Chapter 12
Privacy Penetration Testing: How to Establish Trust in Your Cloud Provider

Christian W. Probst, M. Angela Sasse, Wolter Pieters, Trajce Dimkov, Erik Luysterborg and Michel Arnaud

12.1 Introduction

In the age of cloud computing, IT infrastructure becomes virtualised, and all aspects of the stack of hardware, platform and software take the form of services. Moreover, these services can be offered by different organisations, which may purchase their capacity from again different organisations. The complexity of who owns, possesses, controls and uses information increases dramatically (Floridi and Turilli 2011).

In this sense, cloud computing forms an instance of the broader concept of de-perimeterisation (Jericho Forum 2005; van Cleeff and Wieringa 2009). De-perimeterisation denotes the disappearing of boundaries around the IT infrastructure of organisations. Whereas information security was previously conceived as separating the trusted *inside* from the untrusted *outside*, such a clear delineation is not possible anymore. The question is what can take its place, i.e., how re-perimeterisation would be possible.

Of course, security has never been completely based on a single perimeter. People working for an organisation would leave the perimeter in their private lives, enabling information to cross the boundary between the organisation and its surroundings. This has become more prominent with the use of mobile devices in the workplace, or "*bring-your-own-device*". Also, the inside of the organisation might not have been completely trusted, as there would always be a chance that people inside the organisation would misuse their credentials for their own benefit. This so-called insider threat has become a substantial area of research (Hunker and Probst 2011).

In this sense, it is not surprising that the notion of a security perimeter has broken down. Developments like cloud computing have only made more explicit that such a concept is untenable, and accelerated the emergence of different security architectures. The original idea that the perimeter is as close to the data as possible (data-level security) cannot be the only solution in cloud-computing environments, as full encrypted processing is not feasible (Pieters 2011b). Instead, we are now

M. Arnaud (✉)
Université Paris Ouest Nanterre La Défense, Paris, France
e-mail: Michel.arnaud@u-paris10.fr

S. Gutwirth et al. (eds.), *European Data Protection: In Good Health?*,
DOI 10.1007/978-94-007-2903-2_12, © Springer Science+Business Media B.V. 2012

looking at complicated re-perimeterisation processes, where different layers of protection, in different domains (digital, physical and social) need to be aligned for all organisations involved (Abrams and Bailey 1995). For example, to prevent confidential data from leaving an organisation, we would have to prevent outsiders from gaining access to the building, prevent employees from taking the data home, and check outgoing e-mail traffic. We would need to perform the same checks at the cloud provider. Worse, some attacks may combine weaknesses in different domains to circumvent the carefully crafted multi-perimeters, or "virtual perimeters".

Especially, private and sensitive data requires special protection when being stored or processed in a cloud infrastructure. Organisations want to have some confidence that the benefits of moving to a cloud environment outweigh the risks. People accept risk and uncertainty in exchange for an expected benefit, but as the cloud infrastructure is not transparent for the user, this requires trust in the providers and their security policies. The more sensitive the data in question is, the better and stronger guarantees are required when the data is being stored or processed in a cloud infrastructure—or more trust. Because of the difficulties of cross-organisational security assessment, this trust may be hard to justify.

The questions are thus how to empower cloud users to develop trust in cloud infrastructures, how to align security policies to form a reliable perimeter within one's own organisation, and how to trust and/or verify the security measures in place in other organisations in the cloud network? After discussing security challenges in the cloud in the next section, we first look at the question of trust into cloud infrastructures in Sect. 1. This leads to the suggestion of PuPPeT, a public privacy penetration-testing agency, in Sect. 4. In this section we also discuss how to test security policies and how to verify security measures; since the suggested agency will have to act across organisations, we introduce cross-domain methods for security testing and for modelling organisational security. The present contribution brings together these different factors in securing data in the age of the cloud, for which open questions are discussed in Sect. 1, followed by conclusions in the final section.

To simplify discussion, we will in the following use the term "cloud operator" for organisations offering a cloud infrastructure, the term "cloud user" for organisations running cloud applications operating on their customers' data, and "data owner" for organisations and individuals using cloud applications.

12.2 Security Challenges in the Cloud

When considering the security impact of adopting a cloud-computing environment, opinions regarding the exact nature of the "cloud threat" differ quite substantially.

Some state that there is really nothing new under the sun, and that, especially with respect to a private cloud environment, the security-related issues are the same as those existing in a "non-cloud" enterprise today (Robinson et al. 2011). Some state that, because of the nature of the cloud itself (i.e., difference in scale of entry points to potentially be subject to attacks), the security risks are clearly of a different nature

or type (Jansen and Grance 2011). Others talk more about a difference in scale, not in type of threat (Mitra and Mallik 2010).

