The Cloud: Understanding the Security, Privacy and Trust Challenges

Final Report

Neil Robinson, Lorenzo Valeri, Jonathan Cave & Tony Starkey (RAND Europe)
Hans Graux (time.lex)
Sadie Creese & Paul Hopkins (University of Warwick)

TR-933-EC

30 November 2010
Prepared for Unit F.5, Directorate-General Information Society and Media, European Commission

Preface

The overall objective of *The Cloud: Understanding the Security, Privacy and Trust Challenges* is to advise on policy and other interventions which should be considered in order to ensure that European users of cloud environments are offered appropriate protections, and to underpin a world-leading European cloud ecosystem. Cloud computing is increasingly subject to interest from policymakers and regulatory authorities. The European Commission's recent Digital Agenda[1] highlighted a need to develop a pan-European 'cloud strategy' that will serve to support growth and jobs and build an innovation advantage for Europe. However, the concern is that currently a number of challenges and risks in respect of security, privacy and trust exist that may undermine the attainment of these broader policy objectives.

Our approach has been to undertake an analysis of the technological, operational and legal intricacies of cloud computing, taking into consideration the European dimension and the interests and objectives of all stakeholders (citizens, individual users, companies, cloud service providers, regulatory bodies and relevant public authorities).

This study represents an evolutionary progression in understanding the implications of cloud computing for security, privacy and trust. Starting from an overview of the challenges identified in the literature, the study builds upon real-life case study implementations of cloud computing for its analysis and subsequent policy considerations. As such, we intend to offer additional value for policymakers beyond a comprehensive understanding of the current theoretical or empirically derived evidence base.

As a first informal deliverable of this project, a draft literature review paper was prepared which formulated the study team's common understanding of cloud computing and the associated open questions surrounding some of the important security, privacy and trust issues. This was based on the main academic, research and industry publications available.

In order to test the accuracy and completeness of this common understanding and to bring it to a more mature form, the study team participated in the SecureCloud conference in Barcelona on March 16–17 2010.[2] During this conference, a series of interviews with key stakeholders from various domains were conducted. While these largely corroborated the validity of the draft literature review, a few additional comments and nuances were provided which had not been sufficiently recognised in the draft.

---

[1] European Commission (2010a)

[2] See (as of 20 November 2010) http://www.cloudsecurityalliance.org/sc2010.html

On the basis of these inputs, a consolidated literature review was prepared and then enriched with descriptions of specific cloud case studies. This was presented to a group of stakeholders at an expert workshop in Brussels on 10 September 2010, during which possible recommendations with respect to EU policy actions on cloud computing were identified and discussed.

The present paper represents the final consolidation of all inputs, suggestions and analyses and contains our recommendations for policy and other interventions.

Due to the length and complexity of this paper, it includes a brief 2 page Executive Summary representing an abridged version of the key points, followed by a longer extended summary 'Synthesis' Chapter.

For more information please feel free to contact:

Neil Robinson

Senior Analyst

RAND Europe

Westbrook Centre

Milton Road

Cambridge CB4 1YG

United Kingdom

Tel: +44 (0)1223 353 329

Email: neil_robinson@rand.org

Acknowledgements

The authors would like to thank Marjory Blumenthal (RAND Corporation / Georgetown University), Simone Vernacchia (Rome), Marco Casassa Mont (HP Labs, Bristol) and Project Officers Bart van Caenegem, Valeria Bricola and Mario Scillia from the European Commission, Directorate-General Information Society and Media Unit F.5 (Trust and Security) for their helpful and constructive remarks during the preparation of this report. In addition the authors would like to acknowledge the support of Dr Lorenzo Valeri, who, following his new appointment at the School of Government in LUISS University in Rome in April 2010, maintained his engagement with the study and moderated the expert workshop. Finally, the authors would also like to acknowledge the contributions of other researchers most notably Helen Rebecca Schindler (RAND Europe) and Davide Maria Parrilli (time.lex).

Contents

Glossary 7

Executive Summary 10

Synthesis 13
Defining cloud computing 13
Defining security, privacy and trust 14
Issues arising from the reviewed literature 14
Risk control frameworks 16
Operational challenges 17
Implications from case studies 18
Gap analysis 20
Solving the challenges: observations and recommendations 21
Conclusions 24
Methodology 24
Structure of the report 25

CHAPTER 1 **Introduction** 27

CHAPTER 2 **Definitions and drivers** 28
2.1 Definitions of cloud computing 28
2.2 What's pushing cloud take-up? 30
2.3 The economics of cloud computing: implications for security 31
2.4 Concluding remarks 34

CHAPTER 3 **Understanding the implications for security, privacy and trust** 35
3.1 Defining security, privacy and trust 35
3.2 Growing focus on security, privacy and trust concerns 35
3.3 Identifying key issues and possible enablers for security, trust and privacy in the cloud 39

CHAPTER 4 **Security, privacy and trust challenges stemming from the technological underpinnings of cloud computing** 40
4.1 The linchpin of trust: the hypervisor 40

4.2 Can the distributed models of computation characteristic of grid technology adequately serve the availability and interoperability needs of cloud computing? ... 43
4.3 Current state-of-the-art web services may not be sufficient to establish interoperability for identity management in the cloud ... 45
4.4 Trustworthiness in service-orientated architectures (SOAs) ... 47
4.5 Will web application frameworks (APIs and SDKs) be credible in providing trust across distributed environments? ... 47
4.6 The fragility of current encryption approaches in the cloud context ... 48
4.7 Concluding remarks ... 51

CHAPTER 5 **Security, privacy and trust challenges inherent to the legal and regulatory aspects of cloud computing ... 53**
5.1 Horizontal perspective: applicable law and jurisdiction ... 54
5.2 Vertical issues: main applicable laws ... 55
5.3 The will of the parties: contractual provisions ... 61
5.4 Overcoming legal barriers: key tools ... 62
5.5 Concluding remarks ... 63

CHAPTER 6 **Putting it all together: key risks and operational challenges ... 66**
6.1 Summary of legal and technical issues ... 66
6.2 Migrating to the cloud: the operational challenges ... 68

CHAPTER 7 **Case studies ... 73**
7.1 Introduction ... 73
7.2 Initial classification of case studies ... 73
7.3 Observations on case studies in practice ... 75
7.4 Matrix of case study typology ... 75
7.5 Case Study 1: the Danish National IT and Telecom Agency ... 76
7.6 Case study 2: the City of Los Angeles ... 82
7.7 Case study 3: EU eHealth provider ... 90
7.8 Implications ... 94

CHAPTER 8 **Gap analysis ... 98**

CHAPTER 9 **Solving the challenges: recommendations and actions ... 106**
9.1 Introduction ... 106
9.2 Recommendations ... 108
9.3 Specific actions ... 111

CHAPTER 10 **Conclusions ... 117**

REFERENCES ... 119

APPENDICES ... 129

Glossary

Accountability – the ability of an entity to be held responsible or answerable (including in legal terms) for its actions

ACM – Association for Computing Machinery

AJAX – Asynchronous JavaScript and XML

APEC – Asia-Pacific Economic Cooperation

API – Application Programming Interface

Capex – Capital Expenditure: upfront or one off investment

CDN – Content Delivery Network

CEN – European Committee for Standardization

Cloud Service Provider – the entity offering the cloud service

Cloud Service User – the entity using the cloud service. This may be an organisation (such as a firm) or a private individual and so may be distinct from the end user

CORBA – Combined Object Request Broker Architecture

CPU – Central Processing Unit

Data Controller – an entity (defined under Article 2 of the EU Data Protection Directive 95/46/EC as "natural or legal person, public authority, agency or any other body which alone or jointly with others determines the purposes and means of the processing of personal data"

Data Processor – an entity defined as acting on behalf of the Data Controller for the processing of personal data

Data Subject – commonly understood as the living individual who is the subject of personal data

DCOM – Distributed Component Object Model

DMCA – Digital Millennium Copyright Act

DMTF – Distributed Management Task Force

EC2 – Amazon Elastic Compute Cloud

End-user – the ultimate user or customer of a product or service

ENISA – European Network and Information Security Agency

FedRAMP – Federal Risk and Authorization Management Program

FP7 – Framework Seventh Programme

HaaS – Hardware as a Service

IaaS – Infrastructure as a Service

IDM – Identity Management

IEEE – Institute of Electrical and Electronics Engineers

IPR – Intellectual Property Rights

ITU – International Telecommunication Union

NIST – National Institute of Standards and Technology

OECD – Organisation for Economic Co-operation and Development

Opex – Operating Expenditure – recurring or regular costs

OVF – Open Virtualisation Format

PaaS – Platform as a Service

Privacy – a number of definitions of privacy exist but they commonly revolve around such concepts as non-intrusion upon the self; the right to be let alone or the expression of or adherence to various legal and non legal norms regarding the right to private life

QoS – Quality of Service

S3 – Amazon Simple Storage Service

SaaS – Software as a Service

SAML – Security Assertion Markup Language

Security – in the context of this study, this term refers to information security, to which we mean the confidentiality, availability and integrity of information

SEIM – Security Event and Incident Monitoring

Shared Services – internal sharing of IT services between administratively separate sub-entities within the same organisation

SLA – Service Level Agreement

SMC – Secure Multiparty Computation

SOA – Service-Orientated Architecture

SOAP – Simple Object Access Protocol

SSL – Secure Sockets Layer

Transparency – a generalised concept of appropriate openness

Trust – assurance or confidence that people, data, entities, information or processes will; function or behave in expected ways

TPM – Trusted Platform Module

UDDI – Universal Description Discovery and Integration

Virtualisation – the ability to run variable virtual (separate from actual) versions of operating systems, hardware, servers etc

VM – Virtual Machine

WS-Agreement – Web Services Agreement

WSDL – Web Services Description Language

Executive Summary

Our research investigated the security, privacy and trust aspects of cloud computing and determined whether these were sufficiently distinct to warrant public policy intervention. On the whole, cloud computing brings into acute focus many security and privacy challenges already evident in other domains such as outsourcing or behavioural advertising.

Defining cloud computing

For the purposes of this study, we adopt a definition of cloud computing proposed by the US National Institute for Standards and Technology (NIST) in 2009:

> *Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction*.

The literature identifies four different broad service models for cloud computing:

- *Software as a Service (SaaS)*, where applications are hosted and delivered online via a web browser offering traditional desktop functionality
- *Platform as a Service (PaaS)*, where the cloud provides the software platform for systems (as opposed to just software)
- *Infrastructure as a Service (IaaS)*, where a set of virtualised computing resources, such as storage and computing capacity, are hosted in the cloud; customers deploy and run their own software stacks to obtain services.
- *Hardware as a Service (HaaS)*, where the cloud provides access to dedicated firmware via the Internet

Cloud computing offerings also differ by scope. In private clouds, services are provided exclusively to trusted users via a single-tenant operating environment. Essentially, an organisation's data centre delivers cloud computing services to clients who may or may not be in the premises. Public clouds are the opposite: services are offered to individuals and organisations who want to retain elasticity and accountability without absorbing the full costs of in-house infrastructures. Public cloud users are by default treated as untrustworthy. There are also hybrid clouds combining both private and public cloud service offerings.

Issues, challenges and concerns

We identified a number of issues in the **literature** relating to technological and legal challenges confronting privacy, security and trust posed by cloud computing. Regarding the challenges in the technological underpinnings of cloud computing, we note evidence stemming from: virtualisation (e.g. vulnerabilities in hypervisors could have potential widespread effects on data integrity and confidentiality); whether grid computing models can afford the appropriate level of interoperability ; if Web Services will be effective for identity management in the cloud; establishing trust when using Service Orientated Architectures and web application frameworks and finally whether current technical methods of encryption will remain viable to achieve confidentiality. There are a number of challenges posed by a range of legal and regulatory frameworks relevant to cloud computing. These include the viability of legal regimes which impose obligations based on the location of data; the ex-ante definition of different entities (such as distinguishing between data controllers and processors); establishing consent of the data subject; the effectiveness of breach notification rules; the effectiveness of cyber-crime legislation in deterring and sanctioning cyber-crime in the cloud and finally difficulties in determining applicable law and jurisdiction. From an operational perspective, the study uncovered issues relating to the effectiveness of existing risk governance frameworks, whether cloud customers can meet their legal obligations when data or applications are hosted overseas, how to be compliant and accountable when incidents occur; whether data will be locked into specific providers; the complexities in performing audit and investigations; how to establish the appropriate level of transparency and finally measuring security of cloud service provision.

The **case studies** identified a number of challenges relating to cloud service provision from recent real-world instances. These include the immature and exploratory nature of cloud computing deployments; the necessity that those using cloud services should be versed in their tolerance for risk prior to migrating to the cloud; how to balance the business benefits of cloud computing with achieving security and privacy obligations; the need to integrate cloud security into existing security measures; the importance of understanding untoward dependencies created by cloud computing deployments and finally that tailored and specific security agreements can be achieved but only if the cloud user has sufficient negotiating power.

These identified real-world concerns were supplemented by additional material gathered at an **Expert Workshop**. Participants commented that it was difficult to achieve a high degree of accountability or transparency in the cloud; that there was little awareness raising for either cloud customers or private citizens; little established guidance on expectations for cloud users in meeting their legal obligations and finally lack of harmonisation of relevant legal and regulatory frameworks, potentially presenting an impediment to realising the economic and social benefits of cloud computing for Europe.

Recommendations

Our recommendations are orientated around four themes of current opportunities for policy action:

- **Compliance** - *Greater harmonisation of relevant legal and regulatory frameworks* to be better suited to help provide for a high level of privacy, security and trust in cloud computing environments. For example: *establishing more effective rules for accountability and transparency* contributing to a high level of privacy and security in data protection rules and *expansion of breach notification regimes* to cover cloud computing providers.
- **Accountability** – Improvement of rules enabling cloud users (especially consumers) to *exercise their rights* as well as *improvement of models of Service Level Agreements (SLAs)* as the principle vehicle to provide accountability in meeting security, privacy and trust obligations.
- **Transparency** –improving to way in which levels of security, privacy or trust afforded to cloud customers and end-users can be discerned, measured and managed, including *research into security best practices, automated means for citizens to exercise rights* and *establishment of incident response guidelines*.
- **Governance** – The *European Commission could act as leading customer* by deploying cloud computing solutions as part of its e-Commission initiative and indirectly supporting the improvement of existing operational risk control frameworks. Research funding could be assigned to *improving Security Event and Incident Monitoring* in the cloud amongst other things.

Synthesis

The central hypothesis investigated by this study is whether the complexity associated with cloud computing and the type and character of challenges to security, privacy and trust are sufficiently unique to warrant the development of entirely new public policy approaches.

Defining cloud computing

For the purposes of this study, a definition of cloud computing proposed by the US National Institute for Standards and Technology (NIST) in 2009 has been adopted:

> *Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.*[3]

Building upon this definition, the literature identifies four different broad service models for cloud computing:

- *Software as a Service (SaaS)*, where applications are hosted and delivered online via a web browser offering traditional desktop functionality, eg, Google Docs, Gmail and MySAP.
- *Platform as a Service (PaaS)*, where the cloud provides the software platform for systems (as opposed to just software), the best current example being the Google App Engine.
- *Infrastructure as a Service (IaaS)*, where a set of virtualised computing resources, such as storage and computing capacity, are hosted in the cloud; customers deploy and run their own software stacks to obtain services. Current examples are Amazon Elastic Compute Cloud (EC2), Simple Storage Service (S3) and SimpleDB.
- *Hardware as a Service (HaaS)*, where the cloud provides access to dedicated firmware via the Internet, eg, XEN and VMWare.

The literature also differentiates cloud computing offerings by scope. In *private clouds*, services are provided exclusively to trusted users via a single-tenant operating environment. Essentially, an organisation's data centre delivers cloud computing services to clients who may or may not be in the premises. *Public clouds* are the opposite: services are offered to

[3] Mell & Grance (2009a)

individuals and organisations who want to retain elasticity and accountability without absorbing the full costs of in-house infrastructures. Public cloud users are by default treated as untrustworthy. There are also *hybrid clouds* combining both private and public cloud service offerings.

