integration with social networking solutions, and considerably more in between.

While the flexibility of cloud APIs is not in question, and indeed depending on the source considered one of the driving forces behind the widespread adoption of cloud computing, there does remain considerable security considerations.

Indeed, these security considerations may not even be malicious, whereby an administrator may inadvertently invoke an action that may have significant repercussions. Consider the command available for EC2 customers entitled ec2-terminate-instances. As you can likely guess, this command will terminate an EC2 instance, the implication of this action is that the data stored within the instance will also be deleted.

In order to reduce the risk of such an action being inadvertently carried out, there is an opportunity to implement a safeguard to prevent inadvertent deletion using a feature available through the AWS console, command line interface, or API. Such a feature provides protection against termination with the DisableApiTermination attribute; this controls whether an instance can indeed be terminated using the console, Command Line Interface, or an API.

While such a feature, or rather attribute, is an important step in preventing accidental deletion of a particular instance, it is only one example of where an accidental action can have significant repercussions. A simple error such as mistyping the IP address for an instance is equally likely to result in the unavailability of the provisioned service, and does not have the luxury of an attribute to protect against the error. While of course the latter example is a simpler fix than the deletion of an instance, these examples do demonstrate some of the challenges facing the use of cloud APIs.

Other challenges facing cloud end customers, and their use of APIs, are also malicious attempts to circumvent authorized process. In a recent article published by DarkReading,[28] author Rob Lemos presents the security risks API keys present to their end customers. Such keys are utilized to identify applications utilizing provisioned services; however, should such keys fall into the hands of malicious actors they can be used to capture confidential data or rack up fees

for the end customer. The issue has arisen not due to a weakness in the keys themselves, but rather the manner in which they are managed, whereby in particular implementations they are used to identify users, and as such are not protected by developers as assets that are critical to the business with examples of them being e-mailed and being stored on desktop hard drives.

Recently, the CSA chapter Switzerland (https://chapters.cloudsecurityalliance.org/switzerland) held a chapter meeting focusing entirely on service orientated architecture as it relates to cloud computing in which coauthor Raj Samani recently spoke. This meeting focused on the security challenges relating to APIs within a cloud environment and presented emerging research within this field. Emerging areas of research include the use of technology to enforce access policy, and governance rules as they pertain to the use of APIs. It is therefore recommended for the reader to coordinate with the chapter should they wish to get more detailed information about this very important (and sadly not hugely researched) topic.

Denial of Service

A Denial of Service (DOS) or its now more popular unruly child the DDoS attack is not a new phenomenon, and has plagued information technology (IT) managers for many years. It refers to an attack that aims to overwhelm the victim with network traffic or consume resources (central processing unit, Memory, for example) and subsequently prevent the processing of legitimate requests. The various types of DOS can be broadly defined into two categories:

- Infrastructure-based attacks: These particular attacks reside within layers 3 and 4 of the Open Systems Interconnection model (OSI) stack, but in effect intend to submit large volumes of traffic intended to overwhelm the target, and prevent its ability to respond to legitimate requests. It is now considerably easier to initiate such attacks. In the McAfee report entitled "Cybercrime exposed,"[29] DOS (or DDoS) services are accessible to anybody with access to a search engine, and can be purchased for as little as $2 per hour. Subsequently, the probability of such attacks occurring is increasing and this is reflected in the report published by Prolexic in their "Quarterly Global DDoS Attack report Q3 2013"[30]; compared to Q3 2012 the total number of attacks increased by 58%, with infrastructure-based attacks increasing by 48%.
- Application-based attacks: Unlike the use of traditional infrastructure-based DDoS attacks, the emerging trend has been for the use of layer 7 attacks (OSI stack). What this actually means is that rather than using network traffic to overwhelm the target, it would use traffic that appears legitimate. According to Prolexic, these particular attacks represent around 20% of DDoS attacks, but still a 101% increase on the preceding year.