However, when the exact nature of challenges in the cloud needs to be quantified, there is one thing almost everyone agrees upon: cloud computing does pose a number of real challenges in terms of security, privacy, and trust, both for cloud providers and cloud users.

Indeed, because cloud computing grew out of an amalgamation of technologies, e.g., virtualisation, Web 2.0 and service-oriented architecture, the related security, privacy and trust issues should be viewed as already known problems in a new setting. However, it is the importance of their combined effect that should not be underestimated. Therefore, in order to propose an appropriate response to the threats related to cloud computing, it is necessary to first understand and define properly the exact challenges in this regard.

In general, when people talk about ensuring security, they refer to integrity, access and availability of data, as well as confidentiality and non-repudiation. Privacy, on the other hand, embraces much more; it is often seen as primarily being about compliance with applicable data protection laws and regulations relating to, e.g., data transfer or location, purpose of processing and data subject rights of access and control. But privacy is much more than data protection, for example, it is also about observable behaviour and anonymity. One could say that data protection only provides the means of protecting privacy, but they need to be used in the right way.

When addressing privacy in the cloud, two aspects must be distinguished: on the one hand, applications running in the cloud should protect the privacy of the data they process; on the other hand, cloud providers should protect the data that is stored or that is processed on their infrastructure. These requirements are not new; the first one is the same as privacy protection in every other application, and the second one is the same as for regular hosting companies. In cloud computing, the risk just is amplified by the multitude of outsourced components and, for example, the uncertainty about location of data.

Therefore, the above concepts need to be further refined and clarified in order to be fully understandable in the cloud context. We propose to add the following clarifications to the existing concepts. Please note that some of these can apply many times but for sake of clarity, we have listed them only once. They are also valid for both cloud users as well as cloud providers.

12.2.1 Security Challenges and Granularity

Security challenges in relation to the cloud environment can (non-exhaustively) be categorised as lack of control on the provider's resources, increased exposure of internal infrastructure via new technologies/interfaces, insufficient adaptation of application/platform security and development lifecycle, unclear ownership of security tasks and lack of cloud specific security standards, to list some.

The above demonstrates that the main security challenge can be translated into one of *granularity*. In other words, in order to understand the full scope of the cloud security challenges, one needs to identify at which level of granularity one can identify the relevant security threats. This will largely depend on criteria such as, e.g., the type of data concerned, the scale of outsourcing, the number of third parties involved, the architecture/technology used, etc. Another important factor is the extent in which cloud providers offer customised services as opposed to standardised ones. The customised approach will allow to better master the security issues in a more adapted manner, also addressing the issue of attribution of responsibilities between the different parties involved.

12.2.2 Privacy and Accountability

Data privacy generally refers to a bundle of legal/contractual rights and obligations related to the collection, processing and sharing (transferring) of personal information. Although several definitions exist, one of the most comprehensive definitions of personal information is included in the so-called 1995 European Data Protection Directive 95/46/EC:

> Personal information is any information relating to an identified or identifiable natural person ('data subject'); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity.

One could argue that this definition means for cloud computing that most data stored in the cloud will be personal information, resulting in the above-mentioned directive being applicable. This means that somebody in the conglomerate of maybe several cloud users and cloud providers collectively processing the data is responsible for protecting its privacy. However, this responsibility may be hard to assign in practice.

Typical privacy issues that are mentioned in connection with a cloud environment are data localisation and applicable law, data breach/leakage issues and data transfers. Clearly different concerns exist when outsourcing customer data to the cloud versus, for example, outsourcing an organisation's business records.

Even though the current privacy legislative framework is far from ideal, and even though often very divergent privacy laws and regulations exist, rendering difficult the handling of data in the cloud, in reality all of these hurdles are not insurmountable. They can indeed be summarised in the challenge of "accountability". Given the volume and/or location of the different cloud service providers, it will be crucial to establish clear rules regarding the (contractual) responsibilities for privacy compliance by any of the parties in the (cloud) chain. As such, and using the terminology of the data protection regulations, clearly identifying the data flow as well as the roles of each data controller, processor and sub-processor, and where they/the data are located/restricted to, will go a long way in ensuring compliance with the applicable privacy laws and (contractual) rules.

12.2.3 Trust and Transparency

Finally, one of the most difficult challenges in cloud computing is to enable customers to develop trust in a cloud environment. In a cloud environment, one of the key questions from individuals and companies is: can I trust the cloud with my data? To answer this question, we need to first examine what "trusting" the cloud means. We only need trust in situations with risk and uncertainty—people accept risk and uncertainty in exchange for an expected benefit.