Defining security, privacy and trust

Before assessing the literature dealing with security, privacy and trust in the cloud, it is important to define these terms because their currency and usage can change radically in different contexts:

- Security concerns the confidentiality, availability and integrity of data or information. Security may also include authentication and non-repudiation.
- Privacy concerns the expression of or adherence to various legal and non-legal norms regarding the right to private life. In the European context this is often understood as compliance with European data protection regulations. Although it would be highly complex to map cloud issues onto the full panoply of privacy and personal data protection regulatory architectures, the globally accepted privacy principles give a useful frame: consent, purpose restriction, legitimacy, transparency, data security and data subject participation.
- Trust revolves around 'assurance' and confidence that people, data, entities, information or processes will function or behave in expected ways. Trust may be human to human, machine to machine (eg, handshake protocols negotiated within certain protocols), human to machine (eg, when a consumer reviews a digital signature advisory notice on a website) or machine to human (eg, when a system relies on user input and instructions without extensive verification). At a deeper level, trust might be regarded as a consequence of progress towards security or privacy objectives.

Issues arising from the reviewed literature

Summarised below are the main considerations in respect of the security, privacy and trust challenges associated with the technological or legal domains constituting cloud computing. This is populated from the review of the literature made during the course of this study and reflects key concerns noted from the desk research:

Summary of main challenges arising from the technological or legal domains constituting cloud computing

Area / Technological or legal domain	Security	Privacy	Trust
Virtualisation	Integrity	Segregation of personal data on shared infrastructure	Compromised virtual machines/ hypervisors permit loss of trust
Grid technology	Availability		Interoperability

Web services	Integrity and confidentiality	Security and confidentiality	Interoperability (in the context of identity and access management)
Service-orientated Architectures	Integrity		The reliance of distributed systems on different security credentials
Web application frameworks	Integrity and availability		Trust across distributed environments
Encryption in the cloud context	Confidentiality	Security and confidentiality	
Applicable law (data/service location)		Existence and effectiveness of privacy protection laws/principles	Existence of a clear legal framework as a basis for the service
Dispute resolution		Accountability: can disputes in the cloud be resolved?	Accountability: is the cloud stable enough to inspire trust
Data protection and privacy	Obligation to implement secure data processing approaches	Compliance with privacy principles	Confidence in data protection practices
Protection of intellectual property rights	Confidentiality and availability (data portability)		Confidence in the security/confidentiality of data entrusted to the cloud
Electronic communication in the cloud		Safeguarding communications secrecy	Protection against eavesdropping (surveillance by public and private parties)
Security obligations and cybercrime	Confidentiality, availability and integrity; effective law enforcement	Safeguards against unlawful intrusions in the personal sphere	Balancing privacy safeguards with the need for security
Accountability and liability	Accountability for security breaches and incidents	Accountability for data leaks: can incidents be identified and sanctioned?	Trust that instruments for restitution and sanction will work
Harmful and illegal content	Availability: can the cloud identify and respond to such content?		Trust in jurisdictions to apply transparent standard or approach to illegal content (in line with cloud user expectations and applicable laws)
Consumer protection	Obligation to implement secure data processing approaches	Data Subject participation; restitution	Consumer protection rules must be effective (ie, applied and enforced in practice)
Scope and quality of services (SLAs)	Transparency and security metrics are needed to ensure that security targets are met		Assurance and commitments between two parties
Flexibility	Availability and scalability		Resources must be available when needed, in accordance with agreement between parties
Validity and consent	Transparency and accountability	Transparency must be ensured. Consent from	Assurance and commitments between

		consumers must be free, specific and informed	two parties must be clear and enforceable

This table is intended as a short checklist of areas to be covered in cloud deployments, and is used to support the analysis of case studies, and specifically to determine to what extent each case study has been able to find appropriate answers to the legal and technical challenges in their respective domains. It should be noted, however, that this list is mostly valid with respect to *public* cloud systems (where specific data or services are outsourced to a third-party service provider), but much less so in the case of *private* cloud systems (which are deployed, operated and controlled by the user). In the latter case, the use of cloud technologies may still cause technical and security challenges, but to a lesser extent since they can be managed internally. Internal cloud computing deployments represent limited issues in respect of new legal or operational challenges.

Risk control frameworks

The literature also notes a number of challenges for current best practice in managing and controlling how security, privacy and trust risks. These mainly revolve around the extent to which current management frameworks and best practice guidelines such as those in the NIST 800- series or in the ISO2700x suite of Information Security Management System best practices are suitable and effective in managing the operational risks evident in cloud computing environments. Areas where operational governance tools and best practices may be untenable in a cloud context include:

- Physical access controls – how can the cloud user achieve requirements for physical access control given the cloud service provider establishes and controls the when who, why and how of physical access measures?
- Application development and maintenance – is it possible to assure the development and maintenance of applications in a cloud environment when external parties cloud service provider or other third parties are responsible?
- Vulnerability management – assigning responsibility for patch management and the deployment of software and hardware updates between the cloud service provider and cloud user is especially complex given virtualisation and the dynamic reconfiguration of software and infrastructures.
- Monitoring – how to establish effective, timely and accurate monitoring of levels of security and privacy in business-critical infrastructure when those responsible for the infrastructure may not be prepared to share such information under standard service level agreement.
- Identification and authentication – the integration and control of identity and access management infrastructures in a cloud environment where the cloud service provider might have different approaches and tolerance for risks to identity infrastructure, in addition to the complexities of providing for identity across distributed cloud environments.
- Access control – how can the cloud user govern access control risks when the levels and types of access control to key ICT assets deployed by the cloud service provider are unknown?

- Encryption - how can the cloud user manage encryption and key infrastructures and assign responsibility across the boundary between their own organisation and the cloud service provider.
- Continuity and incident management – how can the cloud user determine appropriate thresholds and criteria for responding to incidents (eg, agreeing on what constitutes an incident) and policies and processes for responding and achieving assurance of the evidential chain.
- People (security team) – can the cloud service user understand and manage the quality and mechanisms for professionalism in the cloud service provider's security team and whether they match with that of the cloud user.
- People (business function) – how incorporating a high culture of security in the organisation is possible when those responsible for strategic decisions regarding the trade off between security/business requirements may be external to the cloud user organisation (and thus may operate under different motivations).
- Security control testing – how can the cloud service user test the validity and effectiveness of security controls when such tests may not be agreed by the cloud service provider or may impact upon other customers of the cloud service provider.
- Accredited components – how to accredit components in the cloud to the required standard imposed by accreditation processes.
- Data remanence – can the cloud service user achieve the appropriate degree of assurance that data have been securely deleted by the cloud service provider and no data shards remain in the cloud environment.
- Asset management – can the cloud service provider gather and exploit data on the entirety of the cloud service provider's physical assets to complete an asset register in accordance with good practice guidance?

The overall theme of these challenges concerning such frameworks for managing privacy, security and trust is whether as currently scoped such guidance is viable to achieve a high degree of assurance in the context of cloud computing. Given the abstracted and dis-intermediated nature of cloud computing deployments, it may be more difficult to achieve such a high degree of assurance that the risks that such controls present in these common risk management frameworks are being adequately addressed.

Operational challenges

Depending on the cloud service the evidence from our study noted that the following operational challenges are of relevance:

- Data or services may be (or become) hosted from another country, even without the end user necessarily being aware.
- Infrastructure may be shared with other customers, leading to data segregation concerns.
- Incidents may cause service interruptions without it being evident where the problem lies, and thus how it may be addressed.

- Data withdrawal might be difficult, in the sense that it can be hard to determine for a cloud user whether deleted data was actually removed from the provider's systems, or whether it was merely made inaccessible to the user.
- Auditing and investigations may be more challenging, due to the complexity of the system (e.g. use of virtualisation technologies may make it harder to determine where data is located, and which systems may be audited without accessing another cloud user's data).

None of these challenges is strictly unique to cloud services: traditional outsourcing models can also be confronted with these same issues. However, the fact that cloud models combine all of these elements and a layer of technical complexity means that they are viewed as more problematic.

Implications from case studies

In order to test the validity of the challenges identified by the literature and interviews conducted at the SecureCloud 2010 conference, three case studies were undertaken. The characteristics of these are listed below:

Characteristics of case studies

Case study	Cloud service model used by case study subject	Service model offered to end users	Risk	Rationale
Danish National IT Agency	IaaS	Traditional (non-cloud)	Low to Medium	This case study represents a social network and knowledge sharing platform. However, there are concerns regarding the use of personal data collected as part of this implementation
City of Los Angeles	SaaS	SaaS	Medium	Although this case study exists in a different regulatory environment to that of Europe, there are still overarching security and sensitivity lessons associated with the storage of criminal record data in the cloud. The LAPD was a major user community in this implementation
EU Member State eHealth provider	IaaS	SaaS	Medium to High	This case study covers the use and processing of health related personal data in the cloud

The selected case studies reflect a paradox regarding the characteristics of cloud computing relevant to security and privacy in that they constitute both a challenge and an opportunity. For example, the degree of granularity regarding the locations of stored data in the cloud may be linked to the negotiating power of the cloud customer and may in certain cases conflict with the prevailing regulatory framework (eg, for European cloud customers, the legal requirements stemming from privacy and data protection law). Equally, the wealth of data created by the logging functionality inherent in the cloud computing model present something of a security opportunity in terms of understanding exactly what is going on in an information and communication technology (ICT) infrastructure, especially for small and medium enterprises (SMEs) or organisations that might otherwise not have the know-how or resources to implement such functionalities on their own. Furthermore, the possibilities of quickly and seamlessly re-provisioning infrastructure also represent an opportunity in respect of availability and resilience. The findings from the case studies are summarised below:

Cloud computing is still experimental

Deployments of cloud computing (in the current understanding of the term) are relatively recent phenomena. Many organisations implementing cloud computing solutions have barely got to grips with the realisation of the intended benefits of the system, let alone the complex issues of security and privacy. The case studies illustrate that organisations are keen to experiment with cloud-based models while relegating compliance and security issues to a secondary role.

Moving to the cloud makes it even more imperative to be fully aware of your own risk tolerance and security level

The case studies illustrate that prior to the implementation of any cloud computing solution, cloud service users must be diligent and effective in identifying their own risks and responsibilities and take appropriate measures where standard terms on offer from the cloud service provider are not adequate.

Can cost drivers override security?

The attractiveness of cloud computing as a way to quickly access capex benefits with respect to ICT service provision would, from analysis of the case studies in this report, appear to outweigh concerns regarding security and specific security requirements of some cloud users (for example, those using or interacting with particularly sensitive types of personal data).

Integration with existing security measures is necessary

The case studies demonstrate the need to integrate cloud computing with existing security mechanisms: some security and trust mechanisms will need to be in place in the cloud customer organisation prior to any deployment. It is not always the case that cloud computing permits a transition from a 'no security' situation to a 'more secure' situation.

Untoward dependencies can be a problem

Cloud computing deployments highlight issues regarding resilience and the need for the cloud customer organisation to consider whether their organisations own access to the Internet is resilient enough (for example, provision of dual-homed Internet connections over physically and logically separate service providers).

Granularity *is* possible (but only if you have the right cloud provider and sufficient negotiating power)

The utility-like characteristics of cloud computing give rise to a weakness in that the standardised products being offered may not always match with the end users requirements, particularly in respect of security and privacy. The case studies illustrate that if the cloud provider sees some incentive in being willing to negotiate over security and privacy measures and the cloud customer has sufficient buying power, then standard terms and conditions may be tweaked to better reflect security and privacy requirements.

Gap analysis

Noting these technological and legal issues along with the operational concerns highlighted by the case studies and further explored in the Expert Workshop, a policy gap analysis was undertaken to see where existing policy instruments or tools appeared to be viable and effective in addressing such challenges or where there were opportunities to improve public policy to obtain better objectives for security, maintain or improve a high level of protection of personal data of citizens and increase trust. The results of this analysis are detailed below:

Gap Analysis for European policy

Domain	Gap or issue
Legislation	
Data Protection Directive 95/46/EC	Viability of regimes detailing specific entities, location determining obligations and rules supporting accountability
e-Commerce Directive 2000/31/EC	Lack of harmonisation of notice and takedown regimes, transparency and linking of applicable law to physical location
BRUSSELS I Council Regulation (EC) 2001/44/EC	Whether the applicable standard EU/EEA rules will work when the cloud user has limited negotiation power
Unfair Commercial Practices Directive 2005/29/EC	Additional and perhaps diverging requirements imposed at the Member State level may undermine economic effectiveness of cloud computing deployments
Electronic Communications Privacy Directive 2002/58/EC, as amended by the Citizen's Rights Directive 2009/136/EC	Breach notification rules do not apply to cloud service providers and communications secrecy obligations may be difficult or impossible to meet in a cloud environment
Data Retention Directive 2006/24/EC	Applicability of data retention rules
Framework Decision on Attacks against Information Systems 2005/222/JHA	Whether existing legal approaches to sanctioning cybercrime will be effective in deterring acts of computer and network misuse in the cloud or addressing their consequences
Implementation and governance	
Cloud users (public and private sector)	Whether Security risk management programmes will continue to be effective in a cloud computing context; awareness raising programmes; low transparency; low accountability; poor management of continuity and resilience
Cloud service providers	Absence of awareness raising programmes; poor transparency; poor accountability
Risk controls	Viability of security controls present in common risk frameworks such as ISO27001 and the NIST standard when shared responsibility for their deployment and management rests

	across two or more parties in a cloud service provision
Guidelines	Little or no guidance on data protection aspects, auditing, industry response or forensics or the specific context of Europe for cloud service users
Research	
ICT research programme	Scarce or non-existent scientific and technical research in supporting investigation into processing of encrypted data, security measurement, security event and incident monitoring in the cloud, identity and access management in cloud access management, interoperability of security policy requirements (across cloud infrastructures) and interoperability between cloud service providers
Policy research	Little policy research into best practice for metrics and guidance for meeting privacy obligations

Solving the challenges: observations and recommendations

Addressing the security and privacy challenges of cloud computing is a complex undertaking since it requires a combination of technological solutions and legal approaches that is capable of addressing operational realities and concerns. The extent of the technical, legal and operational challenges has been presented, and it is clear that an appropriate policy framework is required to address these challenges in a way that reconciles business drivers with public interests and European values, while avoiding the pitfall of imposing burdens that put European businesses at a competitive disadvantage. Without attempting to create a comprehensive single framework for dealing solely with this one domain, we propose recommendations aimed at resolving or progressing the challenges identified previously. The recommendations are based around four themes: compliance, accountability, transparency and governance.

- Compliance: ensuring that a cloud deployment meets the requirements imposed by the applicable normative framework, including general legislation, sector-specific rules and contractual obligations; the challenges in complying with data protection rules are a key example of this.
- Accountability: ensuring that security or privacy breaches in the cloud deployment are correctly addressed, including through appropriate compensation mechanisms towards any victims.
- Transparency: ensuring that the operation of the cloud deployment is sufficiently clear to all stakeholders, including service providers and users, both professional businesses and private consumers; this can be witnessed, for example, in the difficulty of determining who/where a cloud service provider is, and where his responsibilities/liabilities end.
- Governance: ensuring that the European Commission's policy objectives and actions of the European Commission are well aligned with ongoing stakeholder activities, including by actively participating in the establishment and promotion of standards and best practices, and in interactions with cloud service providers.

In the table below proposals are presented for each stakeholder to take forward certain actions to support the achievement of the goals summarised above.