When considering a DOS attack as it pertains to cloud computing, there are two main considerations: (1) the threat of DOS attacks against provisioned cloud services and (2) how cloud computing (and predominantly dedicated Software as a Service (SaaS) services) can be used to reduce the risk of DOS attacks. In this section, we will focus on the first scenario.

Denial of Service against the Cloud

The migration to a cloud computing platform should provide greater protection against such attacks than traditional internally hosted service. This at least is the view taken by the ENISA, their publication entitled "Critical Cloud Computing"[31] takes the view that "Elasticity is a key benefit of cloud computing and this elasticity helps to cope with load and mitigates the risk of overload or DDoS attacks. It is difficult to mitigate the impact of peak usage or a DDoS attack with limited computing resources."

This perspective is of course entirely valid, whereby a typical network-based DOS (or DDoS)-based attack should indeed be better mitigated leveraging a service with redundancy in its resources. Equally, with the probability for a DDoS attack against a CSP likely to increase, the provider will be expected to invest more in providing controls to mitigate the threat. Cloud provider Rackspace,[32] for example, provides specific DDoS mitigation services to customers that can be added as a subscription service, or on demand. Regardless of the pricing model, the service intends to undertake assessment against incoming traffic and in the event malicious traffic is detected transfer to a "sanitation engine" to filter the traffic and forward legitimate traffic to its intended destination. This is one example; other CSPs also offer such mitigation services but the challenge for any potential customer is the effectiveness of the documented solution. In other words, the true test of any paid (or even one that is included and marketed by a provider) solution to mitigate DDoS attacks is during an attack. Such experiences were documented by The Register[33] where code hosting provider BitBucket faced 19 h of downtime due to a DDoS attack on the infrastructure it purchased from AWS. According to Jesper Nøhr, who runs BitBucket; "We were attacked. Bigtime. We had a massive flood of UDP packets coming in to our IP, basically eating away all bandwidth to the box…So, basically a massive-scale DDoS. That's nice." Please note, that as a result of the attack "Peter DeSantis, vice president of Amazon Elastic Compute Cloud (EC2), said that they were definitely taking this lesson about the tardy detection of Bitbucket.org's problem to heart. He said, from Amazon's perspective, the black eye from that smarted, and the company would be changing its customer service playbook and network policies to prevent a reoccurrence."[34]

This encounter can lead the reader to think that the advice from ENISA is not entirely accurate, and this would not be fair. The likelihood is that if BitBucket were using internally provisioned services, and not Amazon, then perhaps their service may have been unavailable for longer, and their ability to withstand traffic

not as resilient. Therefore, the ability of the provider (of course this depends on the provider) to withstand a DDoS attack should be more than that of an internally provisioned service. This is an entirely case-dependent statement. However, one risk that cloud end customers should certainly consider is the noisy neighbor concept. This particular point goes against the universally believed concept that the cloud reduces the threat of a DDoS, whereby the use of a cloud provisioned service means that the customer is sharing resources with other customers. This scenario was documented[35] by Rich Bolstridge of Akmai Technologies, who provided three cases in which the shared services approach negatively impacted cloud customers;

- Case 1: DDoS attack against Brazilian bank subsidiary
  An attack targeting the home page of a Brazilian bank's Brazilian site. However, as the Brazilian Web site utilized a shared network infrastructure, the US banking site was also negatively impacted. Somewhat ironically, the bank had invested in DDoS mitigation for the US Web site, but failed to recognize the threat of the shared network infrastructure.
- Case 2: DDoS attack against a Luxembourg customer of a US exchange
  A US exchange had a market data service used by a customer in Luxembourg to serve its clients. The application, however, came under attack, causing it to be unavailable. The service, however, was also used by the exchange's main applications for desktop clients in the United States, which ultimately failed.
- Case 3: DDoS attack against US subsidiary of European bank
  A DDoS attack against the domain name servers of a large regional bank in the United States resulted in the Web site for the bank across three continents also being impacted.