With cloud computing, the expected benefit for the user of a cloud computing service, e.g., a medium-sized enterprise, is reducing cost and increasing reliability. The risks associated with cloud computing include availability and integrity (*Will I always be able to access the data when I need them, and will they be the data I stored?*) and confidentiality (*Might someone working at the cloud provider or another client get access to my customer's personal data?*). Uncertainties surrounding cloud computing include questions such as whether the provider will do what they promise (such as not transferring the data outside the EU without explicit consent), and whether there is any redress and restitution if they fail to deliver.

When deciding whether to trust someone, humans usually consider two qualities: the trustee's *ability* and *motivation* to deliver their side of the transaction. In terms of ability, cloud providers argue that data storage and processing is their core competence, which means they are better equipped to keep data secure than most of their customers—*trust us, we're the professionals*. In terms of motivation, Pearson and Charlesworth (2009) argue that cloud computing providers should be highly motivated to safeguard their customers' data, since their reputation depends on it. In a system where customers and providers can trust each other to deliver what their transaction partner expects, all parties can expect to benefit (Riegelsberger et al. 2005). So—is it time to stop worrying, and learn to trust the cloud?

Taking the above into account, the trust challenge can be summarised in one word: transparency. Indeed, establishing a level of trust or confidence about a cloud environment largely depends on the ability of the cloud provider to demonstrate clearly and upfront (and on a regular basis thereafter) the provision of security and privacy controls required to protects the client's data and applications, as well as any evidence to demonstrate the effectiveness of such controls.

12.3 A Pragmatic Trust-based Approach

To pick up our earlier question—maybe it *is* time to stop worrying and trust the cloud. But given that enterprises turn to cloud computing to save money, it makes sense for cloud providers to feel compelled to compete on price. Such competition could lead to cloud providers trying to save on parts of the services that are regarded as non-essential. Whilst customers would note problems with availability in day-to-day usage, effective security and privacy protection manifest themself as *absence* of security breaches. A cut in expenditure on protecting security and privacy does not necessarily lead to a breach—or at least not immediately. So, there is a likely

temptation to save on this protection to offer more competitive prices. Once there has been a breach, a cloud provider's reputation will suffer, but by then, for the cloud user that entrusted its data to that cloud provider, the damage to the enterprise, and its customers, is done. So the question is—how can cloud users tell apart the cloud service providers that take good care of their data and safeguard their customers' privacy, and those that do not?

Unfortunately, there are no reliable trust signals that would allow cloud users to infer whether cloud providers are acting as they promise (Riegelsberger et al. 2005). This means that, rather than trusting cloud providers, they have to put assurance mechanisms in place, such as contracts, inspection of facilities, and testing the security and privacy protection measures in place (Flechais et al. 2005). However, such assurance mechanisms introduce cost for both the cloud user and the cloud provider—meaning neither can reap the full financial benefits of a trust relationship.

So, the answer to our earlier question is that we can learn to trust the cloud, but not without investing in the necessary assurance mechanisms. To be effective, these mechanisms need to address the challenges introduced in the previous section.

Adopting a granular approach means demanding a more customised service adapted to the "sensitivity" level of the data processed or services requested. Companies should not only employ specific security controls to verify the correct functioning of the various subsystems in the cloud environment, they should also ensure strong and adapted security management practices adapted to their changed role.

Current privacy laws and regulations are what they are: they are far from ideal and should be further improved and more harmonised. Meanwhile, both cloud users and cloud providers need to comply with this current legal framework. In order to achieve this, both need to clearly attribute accountability to each of the intermediaries for compliance with all relevant (contractual and legal) rules related to, e.g., location of data, data transfer, data usage and data breach procedures in relation to its role and responsibility in the cloud "chain".

Finally, cloud providers must be able to demonstrate in a clear and transparent way that they implement the above-mentioned assurances and approach. At the same time, the cloud user must accept that its role (especially that of its IT department) has changed and entails certain governance responsibilities as well. As such, a trustworthy relationship can be created, not by assuring that everything can be guaranteed in a bullet-proof fashion, but by ensuring that a flexible framework exists whereby data will be protected in a manner consistent with agreed upon policies, procedures and contractual arrangements and that adequate redress or alert procedures are in place. To support this process, we suggest PuPPeT, a public privacy penetration-testing agency.

12.4 PuPPeT—A Public Privacy Penetration-Testing Agency

It is obvious that the principles or concepts of granularity, accountability and transparency apply to any of the above-mentioned security, privacy or trust challenges and are highly intertwined. We believe that they are key to ensuring a properly config-

ured, well balanced and secure cloud environment, thereby allowing both the cloud user as well as the cloud provider to fully exploit the potential benefits of the cloud. They also illustrate that securing the cloud is not only a matter of mere technology, but also a combination of people, processes and technology.