Recommended actions for each stakeholder

Stakeholder	Action
The European Commission, the Article 29 Data Protection Working Party and Member States	Work towards a greater emphasis on international harmonisation of relevant legal and normative frameworks via: • Further efforts to improve consistent application of relevant legal frameworks across the Member States through the conduct of implementation, monitoring and evaluations of relevant EU legislation • Continue to support broad international dialogue fostering harmonisation of relevant legal frameworks (especially regarding privacy) between, eg, the OECD, APEC and the US
The European Commission	Consider making use of cloud services itself, relying on a mixture of private and public cloud systems, and sharing its experiences and best practices with respect to security and privacy with other stakeholders. This could be done by identifying and selecting an appropriate pilot e-Commission service to be delivered via cloud services (eg, launching an European statistical data portal such as www.data.gov in the US)
Regulators in the Member States and cloud service providers	Develop suitable awareness raising mechanisms to help users to become aware of their own privacy and security risks by, for example: • Draft, prepare and issue guidance for cloud users (both organizations and individual consumers) on the benefits, risks and consequences of the storage and use of personal data in the cloud • Draft, prepare and issue guidance for cloud providers on how they should inform cloud users (especially consumers) of their rights in an accessible and understandable manner
The European Commission and regulators in the Member States	Consider introducing/supporting means for collective action against privacy/security breaches in future consumer and data protection policy. This could be achieved via: • The incorporation of such rules in the proposed new European legal framework on privacy and data protection • The dissemination of best practice and common approaches amongst European privacy and data protection regulators in such fora as the London Initiative
The Article 29 Data Protection Working Party	Prepare guidelines (eg, in the form of a policy paper) on expectations of cloud users and obligations imposed on providers in the context of data protection • Prepare and release an Opinion on aspects of meeting data protection obligations in a cloud context relevant for data controllers and data processors
ENISA	Prepare incident response guidelines and response disclosure obligations, if appropriate and necessary • In conjunction with stakeholders such as cloud service providers and practitioners, prepare and release guidance covering how to respond to incidents targeted towards the three main categories of stakeholders (cloud providers, cloud users, and law enforcement). The cross-border nature of cloud computing and effective identification of competent authorities should take priority in this guidance document • Consider the preparation of guidance helping law

	enforcement and cloud providers and users understand the complexities of forensics in a cloud computing environment
The European Commission	Explore if and how cloud users should be made aware of incidents in the cloud that could negatively impact them through consideration of the creation and inclusion of appropriate generalised and breach notification rules in the future legal European framework regarding privacy and data protection
The European Committee for Standardization (CEN) and ENISA	Actively participate in the establishment and promotion of security standards and best practices, and should interact with cloud service providers to foster a culture of security and privacy awareness • Consider the extent to which current globally accepted standards in the realm of security and privacy (such as the ISO2700x Information Security Management System or NIST 800-series standards) might be updated to be more suitable for a cloud environment through stakeholder engagement (eg, via a conference or workshop) • Explore the necessity of new cloud security standards, keeping in mind their potential future use for the purpose of obtaining security certification (eg, as a guideline to obtain ISO/IEC 2700x compliance accreditation) • Explore and support the development of internationally acceptable standards for interoperability (eg, through dialogue with ISO and industry groups such as the Cloud Security Alliance and OpenCloud)
Industry: cloud service providers	Work to make it easier for cloud users and citizens to exercise their rights when using cloud solutions • Be encouraged to participate fully in interoperability and standardisation processes and initiatives to permit users to more easily move their data between providers • Investigate, develop and implement simple, transparent, automated and effective solutions to support the exercise of privacy and data protection rights in a cloud service, permitting data to be easily accessed, erased, moved and rectified, for example
Academia, the European Commission and ENISA	Investigate how policy and practice diverge regarding demand-side factors concerning security, privacy and trust • Conduct policy research into best practice and common approaches to applying security metrics in the cloud • Investigate the creation of a generic maturity model to independently ex-ante measure and assess cloud security provision
Academia, the European Commission and ENISA	Undertake policy and legislative research into standardised SLAs/liability provisions for cloud computing Undertake policy research to identify common and useful liability provisions and templates for SMEs that serve to uphold a high level of security and privacy for cloud users and European citizens
Academia	Attract European research & development funding (through the scientific European 7th Framework Programme) to: • Conduct research to establish effective locations and types of security detection mechanisms (sensors) in cloud architectures to support rapid detection of emergent threats from the cloud • Investigate and pilot secure virtualised architectures and trust domains for cloud computing environments (virtual machines running on specially configured hypervisors

	according to a set of common security domain of controls) • Investigate tools and instruments to support shared responsibility of security controls contained in best practice and guidance such as the ISO2700x suite (eg, configuration of roles and responsibilities for execution of patch management activities between the cloud provider and cloud user) • Design testbeds and pilots for middleware, software interfaces and APIs for Security Event and Incident Management (SEIM) to permit interrogation and collation of all events that might be of interest in measuring security (eg, CPU failure, hardware failure, etc.) as well as establishing alert thresholds. Further explore the design of SEIM 'dashboards' for cloud computing across different use cases (SMEs, consumers) • Explore and investigate appropriate means for interoperability of both data exchange (where possible) and enforcement of security, privacy and business policies attached to data across distributed cloud environments. • Research approaches to implementing interoperable abstraction of security, privacy and trust policies to enable data 'policy stickiness' in untrustworthy cloud environments • Research technical measures to support transparency in security for the SME/private end user of cloud (eg, using SEIM to provide for accessible assurance and compliance checking when the cloud user moves across different trust boundaries in the cloud) • Research frameworks and common specifications for handling identity access and identity assurance in the cloud (eg, mapping of identities across cloud service providers, designing suitable Trusted Third Party mechanisms to establish federated identity management for cloud computing and dealing with the identity lifecycle in the cloud) • Further research into applicability of technical means to provide for increased levels of data security across trust domains (eg, automated data expiry mechanisms and secure movement or deletion of data) • Chart flows of security meta information in cloud computing environments to enable dynamic, verifiable and attestable trustmarks and certificates

Conclusions

Successfully addressing the challenges identified in this study should not be considered an endpoint but rather a process. Furthermore, this process will necessarily involve the active, responsible and engaged participation of all stakeholders: governments, public policymakers at both the European and national level, industry and citizens. Only such a multi-stakeholder approach is a necessary pre-requisite to ensuring that the economic benefits of cloud computing can be secured whilst continuing to respect European values.

Methodology

A multi-stage methodology utilising desk research, interviews, case studies and an expert workshop was employed. These research methods were used to develop the flow of logic

from the theoretical as contained in the literature, to the practical as reported in the case studies and during the expert workshop.

Desk research and literature search

The study began with an assessment of available policy, technical and scholarly literature regarding the technological and legal components that form the basis of cloud computing.

Interviews

At the SecureCloud 2010 conference the research team conducted a number of interviews with experts and practitioners involved in investigating the security and privacy aspects of cloud computing. These were conducted using a semi-structured interview technique and were intended to further validate the issues identified in the desk research by cross-referencing them with the perspectives of practitioners.

Case studies

A framework and set of criteria was designed to support the selection of case studies. The decision was made to focus on public cloud implementations as these would contain the most interesting and useful evidence for later stages in the study. The case studies were conducted via documentary review and interviews with stakeholders. The data collection for these case studies was conducted between May and August 2010.

Expert workshop

In September 2010 an interactive expert workshop was held in Brussels, to which a variety of stakeholders from cloud users, cloud providers and industry to governments, regulators and European institutions were invited. The purpose of the workshop was to act as a 'reality check' on the issues discovered so far and also to propose policy recommendations across a range of stakeholders and taking into account European values. Attendees participated in a facilitated debate where they were asked to comment on whether the identified issues were of concern for them. Participants as were also asked to propose policy options to overcome the gaps identified during the course of the study.

Structure of the report

The remainder of this report is structured in the following way: Chapter 1 lays out the central hypothesis of this work. Chapter 2 provides a definition of cloud computing based on the available literature and also details the operational and strategic motivations for organisations looking to deploy cloud computing. In this chapter the implications for security of the economic drivers regarding cloud computing are also considered. Chapter 3 defines what is meant by security, privacy and trust for the purposes of this study and also briefly describes an increasing body of literature reflecting on the implications for security, privacy and trust. Chapters 4 and 5 delve deeper into the identified technical and legal domains which comprise cloud computing and reflect upon the challenges associated with each of these constituent elements. Chapter 6 discusses how current popular operational frameworks (such as that contained in the ISO2700x suite of Information Security Management System specifications) for managing technological- or compliance-related risk may not be viable in a cloud computing context. Chapter 7 presents three case studies in order to explore these operational challenges in practice. Chapter 8 contains a gap analysis, establishing which elements of the key European policy instruments may not be tenable in

addressing the identified operational challenges. Recommendations for each stakeholder are presented in Chapter 9. Chapter 10 concludes the report. Following a bibliography, appendices consist of a list of references, a list of interviewees, the case study framework and the agenda for the expert workshop.

CHAPTER 1 Introduction

Successfully addressing the privacy, trust and security risks inherent in deployments of cloud computing represents a complex and difficult challenge for one significant reason: the widespread use of cloud technologies brings into sharper focus linkages among issues that formerly were firmly within the purview of 'stove-piped' regulatory or technological approaches or domains.

Although the technologies and legal or compliance requirements associated with the cloud may not be new or particularly innovative (for example, the audit requirements resemble those required for outsourcing arrangements), the public policy challenge hinges on whether their combination in a cloud environment undermines or degrades current policy approaches. For example:

- Widespread take up of cloud computing might contradict or adversely affect globally established data protection and privacy principles. Consent and user 'control' of personal data are key tenets of privacy and data protection frameworks that are particularly pertinent for Europe. However, widespread implementation of cloud computing reduces users' and even cloud providers' control of personal information as it is opaquely and autonomously disassembled and re-assembled across a highly distributed infrastructure.
- The technological challenges of providing for confidentiality in the cloud present potential barriers to establishing trust. For example, encryption and similar technologies are currently the best way to maintain data confidentiality but their effectiveness is reduced in cloud environments because the cloud user has little or no other effective (logical or physical) security controls.
- Cloud computing deployment poses operational challenges in respect of compliance and risk management. For example, verifying compliance with certain regulatory or voluntary standards is significantly more complex in a cloud environment since the point at which a satisfactory level of assurance is achieved might incur greater and greater costs for the cloud user. Cloud users may have to take on trust their providers' compliance with the requisite operational procedures.

The hypothesis investigated by this research is that the complexity associated with cloud computing is sufficiently distinct from what has gone before to warrant the development of entirely new public policy approaches.

CHAPTER 2 Definitions and drivers

A definition of cloud computing based on a consolidation of the available literature is provided here, as well as an examination of the operational and strategic motivations pushing organisations and individuals to exploit the capabilities of cloud computing.

2.1 Definitions of cloud computing

The starting point for any public policy analysis is the definition of terms and concepts. In light of its novelty there are many definitions of cloud computing. Based on an analysis of twenty such definitions Vaquero et al. find that the literature appears to be converging on the following operational definition:

> *Clouds are a large pool of easily usable and accessible virtualised resources (such as hardware, development platforms and/or services). These resources can be dynamically reconfigured to adjust to a variable load (scale) allowing also for an optimum resource utilisation. This pool of resources is typically exploited by a pay-per-use model in which guarantees are offered by the Infrastructure Provider by means of customised service level agreements.*[4]

This definition emphasises a core characteristic of cloud computing: the availability of unbounded computing resources to individuals and organisations upon demand without the need for significant local infrastructure provisioning. Essentially, cloud computing should allow users to pay exclusively for the resources they actually use and reduce up-front infrastructure commitments.

Cloud computing is an evolved form of ICT outsourcing that utilises (in part) technologies developed to meet the needs of grid computing. This leads to some conflation of the use of the terms grid and cloud computing. The main difference is that grid computing was originally driven by scientific research purposes,[5] with the objective of coordinating resources which are not subject to centralised control under standard, open, general purpose protocols and interfaces. Cloud computing, in contrast, involves *provision of services* with varying proprietary and open toolkits and interfaces. Thus, cloud computing may be considered an IT service, defined as:

---

[4] Vaquero et al. (2009)

[5] E.g.[5] Nature 'web matters' (2000). Examples include the Condor Project, SETI@home and the World Community Grid

> *…a paradigm in which information is permanently stored in servers on the Internet and cached temporarily on clients that include desktops, entertainment centres, tablet computers, notebooks, wall computers, handhelds, sensors, monitors, etc.*[6]

As previously indicated, a clear definition (or coherent set of definitions) of cloud computing may be helpful for effective identification of public policy options, especially in the pan-European context across national and linguistic divides. However, while the definition proposed by Vaquero et al. emphasises one core characteristic of cloud computing – its ephemeral approach to defining and deploying computing resources – it does not consider precisely the technological and operational factors. In this context, an interesting alternative perspective is provided by the definition proposed by the US National Institute for Standards and Technology (NIST):

> *Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.*[7]

This definition, adopted for the purposes of this research, emphasises the flexibility and convenience of the cloud which allows customers to take advantage of computing resources and applications that they do not own for advancing their strategic objectives. It also emphasises the supporting technological infrastructure, considered an element of the IT supply chain that can be managed to respond to new capacity and technological service demands without the need to acquire (or expand) in-house complex infrastructures.

Building upon this definition, the literature identifies four different broad service models for cloud computing:

- *Software as a Service (SaaS)*, where applications are hosted and delivered online via a web browser offering traditional desktop functionality, eg, Google Docs, Gmail and MySAP.
- *Platform as a Service (PaaS)*, where the cloud provides the software platform for systems (as opposed to just software), the best current example being the Google App Engine.[8]
- *Infrastructure as a Service (IaaS)*, where a set of virtualised computing resources, such as storage and computing capacity, are hosted in the cloud; customers deploy and run their own software stacks to obtain services. Current examples are Amazon Elastic Compute Cloud (EC2),[9] Simple Storage Service (S3)[10] and SimpleDB.[11]

---

[6] Hewitt (2008)

[7] Mell & Grance (2009a)

[8] Ciurana (2008)

[99] See (as of 21 November 2010) http://www.amazon.com/ec2/

[10] Kirkpatrick (2006.)

[11] Murty, (2008)

- *Hardware as a Service (HaaS)*, where the cloud provides access to dedicated firmware via the Internet, eg, XEN and VMWare.[12]

The literature also differentiates cloud computing offerings by scope. In private clouds, services are provided exclusively to trusted users via a single-tenant operating environment. Essentially, an organisation's data centre delivers cloud computing services to clients who may or may not be in the premises. Public clouds are the opposite: services are offered to individuals and organisations who want to retain elasticity and accountability without absorbing the full costs of in-house infrastructures. Public cloud users are by default treated as untrustworthy. There are also hybrid clouds combining both private and public cloud service offerings. The overall cloud computing environment is evolving and, therefore, its service models may change with time.

2.2 What's pushing cloud take-up?

Due to the immaturity of the cloud ecosystem there is still a significant lack of qualitative and quantitative research about motivations for using cloud computing services. Currently, the most valuable insights come from market research literature.[13] This literature typically emphasises the economies of scale arising from commoditisation of IT services and the removal of complex on-site infrastructure deployment and management. Cloud computing, therefore, can reduce capital investments and limit the risks of over-provisioning, which is the usual response to the management of uncertain demand. The market research literature does not support these claims with detailed quantitative evidence. However, some data are available from cloud computing service providers and government organisations that are in the process of procuring such services.

According to Amazon Web Services,[14] cloud computing allows substantial economies of scale in hardware utilisation in an area where traditional in-house data centres commonly use about 30 percent of available capacity. Migration to a cloud computing environment gives 100 percent utilisation. This is especially beneficial when serving heterogeneous users with highly dynamic demands for usage and capacity. Cloud computing also reduces software updating and licensing costs (especially in the SaaS model). Amazon Web Services has provisionally quantified these savings: an organisation operating over the web having low but steady usage with occasional peaks can reduce costs by 85 percent, rising to 95 percent for organisations needing high-performance computing.

Cloud computing also leads to efficiency savings via the reallocation of staff and other resources. Again, the literature does not provide comprehensive evidence but the assertion is supported by anecdotal evidence from the business case justifying the decision by a US public institution (in this case the City of Los Angeles) to migrate towards a Google-enhanced email service for its employees.[15] The introduction of this cloud computing

[12] VMware Inc -(1999)

[13] For example, see Gartner (2009a)

[14] Amazon Web Services (2009)

[15] See (as of 23 November 2010)
http://www.informationweek.com/news/services/saas/showArticle.jhtml?articleID=221100129

service is expected to save US$6 million in terms of software licence fee and US$500,000 in hardware retirement costs. At the same time, the organisation expects to achieve additional efficiency savings of US$6 million via the reallocation of staff to other tasks and US$1 million by reallocating IT infrastructure to other activities.

Cost reduction is not the only benefit of cloud computing. The commoditisation of the overall infrastructure and related services allows organisations to focus time and resources on mission critical tasks. More importantly, cloud computing provides flexible sourcing[16]; should business environments change, organisations can adapt their infrastructure accordingly. This flexibility can also be used for testing new services or running low priority business applications such as collaboration applications for employees, consultants and suppliers.

Aside from arguments regarding the deployment of cloud computing for efficiency reasons, others have argued that cloud computing has the potential to produce an 'explosion in creativity, diversity and democratization predicated on creating ubiquitous access to high powered computing resources'.[17] In particular the broad adoption of common tools coupled with the emergent possibilities of using technological components of cloud computing (like service-orientated architectures or web services) to mash data together illustrates the potential innovative aspects of a 'cloud' way of thinking.

Nonetheless, some of the reasons why organisations are turning to the cloud reflect different user characteristics, whether they are public or private or whether they seek to obtain back-office functions (eg, accounting), email or other communication or collaboration services, or even 'niche' high-performance computing. Whilst some cloud drivers are consistent across both the public and private sectors, there are nonetheless some unique demand-side factors arising from the different emphasis put on issues such as privacy and security by public sector organisations.