Therefore, to summarize, the threat of a DOS attack can impact not only internally provisioned services, but also that of CSPs. While the general view that the cloud computing provider should provide a greater ability to withstand such attacks, the probability of a DDoS attack will increase when using shared resources with multiple customers. To summarize, the risk will be reduced if the cloud provider has implemented the appropriate controls to withstand such an attack, but the number of attempts (some that may be successful) will increase.

Malicious Insiders

There exist multitudes of varying statistics attempting to quantify the threat of malicious insiders. While the exact number can be, and regularly is, argued, there is no question that the risk does exist; the only question is how big the threat is. According to the CERT Insider Threat Centre,[36] the malicious insider can be defined as

> A malicious insider threat to an organization is a current or former employee, contractor, or other business partner who has or had authorized access to an organization's network, system, or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organization's information or information systems.

To be clear, this particular threat refers to the conscious effort to compromise information, or an information system. While of course this threat can affect individual organizations, within a cloud computing environment there are three types of cloud-related insider threats based upon the CERT Program,[37] Software Engineering Institute Carnegie Mellon University:

- Rogue administrator: An individual employed by the CSP who undertakes an action to affect the confidentiality, integrity, or availability of the service. Examples include theft of sensitive information or sabotage. Of course, there exist multiple examples of rogue administrators undertaking actions that circumvent the policy of their employer. In addition, such actions can exist even after the employee has left the organization, as was the case reported by InformationWeek.[38] The case refers to a former employee of Gucci who was accused of maintaining a Virtual Private Network (VPN) token, and using it to access the network of his former employer and "deleting virtual servers, taking a storage area network offline, and deleting mailboxes from the corporate email server." Within a cloud environment, CERT identifies four levels of administrators, each with differing levels of access, and subsequently the potential impact if they are malicious. The levels of access are, however, hierarchical, where the top-level administrators (hosting company administrators) have the greatest level of access.
  - Hosting company administrator: Has the highest level of access and therefore could cause the greatest impact such as updating the drivers of the virtual machines to compromise the images. Moreover, can implement network taps to perform man-in-the-middle attacks on all hosted systems.
  - Virtual image administrator: Could create alternate images outside of the authorized baseline, and that report they align with such baseline. Could also potentially copy virtual machines/disks, or modify individual instances of a virtual machine in a cloud so that only some of the cloud behaves the wrong way.
  - System administrators: Have the ability to conduct operating system attacks, and could update the virtual machine drivers to vulnerable instances.

  - Application administrators: Have the ability to copy application data, edit the configuration of applications, potentially can gain control of the hosting platform.
- Exploit weaknesses introduced by use of the cloud: The use of cloud computing introduces vulnerabilities that the malicious insider will look to exploit. One particular example of these vulnerabilities includes a difference between the access control model between the local system and the cloud-based system. Also, another threat proposed is the replication lag exploit. In this example, the cloud environment potentially includes multiple systems that synchronize important information (such as pricing, for example). However, due to network latency, or that servers are located in different geographic locations, the replication of these data may take some time. Therefore, while the cloud environment removes the single point of failure issue compared with a single server located on premise, by understanding the replication lag issue the insider may be able to purchase items for less than the corporate agreed price. The example provided by CERT is as follows:
  - Company has server A that is authoritative for all pricing.
  - Server A replicates prices to servers B1 and B2 that have 1 and 2 s of latency, respectively.
  - Server B1 replicates prices to servers C1 and C2, these have 2 s of latency each.
  - Server B2 replicates prices to server C3 with 4 s of latency.

The attacker wishes to buy a $20 item for $10. Therefore, when a price change is scheduled, they will apply a false notice so the price is actually $10 sending to C3. Then by timing the purchase before the correct price is applied they could remove evidence of the incorrect price, and potentially evidence they circumvented the integrity of the system.