Institutional safeguards, such as regulation, could offer protection, but regulation always lags behind technology, and has not caught up with cloud computing (Pearson and Charlesworth 2009). Additionally, cloud computing is an international business, which means that it is often beyond the regulator's reach. One approach is to rely on self-regulation of markets (Hirsch 2011).

Pearson and Charlesworth make a compelling argument that the solution for this problem is accountability of the cloud provider to their customer enterprises (cloud users). In the case of privacy, the elements for accountability for privacy are (Pearson and Charlesworth 2009):

1. *Transparency*: informing data owners how their data is handled in the cloud, and who has responsibility for which parts of processing;
2. *Assurance*: through privacy policies;
3. *Responsibility*: must be clearly allocated, and taken actively by the cloud provider (rather than relying on enforcement by regulators or cloud users); and
4. *Policy compliance*: rather than following the letter of policies, cloud providers must strive to achieve a proportionate and responsive process for reacting to context-dependent privacy risks.

Pearson and Charlesworth further suggest that these privacy-protecting controls should be built into different aspects of the business process, and cloud users and cloud providers must work together over time to develop robust routine protection of privacy in the cloud. This approach mixes trust and assurance, but remains very much on the assurance side, meaning that the cost for both sides remains substantial.

To overcome this, we suggest PuPPeT, a public privacy penetration-testing agency. We envision PuPPeT to be a more economic alternative to the process sketched above. The agency would award a trust symbol for cloud computing providers that cloud users and data owners can use to make an informed decision about whether or not to trust a cloud provider (Balboni 2009). To award the trust symbol, the agency would perform unannounced security audits and checks—a kind of "privacy penetration-testing".

The agency would be funded by enterprises using cloud computing, but be cheaper than traditional assurance through contracts. It would provide an incentive to keep cloud providers honest in the face of price competition, and is likely to detect problems before they lead to a privacy breach. If enterprises have to pay more for this service for the more sensitive data they place in the cloud, it would provide an incentive for them to minimise the amount of sensitive data they put out there, and thereby limiting the amount of risk they take on behalf of their clients.

The biggest issue is how the agency can actually test whether a cloud provider complies with privacy laws. The rest of this section will discuss some aspects of testing socio-technical aspects of security, but this is only part of the story. The other part is an evaluation of the infrastructure, processes in place, etc. One important

requirement is that the agency must ensure, that these evaluations actually are conducted, and repeated at random intervals to ensure the results' validity. The results of agency evaluations must be available publicly, to allow cloud users and data owners to access, e.g., comments and development of evaluations.

It is important to note that the agency would only be able to test and evaluate the security and privacy measures in place at a cloud provider. Questions such as local jurisdiction being able to force a provider to give access to data might be noted in the agency's report, but per se cannot be part of the seal-decision process, since they are independent of the quality of privacy measures.

Other privacy-relevant questions that are related to the application run by the cloud user on the cloud provider's infrastructure cannot be part of the evaluation either.

12.4.1 Socio-Technical Security Testing

When an organisation decides to work together with a cloud provider, thereby investing a certain amount of trust as described above, the organisation needs to adapt its security and privacy protection measures to accommodate for the new scenario that non-organisation owned premises become part of the organisation's premises, and that non-organisation staff becomes enabled to access the organisation's data. These scenarios did already exist before cloud computing, e.g., with hosted computing and outsourcing, but the promise of cloud computing is that outsourcing becomes an easy-to-use service, and that data can relocate between different machines, countries, continents and (at some point) also providers, without the data owner noticing.

To protect their resources, organisations usually develop security and privacy measures in a top-down manner. The high-level policies describe the desired behaviour of the employees (social domain), the physical security of the premises where the employees work (physical domain), and the IT security of stored and processed information (cyber domain). After the high-level policies have been designed, the individual departments, often with help of a company-wide security department, refine these policies into implementable, low-level policies. These should be enforced through physical and digital security mechanisms as well as employee training. For example, to make sure that data stored on laptops does not end up outside the organisation, policies may be put in place on encryption, physical access to offices, as well as bringing in guests.

Assessing whether the organisations' policies address all identified threats, and whether they are correctly enforced, consists of two steps: *auditing* and *penetration testing*. During the auditing process, auditors systematically check whether proper security policies are defined at all levels and ensure that the policies in place address the identified threats. After the auditing process, penetration tests are used to check whether the policies are enforced without error and whether the policies follow the design specifications.

Both auditing and penetration testing are mature fields in information security and follow methodologies that aim for reliable, repeatable and reportable results.