Citizens and consumers also have specific reasons for using cloud computing: accessibility of their data and the almost limitless quantities of storage space, for example, may be regarded as features of cloud-hosted email services such as Gmail likely to be more attractive to individuals than to organisations.

2.3 The economics of cloud computing: implications for security

In this section, some of the basic economics of the cloud in relation to privacy and security are presented. This is not intended to provide a comprehensive guide to 'cloud economics' or to the economic policy issues raised by the deployment of cloud services.[18] Rather, it is an impression of the evolutionary landscape in which cloud businesses and business models must survive and the implications for the allocation and management of privacy and security issues.

[16] Vouk (2008)

[17] Rayport & Heyward (2009)

[18] More general treatments can be found in, for example, Picker (2008), Nelson (2009) and Etro (2009)

Cloud computing has its commercial roots in the convergence of the computing and communications industries. Economic regulation of computing has in the past been based on the 'effective competition' model, while communications (especially fixed-line communications) have generally been regulated as monopolistic or oligopolistic public utilities. There has been relatively little progress towards a hybrid model that can regulate the combined performance of ICT without distorting its development; initial forays into 'converged regulation' aimed at the Internet (eg, net neutrality, broadband spectrum auctions/trading, etc.) have aroused continuing controversy.

Two developments are of particular relevance to cloud computing and to privacy and security. The first concerns business model evolution, in particular the *relation between economic regulation and non-economic (or non-monetised) behaviour*. Economic regulation is generally linked to explicit contracts and to direct exchanges of money for services; in contrast, many new business models involve little or no direct two-party monetised relationships of this nature, but rely instead on third-party monetisation by selling the attention of free service users, bundling services with other services, equipment, and so on, or effectively investing in other value-generating relationships. Each of these can be detected in relation to different 'as a Service' cloud models because each of the 'Services' is used for a variety of purposes by a variety of customer groups. Where there exist concrete and direct monetary contractual relationships (for example, between ISPs and their customers on one side and between ISPs and those who handle their traffic on another), regulatory leverage can be used to encourage compliance with privacy and security rules (where they exist) as well as competition rules. Where no money changes hands, this may be much harder to accomplish.

The second development is the *delivery of services through extended pathways*. The interposition of communications service provider(s) between cloud service provider(s) and user(s) complicates the internalisation and monitoring of both public and private privacy and security obligations. Where the cloud service providers and users have no direct contract or connection – and may not even be aware of each others' identities or locations – it may be very difficult to ensure that costs and benefits not central to the service arrangement are handled appropriately. Cloud services themselves have some private good characteristics in the sense that access can be controlled and in the sense that they are congestible (the 'worst-case' total capacity required to serve a group can be approximated by adding the capacities required to serve each member). Suitably regulated markets *could* deliver them efficiently. But privacy and security concerns associated with overall capability to provide cloud services – or at least with their most innovative and productive uses – resemble public goods in that they can be delivered *en masse* by a 'cloud' of suppliers to a 'cloud' of users, often without a specific contract or means of charging for access. Such open clouds are not the only mode of provision, but the evolutionary contest between different models will influence the identification, pricing and management of privacy and security. Moreover, the privacy and security capabilities inherent in different cloud models will affect their acceptance by different groups and the way innovative applications develop. For instance, groups of users for whom privacy is critical may favour an explicit collective contract with a consortium of service providers who jointly and severally undertake to deliver specific levels of security and privacy as part of a bundle of services. Alternatively, privacy-conscious users may prefer to negotiate individual 'walled garden'

contracts with a telecommunications provider who will, in turn, negotiate terms with cloud service providers in exchange for access to consumers.[19]

The net result will be a more-or-less 'connected' cloud, in which service levels may converge or differ.[20] In effect, economic competition will jointly determine both the extent and structure of cloud service connections and the stringency of precautions by various parties.

This brings us to a consideration of linkages among cloud participants. Privacy and security risks when using cloud services can be propagated and controlled across different roles in the cloud value chain: service providers, platform providers, telecommunications service providers, software and equipment suppliers, intermediary and end users of cloud services, etc. End users, for example, can limit transfers of sensitive material to the cloud, enforce additional contractual or technical requirements or acquire the services of dedicated security service providers to complement their other cloud utilisation. The contracts among cloud users are of particular importance. 'Security' (used here to encompass data security, privacy and access) derives from a host of performance, security (in the ordinary sense), technical operational, cost and revenue criteria. These, in turn, are negotiated and linked through service level agreements (SLAs). SLAs are not vehicles for delivering security in the broad sense; they cover too many things and must apply in too many circumstances. They cannot even provide a comprehensive list or menu, being essentially incomplete and partially enforceable contracts. But a network of SLAs provides a governance framework for managing 'security' in concert with other performance criteria.

The efficiency, scale economies, flexibility and other cloud characteristics contributing to these benefits are all affected to a greater or lesser extent by the dynamics of cloud business models, the emerging style and form of SLAs, sectoral organisation and types of regulation. Benefits are more likely if there is sufficient variety and appropriate 'matching' among these elements; in this respect, the relatively narrow range of business models and SLAs in the 'European cloud' may be seen as at least a temporary obstacle to progress.

It remains to be seen whether proprietary or relatively closed service architectures are more or less supportive of appropriate levels and flexibility of privacy and security than open cluster or commons models, and the degree to which an evolutionary drift towards one or another model aligns security and economic objectives or brings them into conflict.

---

[19] This is a generalisation of what has come to be known as a 'two-sided market' model. See, for example, Rochet & Tirole (2003) and Armstrong (2006)

[20] For example, given the indirect relation between cloud providers and cloud users, an environment in which all parties on one side of the market adopt common or compatible approaches to privacy and security (whether minimal or extensive) would be preferable to a situation in which levels varied widely and undetectably. Models of the 'evolution of conventions' (Young (1993)) have linked dynamic behaviour in such coordination games to the structure of the underlying network. In particular, in fully connected networks, levels of security will converge to a consensus level that may not be efficient, but which is rather the level most robust against innovative, opportunistic and/or accidental departures from the norm. However, in highly clustered networks or those where network connections can be changed easily, stable diversity – even the spontaneous emergence of a two-tier cloud – is possible. In addition, the speed of convergence depends on the network structure; it is much faster in networks with strong local structures than in random or small-world networks

2.4 Concluding remarks

This first chapter has given an overview of cloud computing definitions and summarised the drivers and benefits associated with the use of cloud computing services. It is clear that that these drivers and hoped-for benefits are contingent on appropriate technological and legal responses to the many security, privacy and trust challenges raised by cloud computing.

CHAPTER 3 Understanding the implications for security, privacy and trust

Having reviewed cloud computing definitions and drivers in the previous chapter, the security, privacy and trust implications of cloud computing will now be examined.

3.1 Defining security, privacy and trust

Before assessing the literature dealing with security, privacy and trust in the cloud, it is important to define these terms because their currency and usage can change radically in different contexts. This document uses the following definitions or concepts:

- **Security** concerns the confidentiality, availability and integrity of data or information. Security may also include authentication and non-repudiation.
- **Privacy** concerns the expression of or adherence to various legal and non-legal norms. In the European context this is often understood as compliance with European data protection regulations regarding the right to private life. In the European context this is often understood as compliance with European data protection regulations. Although it would be highly complex to map cloud issues onto the full panoply of privacy and personal data protection regulatory architectures, the globally accepted privacy principles give a useful frame: consent, purpose restriction, legitimacy, transparency, data security and data subject participation.
- **Trust** revolves around 'assurance' and confidence that people, data, entities, information or processes will function or behave in expected ways. Trust may be human to human, machine to machine (eg, handshake protocols negotiated within certain protocols), human to machine (eg, when a consumer reviews a digital signature advisory notice on a web site) or machine to human (eg, when a system relies on user input and instructions without extensive verification). At a deeper level, trust might be regarded as a consequence of progress towards security or privacy objectives.

3.2 Growing focus on security, privacy and trust concerns

The rising profile of cloud computing has, understandably, spawned a plethora of policy statements, speeches, research studies, journal articles and other papers addressing the associated security, privacy and trust challenges. For example:

- The 2nd Annual Internet of Things Europe 2010 conference in June 2010 included presentations on 'Technological convergence between the Internet of Things (IoT) and Cloud Computing' and also an entire session on privacy aspects of the IoT (which shares many technological underpinnings with cloud computing) chaired by the European Data Protection Supervisor, Peter Hustinx.
- Similarly, the 3rd International Conference on Computers, Privacy and Data Protection in Brussels in January 2010 included a day-long workshop session on Security and Privacy in Cloud Computing which considered the challenges of protecting the confidentiality, availability and integrity of information processed outside the control of the data subject.
- A recent meeting of the European Security Round Table (ESRT) entitled 'Moving to the Cloud: Risks and Opportunities – Assessment for Local Entities' held in June 2010 included an extensive session on 'Privacy, Security and Data Sovereignty'.
- Meanwhile, in the United States, Vivek Kundra, the first US Federal Chief Information Officer (CIO), commented in a workshop in May 2010 that he saw security as the biggest barrier to widespread deployment of cloud computing across the US public sector. His remarks stemmed from a recent high level CIO meeting on the economic gains from cloud computing.

These discussions identify the tension between economic benefits and the costs of necessary privacy, security and trust measures as the key challenge facing those seeking to deploy or implement either cloud computing itself or associated public policy, law or regulation. This challenge can be divided into two questions:

- Given that the need for public policy depends on the way cloud computing is designed, deployed and used, what is the optimal combination of cloud service provision and governance, taking into account the existing legal and market contexts and the costs, complexity and uncertainties associated with these issues?
- Taking into account the incentives of different stakeholders, what is required to ensure that this optimal arrangement can be achieved without adversely distorting the impacts of cloud computing?

By way of an introduction to the detailed mapping that follows, some of the (policy-related) literature addressing the security, privacy and trust issues associated with cloud computing are summarised below, with common conclusions drawn out in order to set the scene for further analysis.

In December 2009 the Cloud Security Alliance (CSA), an industry forum promoting security in the cloud, published version 2.1 of its 'Security Guidance for Critical Areas of Focus in Cloud Computing',[21] addressing both cloud operation and governance. Operational areas of focus include traditional security, business continuity and disaster recovery, data centre operations, incident response notification and remediation,

[21] Cloud Security Alliance (2009)

application security, encryption and key management, and finally identity and access management. Governance areas include governance and enterprise risk management, legal and electronic discovery, compliance and audit, information life cycle management, and finally portability and interoperability.

The Cloud Computing Use Case Discussion Group is working on version 4 of its 'Use Cases White Paper for Cloud Computing'. This includes security scenarios for cloud computing including 'cloud-based Development and Testing', 'Computing Power in the Cloud' and 'Storage in the Cloud'. These security use case scenarios are generic descriptions of the security implications for security of cloud computing under different cloud providers and other scenario variables, and illustrate different aspects of security such as authentication and identity.[22]

An Expert Group Report published by European Commission Directorate-General Information Society and Media in 2010 noted that while trust, security and privacy pose technical challenges to any Internet-provided service, specific cloud computing characteristics (such as multi-tenancy and data location) serve to bring these into acute focus. Furthermore, the report argued that clouds can simplify the malicious use of resources (for example, for hacking purposes) and flagged up the need to address data protection concerns.[23]

The World Economic Forum also noted the importance of addressing security and privacy in its 'Cloud Computing Perspectives' study, which pointed out that disruption and anticipated risks might deter many organisations from embracing cloud computing. Examples include security, privacy, interoperability, governance, vendor lock-in, continuing application of outdated or complex laws, and users' lack of understanding of cloud computing.[24]

Responding to the 2009 Consultation on the Legal Framework for the Fundamental Right to Protection of Personal Data,[25] the Article 29 Data Protection Working Party and the Working Party on Police and Justice highlighted the privacy challenges created by technological developments such as cloud computing which, for example, blur the lines between data processor and data controller, and perhaps even undermine the concept of 'individual' uses of personal data (as private individuals in effect become data processors or controllers due to the quantity of personal data they store on cloud-based applications such as social networking sites).[26]

A 2010 Deloitte report on cloud computing commissioned by the UK's Centre for the Protection of National Infrastructure (CPNI) discussed a range of risks and proposed various mitigation strategies – including measures to enhance resilience. Interestingly,

---

[22] Cloud Computing Use Case Discussion Group (2010)

[23] European Commission (2010b)

[24] World Economic Forum (2010)

[25] See (as of 21 November 2010) Consultation on the legal framework for the fundamental right to protection of personal datahttp://ec.europa.eu/justice/news/consulting_public/news_consulting_0003_en.htm

[26] Article 29 Data Protection Working Party & Working Party on Police and Criminal Justice: WP168; (2009)

resilience in the specific context of cloud computing has not been emphasised in the policy (as distinct from scholarly) literature but may be considered important in understanding how risk is transferred in cloud deployments. For cloud users, establishing a stable and resilient connection to the Internet becomes all the more important once data, applications and services are migrated to the cloud. The report concluded that outsourcing of operations to the cloud does not also outsource key risks. It also noted the wide-ranging legal or regulatory issues involved in cloud computing including rights to data, security loopholes, outsourcing and subcontracting. Further questions remain regarding national laws and regulations governing interception and disclosure of data. The report indicated that the large number and geographical dispersion of third parties involved in cloud service provision make risk assessment slower, more complex and more costly. Finally, the report identified many IT data recovery risks associated with hosting data in multi-tenanted data centres including corruption of customer data, overloading of computing resources and proving the service needs meet disparate IT disaster recovery requirements.

A report by industry market research firm Gartner in 2008 differentiated risks common to many outsourced IT services from those specific to cloud computing. The former included data segregation, data privacy, privileged access service provider viability and availability and recovery. Risks unique to cloud deployments derived from location independence and the extensive complexity associated with subcontracting by cloud service providers.[27]

The European Network and Information Security Agency (ENISA) published in late 2009 the results of an extensive study on cloud computing. It indicated that the characteristics of cloud computing represent something of a double-edged sword in respect of security: the economies of scale and flexibility possible with cloud-based defences permit a robust, scalable and cost effective approach to security but the massive concentrations of resources and data in the cloud present a more attractive target to attackers. The report identified ten security benefits and ten security risks by reference to three generic use cases of cloud computing in different contexts. Security benefits included: scale (cheaper security measures when implemented on a larger scale); the emergence of security as a consideration in demand-side preferences; the use of standardised interfaces for managed security services; the ability of the cloud provider to rapidly and smartly scale resources; greater capabilities for audit and evidence gathering; improved management of updates and pre-hardened default images; and security benefits associated with resource concentration (such as cheaper physical perimeterisation). Risks included: loss of governance; data, application and service lock-in; isolation failure; compliance risks (cloud service use undermines costly certification or even precludes participation in certain markets); management interface compromise (particularly acute when combined with remote access and browser vulnerabilities); data protection (specifically for cloud users acting as data controllers required to fulfil data protection obligations such as verification of data handling processes or compliance with certain rules or procedures in respect of multiple onward transfer to jurisdictions outside the EU); insecure or incomplete data deletion (no true wiping of data); and risks from malicious insiders.[28]

[27] Gartner (2008a)

[28] ENISA (2009a)

A report by the Brookings Institution recommended that the US Federal Government take steps to both 'harmonise laws on cloud computing to… reduce inconsistencies in regard to privacy, data storage, security processes' and 'examine rules relating to privacy and security to make sure agencies have safeguards appropriate to their mission'. The report commented upon the difficulty of obtaining the full benefits of cloud computing when there are divergent and inconsistent or contradictory laws. It also suggested that appraisal processes for cloud computing should take into account amounts spent on privacy and security protections. It presents ambitious recommendations, regarding, for example, global resolution of legal inconsistencies or contradictions.[29]

A Gartner report from 2009 listed security, privacy and trust issues among the top inhibitors to cloud computing adoption. Firstly, the report argued that it was difficult to test whether a cloud provider was effectively mitigating other risks, since testing is hard to arrange and conduct and providers may not often agree to it. Another risk was broadly termed 'data location' – being unable to know where data is stored. A third risk specifically related to availability was that data or code in a cloud service provider's system may be difficult to reach, retrieve or remove. A fourth risk was the possibility of simple technical faults that could destroy the stored data. Related to this were issues connected with data retrieval in the event a cloud provider goes out of business. Finally, Gartner mentioned the data security risks of unauthorised access to data.[30] Gartner has also suggested a relationship between the extent of separation between a customer and data and the level of presumed risk.[31]

3.3 Identifying key issues and possible enablers for security, trust and privacy in the cloud

The following chapters will summarise the key challenges associated with cloud computing from technological and legal (including organisational) perspectives. This overview is based on an extensive literature review, validated by key stakeholders. As noted in the introduction to this report, the study team participated in the SecureCloud conference in Barcelona on March 16–17 2010 in order directly to seek feedback via interviews with stakeholders from various domains and enrich the analysis.[32]

It is important to note this overview of challenges is not intended to present solutions but rather to illustrate how they arise from the technical and legal characteristics of cloud computing. In this way, challenges are set up as open questions for further exploration in the final chapters of the report.