- Using the cloud to conduct nefarious activity: This example relates to a malicious insider who utilizes cloud services to conduct attacks against his or her employer. Indeed, research published by TechTarget[39] suggests that the lack of appropriate fraud detection capability within CSPs allows criminals to undertake activities on commercial providers without such activity being detected. The acquisition of services can be conducted using stolen credit cards, or as indicated earlier through account hijacking.

Abuse of Cloud Services

We briefly touched on this subject in the earlier paragraph related to malicious insiders; however, the abuse of cloud services extends beyond malicious insiders and potentially allows cybercriminals the ability to utilize such services for criminal gain. There are multiple ways in which cloud services can be used for malicious purposes.

Resource Intensive Operations—Cracking Passwords

There is no question, that for the malicious actor their job is considerably easier if their intended victims use very simple passwords. Remarkably, analysis from the breach of Adobe Systems found the most common password used was 123456,[40] and was used by 1.9 million users. Should the target not use a simple password, then the attacker will be faced with alternate means to crack a user password, which has in fact become considerably easier (or rather cheaper) with cloud computing. In particular, using the computing resources to undertake a brute force attack (repeatedly trying different passwords to find the right one), is made considerably more efficient with cloud. There have been many demonstrations highlighting the use of cloud computing to brute force passwords; in 2010, for example, German hacker Thomas Roth was reported[41] to have used AWS to have cracked passwords encrypted within a Secure Hashing Algorithm Hash. By using Amazon's graphics processing unit (GPU) instances, Roth was able crack hashes that contained passwords between one and six characters in 49 min, with the GPU instances costing $2.10 per hour at the time. GPU instances are a product designed for high-performance computing jobs that Roth describes as "known to be the best hardware accelerator for cracking passwords."

Other examples of brute forcing passwords via cloud computing include wireless network passwords; for example, in 2009, the service known as WPA Cracker was reported[42] to have checked a password against 135 million entries in 20 min for only $34. Wireless network and SHA1 passwords are, however, only a tip of the iceberg. There exists a multitude of services available offering computing resources for resource-intensive operations to brute force passwords over a cloud service. As we saw in the two earlier examples, some simply provide the core resources, but in other examples there are dedicated companies offering a simple GUI and SaaS service dedicated for the sole purpose of cracking passwords. There are also toolkits that give the potential hacker an interface into cloud resources for the purpose of using cloud services for brute forcing passwords. It is, however, worth noting that the use of commercial cloud services to crack passwords without authorization will breach the acceptable use policy for the provider.

Hosting Malicious Content

There are two elements regarding the hosting of malicious content: (1) using providers that have no issues regarding any (or almost any) hosting malicious content and (2) using providers to host malicious content circumventing the CSPs acceptable use policy.

The concept of using a provider that offers lenient acceptable use policies is known as BulletProof hosting. Such services have been used by malicious actors (e.g., those hosting content such as pornography or sending spam)

for some time. However, the challenge of using such services is that they are often blacklisted by security providers and therefore the emerging trend for many malicious actors is to utilize commercial hosting services that are not blacklisted, and subsequently then able to reach all intended victims without security tools blocking the sending domains. This trend poses a challenge to commercial cloud providers as the implications of hosting malicious content could result in their operations being blacklisted, which will be to the detriment of existing customers, and ultimately impact profitability. There is also the potential for law enforcement action as we saw with MegaUpload discussed in Chapter 1 that may lead to seizure of equipment. The challenge of course will be for the cloud provider to ensure the customer is not using services for malicious purposes, this will be a challenge because signing up is automated without the need to interact with any human operator, all that is required is a credit card.