To address cloud computing they must be extended, e.g., to implement the privacy-penetration testing suggested above. However, the attention paid to the physical and social domain by these methodologies is limited. Unfortunately, adversaries do not limit their actions only to the digital domain, but they use any weak link they can find regardless of the domain. The lack of methodologies for auditing and testing the alignment of security policies across all three domains makes organisation vulnerable to attacks where the adversary combines physical, digital and social actions to achieve a goal. These cross-domain attacks are even more significant in cloud-computing environments than in standard IT infrastructures, since an organisation's perimeter now includes the cloud provider's premises, its IT infrastructure and staff, all providing new attack vectors into the system.

These problems are further aggravated when organisations have to deal with distributed perimeters or the aforementioned de-perimeterisation caused by cloud-computing infrastructures. In these cases policies need to address much more complex scenarios, since the different *domains* now need to be considered in different *perimeters* as well. The same holds for auditing and penetration testing of policies.

A typical example for an attack that cannot easily be found by evaluating policies only at one level is the so-called "road apple attack":

> An attacker leaves a number of dongles with malicious software in front of the premises of an organization. Employees will take dongles, some of them will plug them into their computer, some of which will not be sufficiently protected, and on some of the thus infected machines the malicious payload will find interesting data to encrypt and send with the employee's credentials.

This attack clearly combines elements from different domains (unawareness of employee, inability to check for dongles, inability to check encrypted traffic) that make it hard to detect, but also hard to audit for. To mitigate this attack we need to apply a combination of policies that are coordinated between different stakeholders. Whilst the likelihood of an attack like this on a cloud provider hopefully is rather small (after all, *they are the experts*), a cloud user itself might imagine its data is safe in the cloud, but might still be attackable since the data needs to be transferred to the local machines to work on them.

Once low-level policies have been defined, they need to be enforced using security mechanisms, and this might result in mistakes. Technicians might put the wrong lock on a door, an employee might ignore or forget some of the policies, or some computers might be misconfigured and, for example, might still accept remote connections. Therefore security departments need to be able to test whether the security policies are properly implemented. These tests should include attempts of gaining physical access to restricted areas, as well as attempts in tricking employees to violate a policy (Dimkov et al. 2010b).

Whilst these tests already are hard to apply in a thorough way for traditional scenarios, issues get worse when we consider cloud computing and its additional challenges. We can expect privacy-penetration testing to work well for testing cloud providers' compliance with privacy regulations to a certain extent as discussed above; however, because of likely differences in tools, languages and ontologies used in different organisations, it will in general be impossible to test the alignment of the

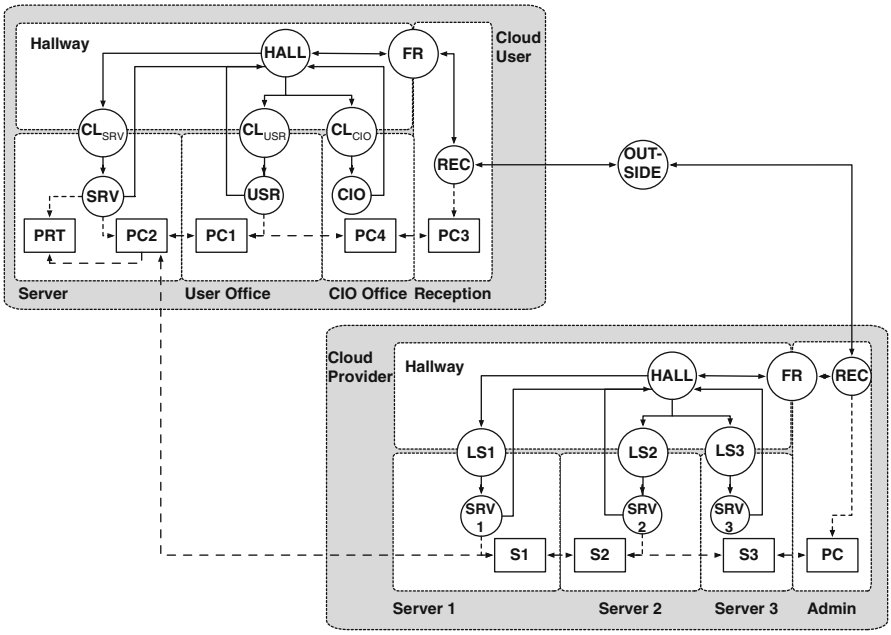

Fig. 12.1 A system model based on ExASyM (Probst and Hansen 2008) illustrating how the physical level (*solid lines*) and the cyber level (*dashed lines*) interact. The model combines a company (*upper left*) with a cloud provider (*lower right*), and represents physical and cyber infrastructure. This model forms the basis of generating attacks based on policies and access control specifications, and can be used for guiding privacy penetration testing

providers' policies with the organisation's policies, and whether the former are in accordance with the latter. On the other hand, providers who are willing to cooperate with organisations to conduct social penetration testing as described above may be able to obtain higher ratings in a quality evaluation.