[29] Brookings Institution (2010)

[30] Gartner (2009b)

[31] Gartner (2008b)

[32] See (as of 20 November 2010) http://www.cloudsecurityalliance.org/sc2010.html

CHAPTER 4

Security, privacy and trust challenges stemming from the technological underpinnings of cloud computing

This chapter provides an analysis of the technological challenges of cloud computing and associated services, and will support the argument that the benefits of using clouds hinge on finding appropriate technological answers to the security, privacy and trust challenges. The starting point of this analysis is that cloud technologies are on the whole not new, but that their development has been revitalised by cloud computing. Development paths as they appear in the literature will be examined.

Taking into consideration the US NIST definition, a number of challenges for security, privacy and trust from the underlying technological drivers of cloud computing can be distilled, namely virtualisation technology, grid computing, web services, service-orientated architectures, web application frameworks and encryption.

4.1 The linchpin of trust: the hypervisor

Virtualisation is a key component in the provision of cloud infrastructure services (computation and data storage) as it enables providers to marry efficient use of hardware resources and multiple customers by using the same physical machine for different applications (as demonstrated by Amazon's EC2 offering). Virtualisation is not a new idea and indeed there is already a range of different concepts of virtualisation ranging from the process level execution of a Java program within a Java Virtual Machine, to the system level virtualisation execution of an operating system such as Windows 2003 within a virtual machine monitor environment like VMware. While there are varying definitions of virtualisation,[33] its importance in cloud computing focuses specifically on system level virtualisation and predominately what is often referred to as full or native hardware virtualisation. In this type, a virtual machine monitor (or hypervisor) replicates the physical machine logically, enabling multiple guest virtual machines (containing one of a number of different operating systems) to run independently. This is referred to as Type 1 hardware virtualisation.

[33] Scope Alliance (2008)

The key concerns identified within the literature concerning virtualisation within cloud computing focus on the security, management of large-scale installations, and interoperability. While virtualisation technology can be considered mature and well utilised within the data centre market, interoperability and virtual machine management are also core issues. A number of vendors, such as Amazon, expose interfaces to enable remote management. They focus on balancing demand within their particular cloud, as opposed to enabling cross-cloud migration of virtual machines and their associated management. However, a number of research activities, such as RESERVOIR[34] are focused on exactly this large-scale demand management across multiple cloud providers, and are demonstrating techniques to migrate live virtual machines between providers and, as such, enhancing open source tools such as OpenNebula[35] to be widely used by the community.

In essence, this technology has been the focus of many cloud computing initiatives. There are, however, a number of key security and trust implications. Seated just above the basic hardware layer, the technology is enabled to intercept (and modify) every processing operation performed further up the software stack, and therefore potentially has the most trustworthy requirements of the cloud layers. At the same time, it also provides a shared platform for multiple customers and their virtual machines. The isolation and separation of those virtual machine activities is critical to guaranteeing the confidentiality and integrity of the operations, without interference from other customers.

Key standards developments are predominately those such as the Open Virtualisation Format (OVF) by the Distributed Management Task Force (DMTF),[36] which underlies a common challenge: that while abstraction supports greater interoperability, it also impacts performance. A number of central processing unit (CPU) manufacturers have for some years been developing virtualisation support in order to accelerate this performance (eg, Intel VT-x and AMD-V).

The security aspects covered by the virtualisation literature are primarily in two areas: the first examines the problems in isolation and the occurrence of vulnerabilities within the configuration and development of hypervisors and how they may be exploited. In this instance, the ability to assure the hypervisor is critical in determining the level of confidence and trust placed in the security solutions. However, the second area, which examines the potential challenges for providing current security controls to a virtualised environment, also requires that there is confidence in the integrity of its behaviour as well as its ability to monitor at a granular and appropriately detailed level security event data.

There is considerable literature relating to the isolation of system processes, with many related papers in areas such as access control, enforcement techniques and policy specifications. Within the specific cloud-based papers, literature focuses on the current weaknesses in the management of virtual resources and their ability to be manipulated by external attackers.[37] Both areas appear to be important routes to compromising a complete

[34] Rochwerger et al. (2009)

[35] See (as of 24 November 2010) http://www.opennebula.org

[36] See (as of 21 November 2010) http://www.dmtf.org

[37] Ristenpart et al. (2009)

cloud, if the hypervisor can be successfully compromised and then the compromise can be propagated through the movement of guest OS machines around a cloud's virtualised environment. However, weaknesses within the hypervisor are a research focus, with vulnerabilities demonstrated through a variety of studies such as those undertaken by Wojtczuk et al.[38] and King et al.,[39] exposed by Cloudburst[40] and described in a comparative analysis by Ormandy.[41] Commentators observe that there have been relatively few vulnerabilities discovered compared to popular operating systems and attribute this to the reduced codebase and therefore its reduced attack surface. They also point out that this gives a greater potential for stricter design separation of functions and the application of formal techniques to give greater assurance and trust, although, as of yet, there appear to be only a small number of technologies to which this has been applied.

In contrast, security researchers including Sailer et al. at IBM[42] and market analysts such as Gartner[43] have looked at methods for establishing trust in the hypervisors using the Trusted Computing Platform Module (TPM) to attest to the integrity of its software, guaranteeing it, for example, to be free from malicious or accidental modification of its core operating system. This is in addition to mechanisms to establish trust within virtual machines themselves as they are transferred between hypervisors and physical servers, as described by Perez et al.[44]

Additional literature demonstrates a focus on emulating and devising specific security controls, such as monitoring, that create potential performance issues should they be implemented using current appliance-based technology.[45] A number of surveys conducted at Stanford University[46] and the University of Illinois at Urbana-Champaign,[47] and by Rosenblum & Garfinkel[48] and van Cleeff et al.[49] have attempted to identify the core security problems associated with virtual machines. While they identify the current

[38] Another Way to Circumvent Intel® Trusted Execution Technology Tricking SENTER into misconfiguring VT-d via SINIT bug exploitation. Wojtczuk et al, . (2009)Rutkowska, Tereshkin. December 2009. http://invisiblethingslab.com/resources/misc09/Another%20TXT%20Attack.pdf

[39] King et al. (2006)

[40] benchai7 (2009)

[41] Ormandy (2007)

[42] Sailer et al. (2005)

[43] Gartner (2007)

[44] Perez et al. (2008)

[45] Sharif et al. (2009)

[46] When Virtual is Harder than Real: Security Challenges in Virtual Machine Based Computing Environments, Tal Garfinkel Mendel & Rosenblum (2005) http://www.stanford.edu/~talg/papers/HOTOS05/virtual-harder-hotos05.pdf

[47] A Survey of Virtualization Techniques Focusing on Secure On-Demand Cluster Computing Nadir Kiyanclar (University of Illinois at Urbana-Champaign, National Center for Supercomputing Applications, November 2, 2005)

[48] Rosenblum & Garfinkel (2005)

[49] Van Cleeff et al. (2009)

important generic issues they require re-examination relative to current architectures and business usage, especially as hardware support is developed.

In essence, the key security and trust themes from the literature are:

- Assurance of the hypervisors' ability to isolate and establish trust for guest or hosted virtual machines is critical, as this forms the root node for multi-tenant machine computing and thus could prove to be a single point of failure, since the hypervisor can potentially modify or intercept all guest OS processing.
- The same properties of the hypervisor, which enable it to inspect and monitor all processing within and between guest OSs, give the potential for enhanced security monitoring, but will require that current security controls based on dedicated appliances can be migrated to virtual machine architectures. They could also lead to a potential loss of individual customer privacy and security.
- For economic purposes, the ability of large-scale instances of virtual machines to be dynamically moved and re-provisioned is vital. It is unclear at this point how adequate the lifecycle management of those instances between hardware and across clouds is, and whether trust can be established to an adequate level, if at all.

4.2 Can the distributed models of computation characteristic of grid technology adequately serve the availability and interoperability needs of cloud computing?

Grid technology is the highly distributed use of loosely coupled computation and storage resources, often across multiple domains. The grid concept has been in existence for some years, and this has resulted in a number of applicable technologies, commercial services and research within this field and numerous publications. As discussed in the previous chapter, some literature suggests that cloud computing is a continuation or extension of grid technology. This is not precisely the case, although there are some similarities to take into consideration.

From the literature, the key technological areas most relevant to cloud computing can be broken down into the following areas:

- Resource allocation
- Contract or service level negotiation
- Orchestration of workflow
- Distributed file system storage
- Federated identity management.

In grid environments the processing of large data sets for computationally intensive applications is highly important. Literature within this area identifies the key issue of ensuring that there is resilience across the computational or storage nodes. As the number of nodes increases the probability of failure of one or more of them across the distributed

cloud is heightened.[50] This state of affairs introduces challenges when considering the availability of cloud services using these techniques. A particularly interesting summary of an ACM workshop with researchers and practitioners in the distributed computing field further highlights the issue of potential scalability and resilience issues that cloud computing providers must tackle.[51] The implementation of loosely coupled and 'eventually consistent' systems may have some implications for sectors which require information within certain expected time limits and may make associated commercial rationale unattractive for cloud users requiring guarantees of service availability. Since many of these techniques are used for large-scale storage and distribution of data, this latter aspect of security (availability) requires careful examination in the context of different cloud computing service platforms.

With significant distributed resources across multiple domains (business or administrative), the ability to manage and distribute the workload is seen as critical to the cloud environment. Grid research and projects have developed important distributed programming techniques and frameworks such as MapReduce and Hadoop. Of particular relevance within the literature is the development of more advanced concepts to orchestrate complex resource requests into a coherent task. This continues to be investigated in research projects by the EU-funded FP7 RESERVOIR (Resources and Services Virtualization without Barriers) project[52] and also as part of other projects in the US. With major manufacturers developing infrastructure components which are policy aware, literature focuses on the ability of the systems to negotiate and renegotiate the service level agreements (SLAs), with projects within the EU developing SLA-aware infrastructure such as SLA@SOI. Additional projects including BEinGRID have piloted and reported on the legal and technological issues arising from practical grid usage. There have also been some attempts to examine the business aspects. However, while the literature and in particular the research has focused on resource planning and execution, there is a limited literature on the encoding of security and privacy parameters into machine readable contracts. Where effort has been expended on encoding these parameters, researchers have run into difficulties. For example, security is difficult to encode as a parameter with most instances relying on time-based measurements. So while it is particularly appropriate to describe availability in the amount of hours of uptime or minimum downtime, the same philosophy is extended to security functions such as patch management. In this case the time taken to identify, test and apply a security patch or new virus signature becomes the predominant metric, when a metric measuring the effectiveness or resistance of the service to attack or compromise may be far more beneficial to the service. The simple fact is that effective metrics to measure and thus encode into contract parameters are not currently available.

Privacy also suffers from related issues, as its attributes, including the data protection principles, do not have established measures. How would one measure consent, for example? Is it based on the number of times consent is sought when personal data are

[50] Vogels (2008)

[51] Birman et al. (2009)

[52] Beco (2008)

processed or when it is obtained or even not given? The problem is further compounded by a reluctance to encode and enable machine readable contracts which essentially are open to interpretation by various courts and regulatory bodies. The consequence is that until these issues are more readily encoded with sensible metrics and the legal basis for automated negotiation understood, progress will still be limited in this area.

The literature, such as that from NIST,[53] frequently identifies the lack of interoperability within cloud environments. Research and subsequently commercial services have been executed that mitigate this situation for consumers of public cloud services. The generic concept used is that of a broker. While this is neither new to grid, or indeed service-orientated architecture (SOA) environments, the attention has focused on developing mechanisms to select different services based upon economic and quality of service (QoS) parameters, as well as potentially replacing traditional expensive content delivery networks (CDNs) with multiple cloud suppliers, for example, the metaCDN project.[54] The salient point to conclude from this literature is that grid technology and indeed research is enabling the creation of global cloud services from potentially large chains of smaller suppliers, and in so doing is developing a marketplace, despite the lack of interoperability, by using and developing still further the resource allocation, negotiation and orchestration areas.

Security provisions within grid technology appear to be dominated by the Globus Toolkit that was created for large-scale scientific applications. Based on the available literature, however, security practices within this field do not appear to have been adopted by current cloud providers.[55] Additionally, distributed programming frameworks such as Hadoop have to date had minimal security controls built into them.[56]

The scientific focus has led to minimal evolution of security techniques and practices in the field of grid computing technology. That is not to say tools and schemes have not been applied or suggested. Privacy and trust, on the other hand, are less frequently referred to, essentially because in the majority of applications the processing of personal or sensitive data is not the main use, and because trust is predominately equated with security and not assurance from the customer and end users perspective as is the case for this study.

4.3 Current state-of-the-art web services may not be sufficient to establish interoperability for identity management in the cloud

From the literature it appears that a number of cloud computing services, such as Amazon (EC2 and S3), Sun and Eucalyptus,[57] provide a web services interface for their offerings. These are based on web service standards such as SOAP, WSDL and UDDI to provide a communication format, interface definition and directory for discovery of provided

[53] Mell & Grance (2009b)

[54] Buyya et al. (2008)

[55] Amazon Web Services (2010)

[56] O'Malley (2009)

[57] Nurmi et al. (2009)

services. Similarly to grid computing, decisions on the particular web service to select are important and thus the issues surrounding negotiation and the specification of SLAs (through standards such as WS-Agreement[58]) appear to be important topics for cloud computing.[59] It also appears that current research and initial services are focused on incorporating functional and non-functional requirements into SLAs with the challenges of specifying security and privacy requirements tackled primarily in other related research projects, such as those conducted by Monahan & Yearworth at HP Labs.[60]

The web services security mechanisms currently documented and deployed are well-established mechanisms and present nothing particularly novel for cloud environments, requiring users to generate public/private key pairs and subsequently verify the integrity of requests using hash functions. Confidentiality is protected via standard security protocols such as SSL, with web services security standards such as WS-Security evidenced in products offered, for example, by Amazon Web Services.[61]

Identity management, authorisation and access control within cloud services are frequently cited as a challenge for cloud computing, for instance in guidelines from the Cloud Security Alliance.[62] The challenge is particularly acute with multiple identity management schemes for the corporate (eg, SAML, WS-Federation, Liberty Alliance) and consumer (eg, OpenID, Information Cards, OAuth) markets reducing the chances of interoperability for identity management. Although federated identity management schemes are commonly seen as the current best solution for authentication and identification on the Internet, the commercial dynamics inherent in cloud computing as regards the current concern over data portability and interoperability may serve to undermine the achievement of a realistic scheme. There are a number of cloud providers that are supporting and developing identity management products for their cloud services (such as Microsoft Azure).[63] The development of identity providers and associated technical mechanisms are of considerable importance to cloud services across the different layers of abstraction (eg, IaaS, PaaS and SaaS). Indeed, establishing identity is not restricted to the service provider verifying the user: the identity of services also has to be established for the user.[64] (For example, at the Defcon 2009 conference the identity of virtual machine images was manipulated to demonstrate an attack on consumers of the Amazon EC2 service.[65]

[58] See (as of 21 November 2010) http://www.ogf.org/documents/GFD.107.pdf

[59] Buyya et al. (2008)

[60] Monahan & Yearworth (2008)

[61] Amazon Web Services (2010); Nurmi et al. (2009)

[62] Cloud Security Alliance (2009); Mather et al. (2009)

[63] Cameron (2008); Cameron (2009)

[64] Cloud Security Alliance (2009)

[65] SensePost (2009)

4.4 Trustworthiness in service-orientated architectures (SOAs)

While SOA is frequently cited as a key technology of cloud computing, most literature appears to refer to its importance as an architectural concept implemented typically by the web services technologies referenced earlier. A more precise definition can be found within the OASIS and Open Group bodies. In the context of this study, SOA is a method of constructing software with reusable services interacting through standard interfaces and protocols. These can be using web services, as is so often the case in a public/private environment, or Distributed Component Object Model (DCOM) or Common Object Request Broker Architecture (CORBA) in a private environment.[66]

Security of the SOA interfaces and protocols has been reviewed in the context of their individual technology specifications (eg, web services, CORBA, etc.) within literature independent of cloud computing. However, at least one discussion identified in the *Security and Privacy* journal of the IEEE in 2007[67] has focused on the ability to handle these within a heterogeneous SOA environment where multiple security credentials between different technology layers can be chained together. In reality, the solution appears flawed at best for a malicious cloud environment; however it is interesting to see discussion on verifying the integrity of these layered credentials in a distributed system reliant upon different security credentials.