Subsequently, providers will need to establish mechanisms to determine whether fraudulent activities are taking place, but according to John Rowell of Dimension Data, "[There are] service providers that…do not have adequate fraud measures in place, and they have to be losing insane amounts of money on it. It's got to have an immense impact to their profitability as well as just the health and cleanliness of their platform.[30] However, the challenge will be the level of scrutiny toward customer operations in the provisioned service; one of the biggest selling features for the use of cloud is its ease of use. Indeed, many providers make establishing their services so simple that in many cases the IT Departments are not even aware (known as Shadow IT). By adding more checks and oversight there is the potential for customers to not see the service as simple, and to migrate to providers that may not be as onerous in their oversight. Therefore, a balance is absolutely necessary between fraud detection and ease of use.

Due Diligence

Migrating to the cloud is a simple and effective way to transfer existing workloads to an external party without the need to rush out and buy new hardware, install the operating system, hire administrators, etc. Indeed, the cloud is one of the most effective mechanisms to outsource the work for an organization; however, sadly the risk cannot be outsourced so easily as the failure to undertake appropriate due diligence will leave the end customer liable.

In particular, where personally identifiable information (PII) is hosted, there will likely be data protection legislation that demands due diligence when using third parties to host such data. In the United Kingdom, this is documented within the Data Protection Act, under Principle 7.[43] Under this principle, it demands that the end customer undertakes appropriate due diligence

to ensure that the data processor (in this case the cloud provider) has the appropriate controls in place to protect the data. Furthermore, in the United Kingdom, the end customer (or data controller) will also need to ensure that under principle 8 the data are not transferred outside of the European Economic Area. While of course there are exclusions to using providers outside of this area, such as Safe Harbor, what these two examples demonstrate is that using cloud computing requires due diligence to ensure compliance against regulatory requirements, and that the risk cannot be outsourced because the end customer remains liable.

The intention may be to leverage certifications that the provider may boast to demonstrate security; however, regulators in many countries across the world have dictated this to be insufficient. The UK Information Commissioners Office recently clarified this position:

> The Data Protection Act does not stop the overseas transfer of personal data, but it does require that it is protected adequately wherever it is located and whoever is processing it, this includes if it is being stored in the cloud outside of the UK. While any scheme aimed at ensuring people's information is adequately protected in line with an organisation's requirements under the Act is to be welcomed, organisations thinking of using CSPs must understand that they are still responsible for the safety of that data. Just because their CSP is registered with such a scheme, does not absolve the organisation who collected the data of their legal responsibilities.[44]

This is only one such example from regulators reminding end customers of their obligation in undertaking appropriate due diligence when acting as data controllers. In the European Union, the Article 29 Working party published guidance[45] outlining the obligations of cloud computing customers in ensuring that providers adhere with data-protection rules. The Working Party, a committee comprising representatives from the 27 data protection authorities within the EU member states, also confirmed that cloud computing poses risks to data security, such as "loss of governance, insecure or incomplete data deletion, and insufficient audit trails or isolation failures."

Due diligence is therefore quite obviously imperative, and while the desire may be to adopt cloud computing just as quickly as the sign-up process allows, it is important to note the obligations to undertake a sufficient assessment of the risks associated with migrating to a third party whether in the cloud or not.

Shared Technology Vulnerabilities

One of the many benefits of cloud computing is the ability to leverage economies of scales by sharing resources across multiple customers. However, this very benefit also represents a significant weakness as this demands strong

isolation to ensure that a vulnerability or misconfiguration in one instance does not affect other instances, and ultimately the entire cloud.

The types of risks associated with this category includes the failure of mechanisms associated with the storage memory, routing, and even reputation between different tenants of the shared infrastructure (e.g., so-called guest-hopping attacks, SQL injection attacks exposing multiple customers' data stored in the same table, and side channel attacks).

According to ENISA, the "likelihood (probability) of this incident scenario depends on the cloud model considered; it is likely to be low for private clouds and higher (medium) in the case of public clouds."

The impact can be a loss of valuable or sensitive data, reputation damage, and service interruption for cloud providers and their clients.