In the end, the cloud user will need to trust to a certain extent in the cloud provider's will and credibility to enforce certain policies—the goal must be to minimise the gap between the real risk faced by the organisation and the risk it is willing to accept (Probst and Hunker 2010). The suggested public penetration-testing agency is one tool for organisations to evaluate how big a risk they need to take, or how much trust they can have in their cloud provider.

12.4.2 Socio-Technical System Models

To allow for systematic approaches to testing of information infrastructures, including cloud-computing service architectures, we need models for describing the interesting aspects of the system in question. In the penetration test described above, such systems models (Fig. 12.1) can be used to automatically develop attack sce-

narios to be executed in the tests. The benefit of this approach is that it takes into account the actual system of technologies, physical infrastructures such as buildings, and people in a systematic way (Dimkov et al. 2010a,b; Probst et al. 2006).

System models are specific tools within the framework of organisational security-policy alignment; making sure security policies adequately address the goals they were put in place for. In cloud scenarios, this involves alignment of policies *between* organisations as well. Policy alignment aims at ensuring that policies are aligned horizontally, with policies at the same abstraction level, and vertically, with policies at different abstraction levels.

When defining a set of high-level policies, two problems arise: the policies might conflict with each other, or there might be situations for which no policy is defined, resulting in threats not being addressed. Horizontal alignment of policies aims at assuring that high-level policies are consistent and address as high a percentage of threats as possible. When introducing new policies they need to be checked for consistency with existing policies, and for adequacy in protecting against the attacks they were meant to address.

Ideally, high-level policies and low-level policies should allow and forbid the same behaviour. Vertical alignment of policies aims at refining high-level policies to low-level policies whilst ensuring that the latter faithfully implement the former. It is this vertical alignment of policies that system models aim to address, by testing the infrastructure with its low-level policies against the targets expressed by the high-level policies. It is then for example verified if, within the constraints represented by the low-level policies, it is possible for sensitive data to leave the premises. By describing policies in system models, it can thus be verified whether higher-level policies are satisfied. When low-level policies allow behaviour that violates a high-level policy, an attack scenario is produced. Such an attack scenario can then be used as input for the penetration tests.

Although the low-level policies developed in the departments may be complete when restricted to a single domain, when combined with policies from other security domains the combination may not necessarily be complete as well. Thus, a number of actions allowed in one domain may lead to an attack when combined with allowed actions from other domains. In order to support attack scenario generation, models need to be able to describe not only the technical aspects of the system, such as infrastructure, policies, and access-control specifications, but also sociological aspects, such as potential behaviour of employees (Probst and Hansen 2008; Dimkov et al. 2010a; Pieters 2011a). Using this additional information, attacks on the modelled infrastructure can be generated. These represent misalignments between high-level policies and low-level policies.

Using models that include likelihood of certain events to happen, it becomes possible to include descriptions of the less predictable human behaviour into reasoning. The models can then be used to estimate the *risk* of attacks, namely the probability of success with the losses incurred when the attack succeeds, and attacks can be ranked based on the risk. The losses incurred, often called the *impact* of an attack, can be calculated based on the value of the information assets that are affected.

The important benefit of using models and tools for generating attacks is twofold. First, tools can explore also large system specifications in systematic ways, and guarantee a thorough investigation, resulting in a list of possible attacks. Second, and this is especially important when considering cloud computing, one can combine models from different sources to obtain a holistic view of the overall system. This guarantees that the penetration tests performed by the PuPPeT agency cover all possible weaknesses.

Again this is a special problem when considering policies that are defined within an organisation, policies that are defined at the cloud provider, and policies that are defined between the two. It is because of the increased possibilities for misalignment in a multi-organisational context that inter-organisation penetration testing becomes even more important. This means that cloud providers could provide a model of their system, which could then be used by privacy penetration testers to guide the testing process. Since such a system model should be considered sensitive information, we expect that it either is shared only between the cloud provider and the agency performing the privacy penetration testing, or it is developed by the agency as part of a kind of certification process.

12.5 Open Questions

So far we have discussed security and privacy protection-related issues in the cloud and cloud-based applications as faced by cloud providers and cloud users. One of the biggest problems is that of trust into the protection of sensitive data, and an awareness of privacy issues when storing data in the cloud. The suggested privacy penetration-testing agency PuPPeT could help in addressing these points. Of course, there still remain a lot of open questions with respect of protection of privacy of data stored in the cloud.