4.5 Will web application frameworks (APIs and SDKs) be credible in providing trust across distributed environments?

Application processing interfaces (APIs) and software development kits (SDKs) exist at varying levels of abstraction within the cloud (eg, IaaS and PaaS). While the APIs are currently well documented and 'open', they are essentially proprietary in nature and often the underlying implementation is closed source (eg, Google Apps and Microsoft Azure). The proprietary nature, closed source and varying abstraction level reduce the ability of applications to be ported between cloud implementations. However, the standardisation of APIs appears to be being pursued by projects such as the Cloud Computing Interoperability Forum initiative,[68] which is providing a metadata description of the APIs as they develop rather than trying to standardise on a particular vendor's API. There appears to be very little literature to date on standardising security APIs.

Interestingly, many providers also provide development environments within the cloud and thus, given the closed nature of the environments, it's not clear how the integrity of code and intellectual property will be protected. However, these issues appear to be general to both cloud environments as well as other collaborative open-source projects. It remains to be seen whether the drivers acting upon code and intellectual property development in collaborative open-source contexts and cloud environments manifest themselves in

[66] Laplante et al. (2008)

[67] Dushin & Newcomer (2007)

[68] See (as of 21 November 2010) http://code.google.com/p/unifiedcloud/

markedly different ways and whether, for example, parallel results may occur in respect of the levels of security.

The emergence of 'rich Internet applications', and various programming languages and models, do not appear to cause particular concern regarding cloud environments per se within the literature. Instead they highlight the vulnerabilities that commonly occur within web applications and new methods, such as Ajax, that make it particularly problematic for individual users. For example, the emergence of applications utilising third party applications, with potentially unknown provenance and security, can result in attacks focused on cross-site scripting/cross-site request forgery vulnerabilities. Ajax, which enables background processing, potentially reduces the user involvement in these transactions and thus increases the chance of success. A comprehensive identification of these issues has been undertaken by ENISA.[69] It may be seen that the stakeholder group most affected by these challenges would be consumers and citizens, since the consumer may be the most appropriate and responsible party to mitigate the risk from the appearance of such challenges.

4.6 The fragility of current encryption approaches in the cloud context

Encryption is a core mechanism for maintaining the confidentiality of all data, whether it consists of business, personal or sensitive information, and it can also be used to establish the integrity of various transactions, code and data. Thus while it is often considered a security control, it is such an important one that it requires special discussion in relation to its limits on maintaining confidentiality and integrity within the cloud computing environment.

The uses of encryption in accessing services in the cloud are broadly similar to any other outsourced service provision currently, and can use as much (or as little) data protection as conventional technologies, or instead rely on binding commercial terms between customer and provider to offload responsibility to the latter. In many public offerings, services are provided via an HTTPS-protocol connection to a web service, and so benefit from Secure Socket Layer (SSL) protection of data in transit. There appears to be no discussion in the literature for deriving additional cryptographic communication protocols that are cloud specific.

However, within the literature the protection of the data during processing or while stored (as described by the Cloud Security Alliance) is of considerable concern.[70] Current uses of the cloud either recommend the storage of non-critical data only,[71] or, in the case of security guidance from the Cloud Security Alliance, that it is encrypted prior to storage and only decrypted when back at consumer premises.

Where the data must be processed, as is so often the case with services at the SaaS and PaaS levels of cloud abstraction, such security techniques are only helpful against an outside

[69] ENISA (2008)

[70] Cloud Security Alliance (2009)

[71] Mather et al. (2009)

attacker, and to some extent a malicious insider within one of the organisations. They provide no real protection against misbehaviour of the provider itself: whether the data are transmitted or stored in encrypted form, processing is necessarily carried out on plain text. This is of course true for conventional outsourcing solutions, too, but if cloud resources become more commoditised and the data may be in the hands of third, fourth and nth parties, the trust inherent in direct business-to-business transactions may become strained to the extent where this becomes a significant concern.[72]

But is it necessary that the data be decrypted for processing? As Kerschbaum & Deitos argue,[73] security against the business partner may be enhanced by technological mechanisms. They advocate secure multiparty computation (SMC), which allows a number of (presumably non-colluding) computing agents to cooperate in performing a calculation, where none is aware of the complete set of inputs, so that any who end up possessing the result are little the wiser. Such techniques[74] are quite inefficient in terms of communication and processing demands, and also add extra administrative and cost burdens in arranging diverse resource provision.

A second technological solution is what Smart & Vercauteren describe as the 'holy grail of cryptography for a very long time', a fully homomorphic public-key encryption scheme.[75] A homomorphic encryption scheme is one where some (arithmetic) operations on plaintext can be executed by an agent possessing the encryption key and ciphertext, but without decrypting the latter. At one extreme such properties can manifest as the vulnerability of an encryption scheme known as malleability; if the operation can be executed without possessing the encryption key that is certainly a weakness.

Some well-known schemes, such as RSA and ElGamal, exhibit the property that the product of ciphertexts is an encryption of the product of the corresponding plaintexts. Other schemes, where the data are effectively used as the exponent in the encryption calculation, have the property that the product of ciphertexts encrypts the sum of the plaintexts. Reasonably efficient schemes are known for supporting addition and multiplication by a known constant (and so affine transformations), but this is rarely adequate for business purposes. This was until recently the state of the art: Fontaine & Galand provide a good baseline survey of that state as of 2007.[76]

In June 2009 a major advance was announced: Gentry's fully homomorphic scheme (supporting both multiplication and addition of unknown plaintexts) based on an encryption using ideal lattices.[77] Other less well-publicised schemes emerged around the same time (before as well as after) which allow calculation of polynomials of any given degree over plaintexts,[78] and a range of other transformations.[79] Smart & Vercauteren have

[72] Chow et al. (2009)

[73] Kerschbaum & Deitos (2008)

[74] Cramer, et al. (2002); Yokoo & Suzuki (2002.)

[75] Smart & Vercauteren (2009)

[76] Fontaine & Galand (2007.)

[77] Gentry, (2009)

[78] Aguilar Melchor, et al. (2008)

followed Gentry's line of development to develop a fully homomorphic scheme with relatively small key and ciphertext size.[80]

These developments are promising, but they remain impractical for immediate deployment: some schemes require a megabit of ciphertext for each number, while Smart & Vercauteren conclude of their 'efficient' scheme:

> *We see that in practice our scheme appears to obtain a better depth of decryption circuit than theory predicts, although still not deep enough to enable fully homomorphic encryption; at least at practical key sizes.*[81]

On the other hand, negative results in the literature are rather weak,[82] meaning that there is no known theoretical bar to achieving a practical fully homomorphic scheme. This remains an active research topic and, given the existence proof provided by Gentry's scheme, is likely to receive even more attention in the next few years. It remains cutting-edge research, so it would be brave to predict when exploitable results may emerge, but any suitably efficient scheme is likely to present little implementation difficulty; so usable fully homomorphic encryption may be just over the horizon.

The numerical case only addresses a small class of data processing tasks, of course. Literature has not introduced any serious research into the word-processing of strongly encrypted documents. Some of the more data-intensive tasks, such as the private information retrieval problem have been studied,[83] but here again the scenario is generally one of data distributed among multiple non-colluding providers.[84]

The presentations and interviews at the SecureCloud 2010 conference largely confirmed the analysis above and helped to identify key priorities and concerns.

A primary observation was that cloud computing service providers largely feel confident in the maturity and security of their solutions, but that it was much harder to define security in terms that could be measured through objective metrics. This is particularly important for potential users of cloud computing services: in the absence of direct control over the available infrastructure, suitable approaches must be found to test the cloud system, and metrics should be found to measure and assess its security.

The availability and accessibility of sufficiently extensive logging systems was mentioned repeatedly as a supporting element to allow end users to identify and examine security incidents. An additional problem here is ensuring the trustworthiness of such logs, since

---

[79] Prabhakaran & Rosulek, (2008)

[80] Smart & Vercauteren (2009)

[81] Ibid.

[82] Ostrovsky & Skeith (2007); Prabhakaran & Rosulek (2008)

[83] Ostrovsky & Skeith (2007)

[84] Molina et al. (2009); Sang & Shen (2009); Vaidya et al. (2008)

they will typically be kept by the cloud service provider (who may have a commercial interest in obfuscating incidents).

For both of these issues (metrics and logging), several guidelines and methodologies are being developed, but there does not appear to be much consensus yet on best practices and common approaches to applying them, and how they are used operationally for such things as incident response. In these areas in particular, there is still much margin for further discussion and progress.

In addition, it was noted that the combination of several flexible but complex technologies leads to additional security challenges. Weaknesses in the hypervisor technology used by a specific cloud service provider, for example, may affect the security of any system that it manages or hosts. In that sense, the fact that cloud computing can scale so easily (and therefore accommodate a large number of systems on a more compact shared infrastructure) creates a new risk: cloud service providers make for attractive attack targets, since a single weakness can be exploited against a potentially large number of entities.

Of course, this risk may be partially offset due to a higher level of security services within cloud providers, whose entire business model depends on a perception of security and availability. For SMEs (whose own security practices are typically less robust), a migration to cloud services may well have a beneficial impact in terms of security.

4.7 Concluding remarks

This chapter has provided an overview of the core technologies underlying cloud computing, and some of options for addressing security and privacy concerns. In Table 1, below, an overall mapping of the challenges posed by each technological building block in respect of security, privacy and trust is presented. It is important to note specific caveats in respect of what is being shown – this mapping is based on those issues identified from the literature reviewed during the course of this study as being particularly noteworthy or important. As such it may not be exhaustive and may not constitute a complete representation given the rapid and dynamic pace of technological change. Furthermore, specific technological flaws or vulnerabilities are clearly part of the story. As will be seen in later chapters, it is the governance and organisational controls applied to the use of these technologies that is the real determinant of the resulting level of maturity in respect of security or privacy. A highly 'secure' technology may be rendered highly insecure by lax security controls and poor governance.

Table 1. Mapping of key challenges in respect of the technological underpinnings of cloud computing arising from the studied literature

Area / Technology	Security	Privacy	Trust
Virtualisation	Integrity	Segregation of personal data on shared infrastructure	Compromised virtual machines/ hypervisors permit loss of trust
Grid technology	Availability		Interoperability
Web services	Integrity and	Security and	Interoperability (in the

	confidentiality	confidentiality	context of identity and access management)
Service-orientated architectures	Integrity		The reliance of distributed systems on different security credentials
Web application frameworks	Integrity and availability		Trust across distributed environments
Encryption in the cloud context	Confidentiality	Security and confidentiality	

As can be seen from the table, the key inherent challenges identified from the literature reviewed during this study in general revolve around issues of availability and integrity: understandable given the technological capability to drive the storage and processing of information (and any subsequent assurance of the way in which the risks to this information are managed) outside of the immediate control or supervision of the commissioning enterprise or end user. This is not to say that other implications, for example, in terms of confidentiality, might arise through vulnerabilities discovered in a particular technology. The challenge of how to provide for identity management across distributed cloud environments whilst maintaining the necessary degree of trust appears to be intractable given the implementation of current technologies. Furthermore the presence of the issue of interoperability thrown up by the challenges associated with some technological underpinnings illustrates a somewhat broader (perhaps operational) interpretation of security – namely that interoperability can help to provide for availability in ensuring data portability. Indeed, this may not even be in respect of data portability but rather any cloud-hosted asset such as a particular machine configuration developed with specific suite of apps.

Technology can only provide part of the response. Legal interventions may be required, and these are discussed in the next chapter.

CHAPTER 5

Security, privacy and trust challenges inherent to the legal and regulatory aspects of cloud computing

Against the backdrop of the technologically orientated challenges introduced in the previous chapters, it is clear that there are also substantial legal aspects to be taken into consideration for the provisions of cloud computing services. While these challenges are global in nature, the normative response may vary substantially from region to region or even from service to service. Diverging interpretations and legal uncertainties could well endanger the development of innovative cloud service models, as they can adversely affect the trustworthiness of such services: how can users invest in the cloud without a clear perspective on the compliance of the chosen solution with the applicable legal framework, or on the guarantees offered by the service provider?

From a European perspective, there is the dual risk that excessively restrictive regulations can place European cloud service providers at a clear competitive disadvantage in relation to their non-European counterparts, and inversely that overly flexible rules could result in serious harm to end users. Ambiguities in the law should at any rate be identified and addressed to ensure that the legal status of cloud computing services and the rights and obligations of each of the stakeholders is clear, especially when dealing with security and privacy concerns.

Legal barriers can arise in a number of areas, each of which have clear privacy and trust implications, as will be commented on in the sections below. As with any type of service, there are the horizontal issues of determining the applicable law and the competent jurisdictions to settle any disputes that may arise. Furthermore, the very nature of cloud computing services will result in specific laws becoming applicable with respect to certain vertical issues, including regulations covering data protection, intellectual property rights, confidentiality, e-commerce, taxation, etc. Most of these regulations will be phrased in general terms which build on the assumption of a fairly static business reality, and will therefore be difficult to apply to the much more fluid environments created through cloud computing services. Finally, the legal framework in relation to such services will also be partially determined by the will of the participants as expressed in specific contracts. Contractual issues will thus occupy a central place in defining the scope of the services and the rights and responsibilities of the parties involved. In the following sections, each of these issues is to be examined in light of the analyses provided by the available literature.

5.1 Horizontal perspective: applicable law and jurisdiction

A first question is the horizontal perspective, and specifically the identification of applicable laws and competent jurisdictions for cloud computing services. These are not trivial issues, considering the fact that they are traditionally addressed mainly on the basis of the geographical location of one or more of the stakeholders involved. In a cloud computing model however, there is likely to be a multitude of stakeholders, some of whom may not even be clearly identified, and all of whom may be located in different jurisdictions.

This issue is fundamental when addressing trust and privacy in the cloud from a legal perspective: the rights and obligations of each stakeholder are determined by applicable law, as are the available avenues for legal recourse. To ensure adequate legal certainty, rules need to be either sufficiently harmonised or based on a clear and transparent contractual framework.

5.1.1 Establishing the legal foundation of trust: how to determine applicable law in the cloud?

The identification of applicable laws in the absence of any explicit choice by the parties involved is difficult in relation to any information society service, and cloud computing service models are certainly no exception.[85] In a European context, the provisions of the eCommerce Directive play a central role, as it contains specific rules on applicable law for information society services (in Article 3).[86] However, it is clear that this will be insufficient to address all questions in this domain: the rules established by the Directive obviously apply only in Member States, and in a non-European international context will not be able to solve conflicts of law. In addition, applicability of the law remains linked to the geographical location of the information society service provider, and in a cloud model it may be difficult to identify this entity or its geographical location. Finally, certain issues – including contractual consumer protection clauses and intellectual property protection – are excluded from the Directive's scope, meaning that answers to conflicts of law in these domains will have to be sought elsewhere. Thus, it is already very complicated to identify the starting point for the establishment of trust, namely the specific laws that will apply in the absence of a choice by the parties. Globally, voluntary choice of applicable law by the stakeholders in a cloud service model may be the only viable solution to identify applicable law. In practice, the importance of this issue should not be overstated, as the choice of an applicable legal system on a contractual basis has indeed become standard practice in information society service contracts.

5.1.2 Handling disputes in the cloud: how to reinforce trust by building in a mechanism for accountability

When disputes arise between parties that they are unable to resolve amicably, they may need to bring their issues before the competent court. Again, the unique characteristics of cloud computing services can make it very difficult to identify the competent jurisdiction. Nonetheless, this issue is crucial in order to ensure the trustworthiness of cloud services: in

[85] Mills (2009)

[86] See (as of 22 November 2010)
http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32000L0031:EN:HTML

the absence of a framework for handling disputes, there may be little accountability behind cloud services, which would fundamentally hamper their usefulness in practice.