ADDITIONAL CLOUD THREATS

Of course, the notorious nine are the top threats as seen by experts. Potential customers should ensure that they undertake a comprehensive risk assessment to determine what "other" threats may exist. An excellent source is the ENISA Cloud Computing Security Risk Assessment[46]; this document outlines the key risks associated with cloud computing (with contributions from coauthor Raj Samani, and edited by Daniele Catteddu now of the CSA). The document outlines the following areas as the key areas of risk for cloud computing:

- Loss of governance: Where the use of cloud computing results in the end customer handing control to the CSP.
- Lock-in: Where it becomes difficult for the end customer to migrate from their cloud provider.
- Isolation failure: Relates to the risk of a failure in mechanisms that are intended to separate storage, memory, routing and even reputation between different tenants.
- Compliance risks: Migration to the cloud may result in compliance failure for the potential cloud customer, for example, the migration of personally identifiable data outside of specific regions.
- Management interface compromise: As the interface to the cloud service is externally accessible (via the Internet) and provides access to large sets of resources, the risk is therefore increased.
- Data protection: This relates to the due diligence threat as defined within the notorious nine, as it may be difficult for the end customer to "effectively check the data handling practices of the cloud provider and thus to be sure that the data are handled in a lawful way."

- Insecure or incomplete data deletion: This threat relates to the deletion of a cloud resource, particularly as it may not be possible to entirely delete the data. This may either be because the physical disk to be destroyed may store data from other clients or the additional copies are not available.
- Malicious insider: As covered under the notorious nine.

Cloud computing presents a hugely efficient and potentially cheaper option for many organizations. However, as this chapter indicates the risk to customers will always be theirs, and therefore it is imperative to undertake a thorough risk assessment to ensure that the appropriate controls are in place. As said previously, "if you are going to put your eggs in the cloud basket, you have to have the mechanisms in place to really solidify that environment."

END NOTES

1. European Network Information Security Agency, *Glossary* [cited July 2014]. Available from: https://www.enisa.europa.eu/activities/risk-management/current-risk/risk-management-inventory/glossary.
2. McCaughey and Ayers, *CyberActivism: Online Activism in Theory and Practice*, (February 2003), [cited July 2014]. Available from: http://www.amazon.com/Cyberactivism-Online-Activism-Theory-Practice/dp/0415943205.
3. Fred Cohen, *Computer Viruses: Theory and Experiments*, [cited July 2014]. Available from: http://web.eecs.umich.edu/~aprakash/eecs588/handouts/cohen-viruses.html.
4. F-Secure, *Cascade* [cited July 2014]. Available from: http://www.f-secure.com/v-descs/cascade.shtml.
5. F-Secure, *Ping-Pong* [cited July 2014]. Available from: http://www.f-secure.com/v-descs/pingpong.shtml.
6. Ralph Langner, Langer. *To Kill a Centrifuge* (November 2013), [cited July 2014]. Available from: http://www.langner.com/en/wp-content/uploads/2013/11/To-kill-a-centrifuge.pdf.
7. Ponemon Institute, *2013 Cost of Data Breach study* (May 2013), [cited October 2013]. Available from: https://www4.symantec.com/mktginfo/whitepaper/053013_GL_NA_WP_Ponemon-2013-Cost-of-a-Data-Breach-Report_daiNA_cta72382.pdf.
8. Beazley, *Beazley Small business spotlight; Data Breach costs* (January 2012), [cited October 2013]. Available from: https://www.beazley.com/.../Data%20Breach%20costs%20June%202012 .
9. Bloomberg Law, *Complicated Compliance: State Data Breach Notification Laws* Cited October 2013. Available from: http://about.bloomberglaw.com/practitioner-contributions/complicated-compliance-state-r.
10. Official Journal of the European Union, *Commission Regulation (EU) No 611/2013* (August 2013), [cited October 2013]. Available from: http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ: L:2013:173:0002:0008:EN: PDF.
11. The Christian Science Monitor, *Data Theft: Top 5 most expensive data breaches* Cited October 2013. Available from: http://www.csmonitor.com/Business/2011/0504/Data-theft-Top-5-most-expensive-data-breaches/3.-TJX-256-million-or-more.
12. Wired.com, *Data Breach Will Cost TJX $1.7B, Security Firm Estimates* [cited October 2013]. Available from: http://www.wired.com/threatlevel/2007/03/data_breach_wil/.