There seems to exist an inherent contradiction between the rationale of cloud computing—to compute on data where it can be done more efficiently and therefore cheaper—and the requests to ensure privacy of personal and confidential data. Constant tracking to avoid *any* leak or abusive use is technically unfeasible. The big goal remains to combine the seemingly contradicting goals of reducing costs and ensuring security and privacy. As discussed before, the incentive to be able to save money often will be more tempting than the obligation to protect privacy. To overcome this, adjusting the cost for obligatory privacy penetration testing based on the sensitivity of data may be a promising approach.

Data location is an important issue for legal protection. The European Commission seems inclined to keep personal data being processed in cloud computing on European territories, but the question is, how feasible this requirement is—making the location of data redundant is one of the big promises of the cloud, and in most frameworks it is difficult at best to limit data's location (van Cleeff et al. 2010). This is also difficult from an auditing point of view.

However, even if we were able to solve these two problems, there remain other, equally important questions. Once data is stored in the cloud, how do we secure cloud-computing systems against breaches? Because of the stack of technologies used in cloud-computing infrastructures, they also offer new, increased attack surfaces, and as before we need to develop security procedures that can mitigate the threats resulting from these.

The ultimate goal, however, must be to enable end users to have confidence that their data is protected when being stored in the cloud—either by them or by organisations. To this end, there is a significant need for privacy frameworks for cloud applications that ideally should embrace different cloud providers. In the long run this would help to ensure that storing data in the cloud could be considered safe.

A public cloud that offers on-demand services to a wide population of users must take relevant compliance mandates with utmost responsibility to minimise the risk of breaches of data privacy—or risk loss of business due to bad publicity and lack of trust. To achieve this high level of data protection, identity management technologies such as authentication, authorisation, user management, compliance and others are paramount:

- Users must be strongly authenticated to validate their identity;
- Web-based Identity Federation to ease the authentication process should be available;
- Up to date access rights must be checked against cloud application's access control policies;
- All user interactions must be logged to ensure non-repudiation;
- User accounts must be de-provisioned in a timely manner;
- Dormant accounts must be identified and removed quickly; and
- Access permissions must be certified on a continuous basis.

To date, many of these points require explicit actions, which results in untimely responses and consequently vulnerabilities.

Future research clearly should address these points, and try to proactively develop protection and detection mechanisms. We expect to see an increasing number of vulnerabilities in cloud computing that we need to be able to handle. On the one hand systems will be ported to the cloud, which have not been developed for cloud computing, thereby being exposed to threats that were not relevant in the original development. On the other hand, once we know how to address infrastructure vulnerabilities, we expect to observe new threats on the application level, threats that are enabled by the cloud infrastructure.

12.6 Conclusions

Cloud computing is offering new opportunities for public and private organisations to get access to IT infrastructure. A traditional, cloud-based environment offers quick and cost-effective access to technology using a browser. This brings agility

to enterprises and improved satisfaction to end users, whilst lowering overall costs. We have argued above that this promise of cost reduction is very likely to result in increased uncertainty about security as well. This is the case since reduced cost is easy to identify (by reading ones invoice), whilst reduced security and privacy protection first becomes obvious once a breach has happened.

We therefore see the need of strict controls and rules to be applied in cloud computing to meet the requirements for efficient personal data protection. European laws must evolve to regulate this new computing approach, and we propose PuPPeT, a privacy penetration-testing agency, to facilitate this.

An independent agency would be an important step in the right direction. It would signal industry and their clients that public agencies realise the risks and take them serious. It would signal customers, by means of a privacy seal, whether or not they should trust cloud providers. And it would signal end users whether the companies they interact with use trustworthy providers. The combined effect of these signals would be that customers are empowered to decide whether or not they are willing to trust in a cloud provider.

Whilst we strongly believe that the suggested privacy seal issued by a privacy penetration-testing agency is an important step in the right direction, a word of warning seems in order. As Edelman (2011) has noted, among online shops the services accredited with trust certificates are more than twice as likely to be untrustworthy as uncertified sites. This is why we believe that a European agency with comprehensible, publicly documented standards and publicly available testing results is essential in guaranteeing privacy of data stored in the cloud.

Another challenge of cloud computing is the increased amount of third-party infrastructure that organisations need to rely on. By letting go of the infrastructure, managing security risk becomes thus an even more important task than before, requiring a joint effort between the client and cloud provider. Here the modular approach described above might be a viable solution to enable companies to evaluate the risk of including a certain provider's infrastructure.

References

Abrams, Marshall D., and David Bailey. 1995. Abstraction and refinement of layered security policy . In *Information security—An integrated collection of essays*, ed. Abrams, Marshall D., S. Jajodia and H.J. Podell, 126–136. New York: IEEE Computer Society Press.