In Europe, this question is largely governed by the BRUSSELS I Regulation, but this legal framework is confronted with the same inherent difficulties as the eCommerce Directive: it is not applicable outside the EU, the final outcome is linked to the physical location of the stakeholders (typically the place of establishment of the defendant), and certain areas of law are excluded from its scope. Thus, here too, alternative mechanisms of deciding the competent jurisdiction (principally voluntary choice by the parties) will need to be considered, as well as alternative conflict resolution mechanisms, including mediation and binding or non-binding arbitration. Again, from a purely practical perspective, the contractual identification of dispute resolution mechanisms (primarily by explicitly designating competent courts) is a standard practice in information society service contracts, so this cannot be considered a crucial issue in practice.

5.2 Vertical issues: main applicable laws

Apart from the horizontal issues related to addressing conflicts of law and jurisdiction, vertical issues – related to a specific domain of law – also need to find an appropriate legal answer. Without attempting to be exhaustive, the domains below will have to be addressed by any cloud service provider.

5.2.1 The applicability and application of current data protection and privacy frameworks in the cloud context

One of the most manifest and visible problems in relation to cloud computing relates to the protection of individuals against abuses of their personal data. Europe has put in place a relatively strict data protection framework via the European Data Protection Directive 95/46/EC,[87] but there are several difficulties in applying these rules in a cloud computing environment as described in a 2009 report by ENISA.[88]

The first issue is determining applicable data protection law. While the Data Protection Directive has provided a certain degree of harmonisation in the EU, differences still remain between national legal frameworks, for example, in relation to formalities to be completed by the data controller or specific data protection regimes (health data, financial data, etc.). The Directive specifies that national laws will generally apply when personal data processing is carried out in the context of the activities of an establishment of the controller on the territory of the Member State. It should be noted that an 'establishment' doesn't necessarily require a formal legal seat: whether an 'establishment' exists should be determined under national applicable law. In practice, having a local office or other type of local presence targeting the local market is usually sufficient.

If the controller is not established on EU territory, the law of a Member State still applies if he make makes use of equipment situated on the territory of that Member State for the purposes of processing personal data, other than for strict transit purposes (Article 4 of the

[87] See (as of 22nd November 2010) http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:en:HTML

[88] ENISA (2009a)

Directive). The presence of local data centres would be sufficient to meet this requirement, but the Article 29 Data Protection Working Party also explicitly noted cookies or locally deployed software as examples of 'equipment' in the sense of this provision.[89] As a result, the Data Protection Directive has a very broad sphere of applicability, at least if the (authoritative but not binding) interpretation of the Working Party is followed.

This ensures that EU data protection rules will often apply to cloud services, which should provide a relatively high level of privacy protection for end users, promoting the trustworthiness of cloud services. However, as a side effect, the data protection laws of multiple Member States may also apply to the activities of a single cloud service provider. This raises the complexity and costs of compliance to a very significant degree, and thus provides a strong incentive for non-compliance if penalties are perceived as unlikely or too low. Compliance costs may otherwise prove to be a real barrier to the creation of cloud computing services in the EU.[90]

Secondly, the basic concepts defined by the Directive may prove exceedingly difficult to apply in such an environment, including the notions of data controller and data processor. These roles are defined rather statically in the Directive, while the reality in a cloud computing service may well be that individual entities change roles quickly as personal data are moved, restructured and re-used continuously, as described in a recent paper by the influential Article 29 Data Protection Working Party.[91] For that matter, the concept of personal data itself is ambiguous in such an environment, as two sets of data that were previously anonymous (and therefore not personal data) may become personal data once combined and processed within the cloud, suddenly changing applicable laws. The fact that the legitimacy of specific acts of data processing must be judged in a specific context can also lead to difficulties, as the use of personal data within a cloud can evolve and change drastically.[92] More fundamentally, the question has been raised (especially in the US) whether a reasonable expectation of privacy can exist with respect to data entrusted to a cloud computing service.[93]

Finally, at the international level alternative regulatory frameworks like the Organisation for Economic Co-operation and Development (OECD) guidelines or the Asia-Pacific Economic Cooperation (APEC) framework may have had a strong influence, but it is unclear whether these rules will be effective to protect individuals in a cloud computing environment, or to what extent they will prove to be enforceable.[94]

Thus, the data protection challenges in relation to cloud computing are significant, certainly in a European context, given the specific and relatively strict data protection framework that has been put in place. Indeed, confronting these expectations with

[89] Article 29 Data Protection Working Party (2002)

[90] Hogben (2009)

[91] Article 29 Data Protection Working Party & Working Party on Police and Justice (2009)

[92] Schwartz (2009)

[93] Couillard (2009)

[94] Mills (2009)

practical realities for cloud service providers and users at the SecureCloud 2010 conference and at the study's expert workshop in September 2010 clearly confirmed this analysis: certain aspects of European data protection law (notably the rigid controller-processor model and the reliance on geographic location of data processing as an important factor in determining applicable rules) lead to substantial difficulties in practice.

However, it was also noted that there appeared to be a substantial degree of poor compliance with data protection laws, especially with respect to transfers to third countries (as many cloud service providers are established outside the EU) or data subject rights. For both of these topics, consistent and clear responses with respect to compliance with EU data protection laws were frequently unavailable.

The more restrictive regulatory framework with respect to sensitive data (mainly health and financial data) was often said to impede migrations, due to a requirement to ensure the segregation of data from other parties or to only process it locally. Whether or not this constitutes a barrier depends also on the transparency of data processing practices (including the location and permitted processing of the data), since the user of the cloud service must typically be able to ensure that such practices are compliant with applicable rules.

Data protection thus remains a sensitive issue, and will be a key component in the analysis of use cases in Chapter 7 of this study.

5.2.2 Can the cloud offer the same protection of intellectual property rights and provision of confidentiality and data portability?

The protection of intellectual property is often seen as a crucial building block for an information-society-based economy,[95] and in that respect the use of cloud computing services offers specific challenges as well. Especially in relation to the protection of copyrights, protected or valuable know-how and confidential information, users may be reluctant to entrust their data to service providers without specific and clear guarantees as to their continued and exclusive ownership and/or control of this information, as well as in relation to its confidentiality.[96] This is especially important in collaborative development services, where multiple users may work concurrently on a single product, giving rise to the question of ownership of this collective effort.[97]

The EU has implemented a significant regulatory framework in recent years to strengthen intellectual property rights, including the Software and Database Directives, the IP Rights Enforcement Directive, and several parts of the aforementioned eCommerce Directive, but it remains to be seen whether these are capable of addressing all of the questions presented by cloud computing environments, and – perhaps more importantly – whether these regulations (by focusing more on paradigm of content consumption rather than content re-use[98]) will not have an unintended stifling effect on the development of innovative service models. It is also worth noting that these regulatory measures primarily address

[95] Werbach (2008)

[96] Murley (2009)

[97] Lee (2008)

[98] Gervais, (2009)

intellectual property rights, with rules relating to know-how, trade secrets or confidentiality still being determined largely at the national level.

Another important concern is data portability: when information is injected in the cloud by an end user, can he or she be certain that it will later on be possible to migrate to a separate system (cloud or non-cloud) and remove all traces in the prior system?[99] If a cloud model offers no guarantee of this, then the possibility of left over remnants of data are an inherent design weakness that may make it legally impossible to use the service, especially for legally protected information (medical records, credit history, etc.).[100]

Based on the interactions with cloud stakeholders in the course of this study, the validity of these concerns (especially with respect to data segregation and lock-in) was confirmed. However, it is also worth noting that these were generally seen as technological and business challenges rather than as legal or policy problems. Technological improvements in security and interoperability appear to be seen as the most effective mechanisms for intellectual property protection, confidentiality and data portability in the cloud.

5.2.3 Pragmatically addressing the applicability of electronic communications regulations in the cloud

A crucial aspect of cloud services is the connectivity between all of the stakeholders, ensuring that the services are continuously available and that they can adapt in accordance with the needs of the users. However, this could also mean that certain regulations in relation to electronic communications – such as the EU Telecoms Package, recently revised – becomes applicable to certain parts of the cloud. This may lead to some cloud service providers being required to observe the specific data protection rules of the Directive on Privacy and Electronic Communications, or to log the activities of the users in accordance with the Data Retention Directive. Careful evaluation of how this can be done in a pragmatic manner will be needed to ensure that European cloud service providers do not suffer undue competitive disadvantages.

5.2.4 Meeting security obligations and responding to cybercrime

In order for end users to trust cloud services, they must be secure, which implies robustness, reliability and availability.[101] Cloud service providers will need to offer the required guarantees in this regard, by protecting their services against internal threats and against external attacks.[102] From a legal perspective, measures will need to be taken in order to meet the security requirements imposed by the Data Protection Directive (when applicable), including by protecting personal data against accidental loss or theft by third parties.[103]

---

[99] ENISA (2009a)

[100] Picker (2008)

[101] Cloud Security Alliance (2009)

[102] Murley (2009)

[103] Article 29 Data Protection Working Party & Working Party on Police and Justice (2009)

External threats will require cloud service providers to be able to collaborate efficiently with law enforcement bodies,[104] which may be difficult in a global environment where such bodies may be unknown. Effective international collaboration between these bodies is also needed, and while some regulatory harmonisation has taken place in recent years, including via the European Framework Decision on Attacks against Information Systems and the Council of Europe Convention on Cybercrime, it remains to be seen whether these will prove to be effective, especially in a cloud context where applicable law may be hard to identify.[105]

Indeed, cloud stakeholders at the SecureCloud 2010 conference repeatedly noted the difficulty in cooperating with national forensic teams in investigations related to cloud computing services. This is strongly linked to the issue of data protection: for cloud service providers, it is difficult to determine appropriate procedures for cooperating with law enforcement bodies, especially in an international context where the capacity and competence of claimed offices may be hard to validate.

5.2.5 Achieving clear and effective accountability and liability in the cloud

Cloud services can only achieve significant uptake in a commercial environment if end users have sufficient means at their disposal to seek and obtain redress in case of incidents; this is a prerequisite for ensuring trust. This means that it will need to be possible for an end user to identify a specific service provider, and for that service provider to be held accountable for its actions.[106] In a cloud context this can be complicated, due to the fact that it is not always transparent to the end user who the responsible party is in the case of specific incidents. Conversely, when an end user oversteps his legal boundaries, it will be necessary to identify him or her as well, even though this may be hard to do in a cloud environment.[107]

This specific element could be addressed (at least at a European level) through a clear and effective application of the eCommerce Directive's transparency and liability rules, insofar as these are applicable to the cloud service provider. However, the provider might also be confronted with a transparency challenge, as its own service could well be built on the basis of infrastructure providers that may be unknown to the service provider or which prove to be insolvent.[108] Care should be taken from a legal perspective that the responsibility and liability model behind a cloud service is not a house of cards, in which blame can be passed on and accountability is never achieved.[109] Such future possibilities were described in a presentation by ENISA.[110]

---

[104] Scolnik (2009)

[105] ENISA (2009a)

[106] Cloud Security Allliance (2009)

[107] Gervais (2009)

[108] ENISA (2009a)

[109] Bradshaw, et al. (2010)

[110] Hogben (2009)

The existence of this accountability issue is not different from traditional outsourcing, where it is also crucial to obtain sufficient liability guarantees from a service provider (without being confronted with endless disclaimers and waivers). However, transparency is again a substantial complicating factor with cloud computing: due to the greater complexity of services, it can also become significantly harder to define responsibilities and liabilities in case of incidents.

5.2.6 Divergent jurisdictional norms on harmful and illegal content and censorship in the cloud may affect availability

As data will be shared across multiple jurisdictions within the cloud, it is also possible that certain regimes may attempt to filter or suppress certain kinds of information, based on nationally applicable rules.[111] There are also deeper questions about the extent to which any such action is consistent with the values of the cloud customer. This is a potential legal threat to cloud computing services, as the reliability of the cloud – including the uninterrupted availability of data and services – is a precondition for its uptake. The responsibilities and liabilities of the cloud service providers will therefore need to be carefully examined, in particular to make sure that information cannot be removed for trivial reasons.

In a European context, the main influence in this respect will be the liability provisions of the eCommerce Directive – similar to those under the US Digital Millennium Copyright Act (DMCA) which have carved out liability limitations for certain information society service providers, most notably for hosting, caching and mere conduit services.[112] However, it will need to be evaluated whether this regime is effective in an international context, whether cloud service providers benefit from these provisions, and whether the notice-and-takedown approach embraced by this regulation does not stifle innovative cloud services. In practical terms, this study has found no substantial indication of this being considered as a crucial impeding factor for cloud computing by stakeholders.

5.2.7 Meeting consumer protection requirements for 'end user' cloud services

Cloud services will not only be used by professionals but also by everyday consumers, via their mail services or through social networking, for example. This makes it especially important that the precise scope and terms of the service are made clear and accessible, and that the rights of the consumer are fully respected. From a European perspective, it will need to be evaluated whether consumer protection rules such as the Unfair Commercial Practices Directive are applicable to cloud computing services and if they are likely to be effective in protecting consumers against abuses. Indeed, as noted in the 2010 Digital Agenda for Europe:

> *People's enjoyment of digital technologies, be it as citizens, consumers or workers, is marred by privacy and security concerns, by insufficient internet access, insufficient usability, by lack of relevant skills or by lack of accessibility for all.*

[111] Werbach (2008)

[112] ENISA (2009b)

> *…Consumers and businesses are still faced with considerable uncertainty about their rights and legal protection when doing business on line.*[113]

So far, there is no strong indication that consumers show a great deal of additional concern with respect to privacy and security in the cloud (as compared to non-cloud services); however, this perception may be due to a lack of understanding (ie, consumers are unaware of the differences between cloud and non-cloud services) or simply due to the difficulty consumers have making their voice heard (ie, consumers are concerned about cloud services, but lack clear representation in policy discussions). Either way, it would be beneficial from a policy perspective to ensure that consumer awareness is increased (with respect to the characteristics of cloud services and their rights), and that the technological and practical means to enforce their rights are improved (for example, with easy tools to exercise their privacy rights as data subjects, or easier access to ombudsmen/consumer organisations which can mediate in disputes). The Digital Agenda already proposes several actions on these points which may prove to be beneficial.

5.3 The will of the parties: contractual provisions

5.3.1 Matching the scope and quality of cloud services (including SLAs) to the requirements of the user

The precise scope and the quality of the services must be agreed between the service provider and the customer. For advanced or business-critical cloud computing services, service level agreements (SLAs) are often concluded as well, in which certain commitments can be determined (for example, in relation to service availability, response times, speed, data integrity, etc.).[114] The conclusion of such agreements can be complicated, due to the differing needs of specific users. For this reason, services are often organised into different tiers, so that customers can choose well-defined service packages. Care must be taken that these are sufficiently comprehensive, accessible and balanced, to ensure that the intentions of both parties are accurately represented.[115]

5.3.2 Is the current legal framework robust enough to stand up to the flexibility inherent in cloud computing solutions?

One of the main drivers behind cloud computing is the possibility of easily extending or reducing the resources used by the customer to match his or her needs, by acquiring additional storage or computing power on the fly, for example.[116] This flexibility must be reflected by the legal framework as well: certain parameters of the service may need to change, without, however, invalidating the core of the agreement (Simmendinger has characterised these issues in the context of the grid[117]). This is a complex exercise since it

---

[113] European Commission (2010)

[114] Armbrust et al. (2009)

[115] ENISA (2009a)

[116] Murley (2009)

[117] Christian Simmendinger (2008), Whitepaper Support for Client-Server based License Management Schemes in the Grid, http://www.it-tude.com/fileadmin/gridipedia/lm_cluster/LM_WHITEPAPER_V0.1-1.pdf

requires a distinction to be made between the crucial unchangeable elements of the contract and the modalities that need to evolve as usage of the service changes.

5.3.3 Validity and consent: establishing common boundaries for negotiation

Flexibility also has an impact on the management of the contract, since many (though not all) cloud services need to foresee the ability to easily change the applicable terms – this impact was described in a University of Berkeley paper.[118] Unilateral modifications of crucial elements are not acceptable, certainly if they would be applied automatically. Therefore, the parties need to agree beforehand under which conditions the terms of a contract may change.[119] In highly advanced models, this can include negotiation procedures (such as described by Simmendinger[120]), where both parties have established automated systems to communicate their demands and desires, and that can maintain and update the agreement between the parties as the needs dictate.[121] However, such systems are still rare, given the conceptual difficulty and legal challenges of exercising one's negotiating powers through an automated process.