13. Jim Lewis and Stewart Baker, CSIS, *Economic Impact of Cybercrime* (July 2014), [cited July 2014]. Available from: http://www.mcafee.com/hk/resources/reports/rp-economic-impact-cybercrime2.pdf
14. Wired.com, *IBM Cloud Storage Partner Nirvanix Files for Bankruptcy* (October 2013), [cited October 2013]. Available from: http://www.wired.com/wiredenterprise/2013/10/nirvanix-bankrupt/.
15. NetworkComputing.Com, *The Nirvanix Failure: Villains, Heroes and Lessons* (October 2013), [cited October 2013]. Available from: http://www.networkcomputing.com/storage-networking-management/the-nirvanix-failure-villains-heroes-and/240162664.
16. Molitor Legal, *New Luxembourg law on the right to claim back data from bankrupt IT and cloud service providers* (July 2013), [cited October 2013]. Available from: http://www.molitorlegal.lu/news/new-luxembourg-law-right-claim-back-data-bankrupt-it-and-cloud-services-providers-0.
17. GigaOm. *Power Outages are the Most Pervasive Reasons for Cloud Outages* (March 2013), [cited October 2013]. Available from: http://research.gigaom.com/2013/03/power-outages-are-the-most-pervasive-reasons-for-cloud-outages/.
18. ZDNet, *AWS cloud Accidentally deletes customer data* (August 2010), [cited October 2013]. Available from: http://www.zdnet.com/aws-cloud-accidentally-deletes-customer-data-3040093665/.
19. Beth Pariseau, TechTarget.com, *Code Spaces goes dark after AWS cloud security hack* (June 2014), [cited July 2014]. Available from: http://searchaws.techtarget.com/news/2240223024/Code-Spaces-goes-dark-after-AWS-cloud-security-hack.
20. Ibid.
21. BBC.com, *Burger King Twitter account hacked with McDonalds logo* (February 2013) [cited November 2013]. Available from: http://www.bbc.co.uk/news/world-us-canada-21500175.
22. Wired.com, *How Apple and Amazon Security Flaws Led to My Epic Hacking* (August 2012), [cited November 2013]. Available from: http://www.wired.com/gadgetlab/2012/08/apple-amazon-mat-honan-hacking/.
23. LifeHacker ITPro, *Why Cloud Services are so easy to hack* (February 2013), [cited November 2013]. Available from: http://www.lifehacker.com.au/2013/02/why-cloud-services-are-so-easy-to-hack/.
24. Warwick Ashford, ComputerWeekly.com, *FBI warns of increased spear phishing attacks* (July 2013), [cited November 2013]. Available from: http://www.computerweekly.com/news/2240187487/FBI-warns-of-increased-spear-phishing-attacks.
25. MindTools, *Caldini's six principles of influence*, [Cited November 2013]. Available from: http://www.mindtools.com/pages/article/six-principles-influence.htm.
26. Mashery, *APIs: the key to a thriving cloud* (February 2009), [cited November 2013] Available from: http://readwrite.com/2009/02/26/mashery_apis_the_key_to_a_thriving_cloud.
27. *Amazon Elastic Compute Cloud: API Reference* (October 2013), [cited November 2013]. Available from: http://awsdocs.s3.amazonaws.com/EC2/latest/ec2-api.pdf.
28. Rob Lemos, DarkReading.com, *Insecure API Implementations Threaten Cloud* (April 2012), [cited November 2013] Available from: http://www.darkreading.com/cloud/insecure-api-implementations-threaten-cl/232900809.
29. Raj Samani and Francois Paget, *CyberCrime Exposed: Cybercrime-as-a-Service* (September 2013), [cited November 2013]. Available from: http://www.mcafee.com/uk/resources/white-papers/wp-cybercrime-exposed.pdf.
30. Prolexic, *Prolexic Quarterly DDOS Threat Report. Q3 2013* [cited November 2013]. Available from: http://www.prolexic.com/knowledge-center-dos-and-ddos-attack-reports.html.
31. ENISA, *Critical Cloud Computing: A CIIP Perspective on cloud computing services* (December 2012), [cited November 201s3]. Available from: https://www.enisa.europa.eu/activities/Resilience-and-CIIP/cloud-computing/critical-cloud-computing.