Balboni, Paolo. 2009. *Trustmarks in e-commerce*. The Hague: Cambridge University Press.

van Cleeff, André, Wolter Pieters, and Roel J. Wieringa. 2010. Benefits of location-based access control: A literature study. *Proceedings of the 3rd IEEE/ACM International Conference on Cyber, Physical and Social Computing* (*CPSCom 2010*). Hangzhou: IEEE Computer Society.

van Cleeff, André, and Roel J. Wieringa. 2009. Rethinking de-perimeterisation: Problem analysis and solutions. *IADIS International Conference Information Systems*, 105–112. Barcelon: IADIS.

Dimkov, Trajce, Wolter Pieters, and Pieter H. Hartel. 2010a. Portunes: representing attack scenarios spanning through the physical, digital and social domain. *Proceedings of the Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security*

(ARSPA-WITS'10). Revised Selected Papers, 112–129. Lecture Notes in Computer Science (6186). Springer Verlag.

Dimkov, Trajce, André van Cleeff, Wolter Pieters, and Pieter H. Hartel. 2010b. Two methodologies for physical penetration testing using social engineering. *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*, 06–10 Dec 2010, Austin, 399–408.

Edelman, Benjamin. 2011. Adverse selection in online "trust" certifications and search results. *Journal Electronic Commerce Research and Applications* 10, (1):17–25.

European Data Protection Directive. 1995. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. *Official Journal* L 281.

Flechais, Ivan, Jens Riegelsberger, and M. Angela Sasse. 2005. Divide and conquer: The role of trust and assurance in the design of secure socio-technical systems. *Proceedings of the 2005 Workshop on New Security Paradigms, NSPW' 05.* New York: ACM.

Floridi, Luciano, and Matteo Turilli. 2011. Cloud computing and its ethical challenges. Paper presented at the Workshop on New Ethical Responsibilities of Internet Service Providers. Hatfield.

Hirsch, Dennis D. 2011. The law and policy of online privacy: Regulation, self-regulation, or co-regulation? *Seattle University Law Review* 34 (2). http://ssrn.com/abstract=1758078. Accessed 1 Sept 2011.

Hunker, Jeffrey, and Christian W. Probst. 2011. Insiders and insider threats, an overview of definitions and mitigation techniques. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications* 2 (1): 3–25.

Jansen, Wayne, and Timothy Grance. 2011. Guidelines on security and privacy in public cloud computing, Draft NIST Special Publication, National Institute of Standards and Technology.

Jericho Forum. 2005. Jericho whitepaper. http://www.opengroup.org/projects/jericho/uploads/40/6809/vision_wp.pdf. Accessed 1 Sept 2011.

Mitra, Sramana, and Saurabh Mallik. 2010. Thought leaders in cloud computing: Interview with Mark White, CTO of Deloitte (Part8). www.sramanamitra.com. Accessed 1 Sept 2011.

Pearson, Siani, and Andrew Charlesworth. 2009. Accountability as a way forward for privacy protection in the cloud. Proceedings of the 1st International Conference on Cloud Computing, CloudCom' 09. Berlin: Springer.

Pieters, Wolter. 2011a. Representing humans in system security models: An actor-network approach. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, 2 (1): 75–92.

Pieters, Wolter. 2011b. Security and privacy in the clouds: A bird's eye view. In *Computers, privacy and data protection: An element of choice*, ed. Serge Gutwirth, Yves Poullet, Paul De Hert and Ronald Leenes, 445–457. Dordrecht: Springer.

Probst, Christian W., Rene Rydhof Hansen, and Flemming Nielson. 2006. Where can an Insider attack. Proceedings of the 4th international conference on Formal aspects in security and trust, FAST' 06. Springer.

Probst, Christian W., and Rene Rydhof Hansen. 2008. An extensible analysable system model. *Information Security Technical Report*, 13 (4): 235–246.

Probst, Christian W., and Jeffrey Hunker. 2010. The risk of risk analysis and its relation to the economics of insider threats. In *Economics of information security and privacy*, ed. Tyler Moore, David Pym and Christos Ioannidis, 279–299. Springer.

Riegelsberger, Jens, M. Angela Sasse, and John D. McCarthy. 2005. The mechanics of trust: A framework for research and design. *International Journal of Human-Computer Studies* (Elsevier) 62 (3): 381–422.

Robinson, Neil, Lorenzo Valeri, Jonathan Cave, Tony Starkey, Hans Graux, Sadie Creese, and Paul Hopkins. 2011. The cloud: Understanding the privacy and trust challenges, RAND Europe, Technical Report, 2011.