32. Rackspace, *DDOS mitigation, prevention and Network security tool* [cited November 2013]. Available from: http://www.rackspace.com/managed_hosting/services/security/ddosmitigation/.
33. Cade Metz, The Register. *DDoS attack rains down on Amazon cloud* (October 2009), [cited November 2013]. Available from: http://www.theregister.co.uk/2009/10/05/amazon_bit-bucket_outage/.
34. Carl Brooks, TechTarget, Amazon EC2 attack prompts customer support changes. October 2009 [cited November 2013. Available from: http://searchcloudcomputing.techtarget.com/news/1371090/Amazon-EC2-attack-prompts-customer-support-changes.
35. Rich Bolstridge, Asia Cloud Forum, *Risks underlying shared services DDoS* (August 2013), [cited November 2013]. Available from: http://www.asiacloudforum.com/content/risks-underlying-shared-services-ddos.
36. CERT, *The CERT Insider Threat Center* [cited November 2013]. Available from: http://www.cert.org/insider_threat/.
37. CERT. Insider Threats to Cloud Computing: Directions for New Research Challenges. [cited November 2013]. Available from: www.cert.org/archive/pdf/CERT_cloud_insiders.pdf.
38. Matthew Schwarz, *Fired employee indicted for hacking Gucci Network* (May 2011), [cited November 2013]. Available from: http://www.informationweek.com/infrastructure/networking/fired-employee-indicted-for-hacking-gucci-network/d/d-id/1097007?.
39. Jessica Scarpati, TechTarget.com, *For cloud providers, fraud detection is integral part of business plan* [cited November 2013]. Available from: http://searchcloudprovider.techtarget.com/feature/For-cloud-providers-fraud-detection-is-integral-part-of-business-plan.
40. BBC News, *Analysis reveals popular Adobe passwords*. (November 2013), [cited November 2013]. Available from: http://www.bbc.co.uk/news/technology-24821528.
41. Jack Clark, ZDNet, *Hacker uses cloud computing to crack passwords* (November 2010), [cited November 2013]. Available from: http://www.zdnet.com/hacker-uses-cloud-computing-to-crack-passwords-4010021067/.
42. Dan Goodin.*The Register Service cracks passwords from the cloud.*(December 2009), [November 2013] Available from:.http://www.theregister.co.uk/2009/12/07/cloud_based_password_cracking/.
43. Louise Kidney. *The Guardian. Navigating a tricky airspace: information governance in the cloud* (July 2011), [cited November 2013]. Available from: http://www.theguardian.com/local-government-network/2011/jul/14/information-governance-cloud.
44. Out-Law.com, *Certifications of cloud provider services welcome, but users cannot rely on them for own data protection compliance, says ICO* [cited November 2013]. http://www.out-law.com/en/articles/2012/july/certifications-of-cloud-provider-services-welcome-but-users-cannot-rely-on-them-for-own-data-protection-compliance-says-ico/.
45. Article 29 Working Party, *Opinion 05/2012 on Cloud Computing* (May 2012), [cited November 2013]. Available from: http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2012/wp196_en.pdf.
46. European Network Information Security Agency, *Cloud Computing Security Risk Assessment* (2009), [cited July 2014]. Available from: https://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment.

This page intentionally left blank