In practical terms, the main challenge is the difference in negotiating power between larger cloud service providers and smaller cloud users. When this is the case, the standard terms and conditions will apply, which often result in limited guarantees for cloud users. Clearly, this issue is not unique to the cloud, and occurs to some extent in every business agreement. However, it should be noted that consumers also use cloud services, and their negotiating power will frequently amount to very little.

5.4 Overcoming legal barriers: key tools

While the overview above has made it clear that there are a number of legal issues to be addressed, it is important to keep in mind that there are many instruments available to do so. Direct regulatory intervention may not prove to be the only or the most effective way to seek an appropriate answer.

5.4.1 Role of 'hard law': EU, national and international laws and treaties

The overview above has indicated that applicable regulatory texts exist for at least some of the key aspects to be addressed. Doubtless they will prove to be useful as general frameworks to handle some of the issues, but their restrictions must also be recognised.[122] In particular, it is clear that their applicability will generally be linked to a specific regional context, and that a uniquely European approach may not prove to be a very attractive option in a domain such as cloud computing, which owes part of its appeal to its international business potential. In addition, the risk of creating competitive disadvantages to service providers established in more demanding regulatory regimes is very real. Therefore, alternative approaches must also be considered.

---

[118] Armbrust, et al. (2009)

[119] D'Andria et al. (2006)

[120] Simmendinger (2008)

[121] Hasselmeyer, et al. (2007)

[122] Werbach (2008)

5.4.2 Soft law: guidelines, good practices and self- and co-regulatory approaches

As cloud computing is still a relatively young domain, not many examples of self- and co-regulatory approaches are likely to have emerged yet. Nonetheless, in the case studies sections below will examine if any initiatives have been adopted by leading cloud computing services that may help to address some of the legal issues identified above.[123]

5.4.3 Contractual frameworks and terms and conditions (T&Cs)

In most business relationships contractual agreements (including in the form of standardised T&Cs) play a dominant role in guiding such issues as conflicts of law, choice of forum, and outlining rights, responsibilities and guarantees. One of the crucial questions will be to what extent contractual frameworks are already being used to eliminate legal ambiguities, and to what extent these approaches can be promoted further to offer a more comprehensive solution to all of the legal issues discussed above. This will be examined through the use case studies in Chapter 7 below.

5.5 Concluding remarks

This chapter has provided an overview of how the literature approaches the main regulatory and legal challenges to assure the security, privacy and trust of cloud computing. As already argued, technology can only provide a partial response. The development of legal responses is not a trivial exercise and requires a comprehensive public policy approach where the interests of all stakeholders are taken into consideration.

Table 2, below, maps the challenges associated with the legal characteristics of cloud computing to the domains of security, privacy and trust. As noted in the previous chapter, this mapping represents an overview of those challenges arising from the literature studied for this report – where there are gaps it does not indicate that no issue yet exists but rather that the issue was not evident in the literature.

Table 2. Mapping of key challenges in respect of the legal and regulatory aspects of cloud computing arising from the studied literature

Area / Legal Domain	Security	Privacy	Trust
Applicable law (data/service location)		Existence and effectiveness of privacy protection laws/principles	Existence of a clear legal framework as a basis for the service
Dispute resolution		Accountability: can disputes in the cloud be resolved?	Accountability: is the cloud stable enough to inspire trust
Data protection and privacy	Obligation to implement secure data processing approaches	Compliance with privacy principles	Confidence in data protection practices
Protection of intellectual property rights	Confidentiality and availability (data portability)		Confidence in the security/confidentiality of data entrusted to the cloud
Electronic		Safeguarding	Protection against

[123] Cloud Security Alliance (2009)

communication in the cloud		communications secrecy	eavesdropping (surveillance by public and private parties)
Security obligations and cybercrime	Confidentiality, availability and integrity; effective law enforcement	Safeguards against unlawful intrusions in the personal sphere	Balancing privacy safeguards with the need for security
Accountability and liability	Accountability for security breaches and incidents	Accountability for data leaks: can incidents be identified and sanctioned?	Trust that instruments for restitution and sanction will work
Harmful and illegal content	Availability: can the cloud identify and respond to such content?		Trust in jurisdictions to apply transparent standard or approach to illegal content (in line with cloud user expectations and applicable laws)
Consumer protection	Obligation to implement secure data processing approaches	Data subject participation; restitution	Consumer protection rules must be effective (ie, applied and enforced in practice)
Scope and quality of services (SLAs)	Transparency and security metrics are needed to ensure that security targets are met		Assurance and commitments between two parties
Flexibility	Availability and scalability		Resources must be available when needed, in accordance with agreement between parties
Validity and consent	Transparency and accountability	Transparency must be ensured. Consent from consumers must be free, specific and informed	Assurance and commitments between two parties must be clear and enforceable

As can be seen, the watchwords from this mapping exercise in the legal domain are transparency, availability and accountability. Transparency is an important challenge to meet security, privacy or trust obligations, since this supports the contractual will of the parties as well as helping to fulfil one of the globally accepted privacy principles. Availability arises as an issue since SLAs and other instruments are key in providing for a governance framework or assurance in respect of availability requirements. Finally, accountability is an important factor arising from one of the main legal challenges of cloud computing: namely that commitments from two parties must be clear and enforceable in practice. This in turn stimulates trust across both sides of the cloud user/provider relationship.

The assessment of the legal viability of migrating to a cloud computing service was by a wide margin the most frequently quoted legal challenge from the research conducted with stakeholders for this study. As a part of such assessments, the compliance of cloud projects with applicable laws and contractual obligations needs to be tested, and this frequently resulted in difficulties. Data protection rules were a common example of this issue, due to the difficulty of determining applicable law(s) and remaining fully compliant. Other laws

and rules however also played a role, especially for more sensitive data such as health records, financial information, etc. In those cases, problems were not only presented by data protection laws themselves (which include stricter rules for sensitive data as per Article 8 of the Data Protection Directive), but also by more stringent sector-specific rules (eg, patient's rights, administrative regulations, etc.) which were seen as difficult to reconcile with flexible cloud models.

Given that a legal/contractual framework has to be sufficiently stable and comprehensive to support the trustworthiness of a legal relationship, how can these requirements be reconciled with the inherent flexibility of a cloud computing model? In practical terms, this can only be done by ensuring that the rights, responsibilities and liabilities of each party are sufficiently outlined, and that the transparency of the legal relationship between service provider and end user is adequately ensured. If these conditions can be met, then compliance and accountability are facilitated, a crucial enabler for trust, privacy and security in the cloud.

CHAPTER 6 Putting it all together: key risks and operational challenges

6.1 Summary of legal and technical issues

The two preceding chapters provided a summary overview of the main legal and technical issues that can present challenges to the deployment and use of cloud services, and also examined what the impact of these issues was on cloud services in practice.

Table 3, below, combines the tables from the previous chapters and provides an overview of many of the key legal and technical concerns found in the literature reviewed, and maps them against the resulting security, privacy and trust implications. As can be seen, some issues such as availability, accountability and integrity appear relatively frequently. Unsurprisingly, each specific identified issue has trust implications: any security or privacy risk in the cloud by definition negatively impacts the trustworthiness of cloud services.

Table 3. Mapping of the legal and technical issues to their respective security privacy and trust implications

Area / Technological or legal domain	Security	Privacy	Trust
Virtualisation	Integrity	Segregation of personal data on shared infrastructure	Compromised virtual machines/ hypervisors permit loss of trust
Grid technology	Availability		Interoperability
Web services	Integrity and confidentiality	Security and confidentiality	Interoperability (in the context of identity and access management)
Service-orientated architectures	Integrity		The reliance of distributed systems on different security credentials
Web application frameworks	Integrity and availability		Trust across distributed environments
Encryption in the cloud context	Confidentiality	Security and confidentiality	
Applicable law (data/service location)		Existence and effectiveness of privacy protection laws/principles	Existence of a clear legal framework as a basis for the service

Dispute resolution		Accountability: can disputes in the cloud be resolved?	Accountability: is the cloud stable enough to inspire trust
Data protection and privacy	Obligation to implement secure data processing approaches	Compliance with privacy principles	Confidence in data protection practices
Protection of Intellectual property rights	Confidentiality and availability (data portability)		Confidence in the security/confidentiality of data entrusted to the cloud
Electronic communication in the cloud		Safeguarding communications secrecy	Protection against eavesdropping (surveillance by public and private parties)
Security obligations and cybercrime	Confidentiality, availability and integrity; effective law enforcement	Safeguards against unlawful intrusions in the personal sphere	Balancing privacy safeguards with the need for security
Accountability and liability	Accountability for security breaches and incidents	Accountability for data leaks: can incidents be identified and sanctioned?	Trust that instruments for restitution and sanction will work
Harmful and illegal content	Availability: can the cloud identify and respond to such content?		Trust in jurisdictions to apply transparent standard or approach to illegal content (in line with cloud user expectations and applicable laws)
Consumer protection	Obligation to implement secure data processing approaches	Data subject participation; restitution	Consumer protection rules must be effective (ie, applied and enforced in practice)
Scope and quality of services (SLAs)	Transparency and security metrics are needed to ensure that security targets are met		Assurance and commitments between two parties
Flexibility	Availability and scalability		Resources must be available when needed, in accordance with agreement between parties
Validity and consent	Transparency and accountability	Transparency must be ensured. Consent from consumers must be free, specific and informed	Assurance and commitments between two parties must be clear and enforceable

This table is intended as a short checklist of key areas to be covered in cloud deployments, and will be used in the next chapter of the study to support the analysis of case studies, and specifically to determine to what extent each case study has been able to find appropriate answers to the legal and technical challenges in their respective domains. It should be noted, however, that this list is mostly valid with respect to *public* cloud systems (where specific data or services are outsourced to a third party service provider), but much less so when using *private* cloud systems (which are deployed, operated and controlled by the

user). In the latter case, the use of cloud technologies may cause technical and security challenges, but to a lesser extent since they can be more effectively managed internally. It would also appear that with private cloud deployments there are a limited number of new legal challenges.

The SecureCloud 2010 conference and the expert workshop held in Brussels towards the end of this study confirmed the validity of this analysis with respect to the legal and technical issues behind the use of cloud computing services. However, an additional topic that was discussed extensively was the operational challenges and risks in migrating to cloud systems, at least when these systems are controlled by third parties. This warrants some further elaboration.

6.2 Migrating to the cloud: the operational challenges

Many of the legal and technical challenges discussed above are not entirely unique to cloud systems, and can also be found with traditional outsourcing solutions. However, from an operational perspective, cloud models are more fundamentally different from traditional outsourcing.

One of the key observations voiced during the SecureCloud 2010 conference and in conversations with cloud stakeholders was the importance of control when moving to cloud-based systems controlled by third parties. More so than technical or legal barriers, the loss of control when entrusting data or services to a cloud service provider was identified as an impediment. This is especially true in terms of transparency or visibility: in a cloud model, it may not be clear to the user precisely where data or services are hosted, or whose control they are under. Similarly, the inherent greater technological complexity of cloud solutions (due to the need for flexibility and scalability) make it harder for the user to determine the objective security of cloud systems.

This issue also exists for traditional outsourcing models, but only to a much smaller extent: the user knows that data/services are hosted on a specific system at a specific location operated by a specific service provider. Even if in reality the traditional approach still permits certain flexibilities and complexities (due to subcontracting, offshoring, co-location, etc.), the customer knows who he can address and where audits and checks can be conducted. In cloud models, these issues can be much more complicated because of the cloud's inherent nebulous and fluid nature: while a customer no longer has to concern himself with exact implementation details (benefit), he also loses part of his ability to assess (cost). As a result, new operational challenges arise.

There are international standards, such as ISO27001/2[124] and from ISF[125] and NIST,[126] which provide guidance on best practice in terms of the operational controls which should

[124] ISO27001: Information Technology – Security Techniques – Information Security Management ISO27001Systems – Requirements; ISO27002: Information Technology – Security Techniques – Code of Practice for Information Security Management

[125] ISF (2005)

be used to mitigate information risks. In general the security community is seeking to adopt these controls within a cloud environment (see for example the activities of ENISA and in particular the ENISA Cloud Computing Information Assurance Framework[127]). The Cloud Security Alliance have also launched a certification programme for professionals wishing to demonstrate their cloud security practice knowledge. However, there remain concerns that not all existing risk controls will easily scale into a mature cloud ecosystem. Controls which may be at risk include the following:[128]

- Physical access controls – how can the cloud user achieve requirements for physical access control given the cloud service provider establishes and controls the when who, why and how of physical access measures?
- Application development and maintenance – is it possible to assure the development and maintenance of applications in a cloud environment when external parties such as the cloud service provider or other third parties are responsible?
- Vulnerability management – assigning responsibility for patch management and the deployment of software and hardware updates between the cloud service provider and cloud user is especially complex given virtualisation and the dynamic reconfiguration of software and infrastructures.
- Monitoring – how to establish effective, timely and accurate monitoring of levels of security and privacy in business-critical infrastructure when those responsible for the infrastructure may not be prepared to share such information under standard service level agreements.
- Identification and authentication – the integration and control of identity and access management infrastructures in a cloud environment where the cloud service provider might have different approaches and tolerance for risks to identity infrastructure, in addition to the complexities of providing for identity across distributed cloud environments.
- Access control – how can the cloud user govern access control risks when the levels and types of access control to key ICT assets deployed by the cloud service provider are not known.
- Encryption – how can the cloud user manage encryption and key infrastructures and assign responsibility across the boundary between their own organisation and the cloud service provider?
- Continuity and incident management – how can the cloud user determine appropriate thresholds and criteria for responding to incidents (eg, agreeing on what constitutes an incident) and policies and processes for responding and achieving assurance of the evidential chain.

---

[126] NIST (2005)

[127] ENISA (2009b)

[128] For a detailed discussion, see Creese et al. (forthcoming) Inadequacies of Current Risk Controls for the Cloud, to appear in the proceedings of CloudCom2010 November 2010.

- People (security team) – can the cloud service user understand and manage the quality and mechanisms for professionalism in the cloud service provider's security team and whether they match with that of the cloud user.
- People (business function) – how incorporating a high culture of security in the organisation is possible when those responsible for strategic decisions regarding the trade off between security/business requirements may be external to the cloud user organisation (and thus may operate under different motivations).
- Security control testing – how can the cloud service user test the validity and effectiveness of security controls when such tests may not be agreed by the cloud service provider or may impact upon other customers of the cloud service provider.
- Accredited components – how to accredit components in the cloud to the required standard imposed upon by accreditation processes.
- Data remanence – can the cloud service user achieve the appropriate degree of assurance that data have been securely deleted by the cloud service provider and no data shards remain in the cloud environment.
- Asset management – is it possible to gather and exploit data on the entirety of the cloud service provider's physical assets to complete an asset register in accordance with good practice guidance?

Some of these controls will require the development of new process, while others may require innovative technologies – for example, it is not clear how a cloud might be monitored for intrusions, where sensors should be placed (such as between virtual machines residing on a single piece of hardware, at the perimeter of a cloud service providers network, or between clouds within a supply chain), to what degree it will be desirable to monitor threats propagating across a cloud ecosystem (which would involve interoperability and information sharing on a level not currently practiced), or whether we need to concern ourselves with attacks emanating from a cloud. This in turn could have legal issues associated with liability and to what degree a cloud service provider might be held accountable for malicious activities hosted on it.

Depending on the cloud service the following operational challenges may be of relevance:

- Data or services may be (or become) hosted from another country, possibly without the end user's knowledge. As noted earlier, this may also cause legal challenges, such as a change in applicable law or competent courts, or greater complexity in ensuring compliance with data protection regulations (for example, new notifications may be required by data protection authorities, and auditing of the service provider's infrastructure to ensure appropriate security may become very problematic).
- Infrastructure may be shared with other customers, leading to data segregation concerns. Similarly, can confidentiality be appropriately ensured?; is intellectual property adequately shielded against infringements?; is personal data properly protected against accidental leaks?; and so on. Again, these issues will trigger the aforementioned legal questions.

- Incidents may cause service interruptions without it being evident where the problem lies, and thus how it may be addressed. From a legal perspective, this will trigger questions of compliance with SLAs or other quality standards integrated in the contractual framework: is the service provider responsible, and can he be held accountable?
- Data withdrawal might be difficult, in the sense that it can be hard for a cloud user to determine whether deleted data has actually been removed from the provider's systems, or whether it has merely been made inaccessible. From a data protection perspective, the distinction is crucial, given that European data protection law requires personal data to be deleted when it is no longer necessary for the purposes of the data processing. Merely allowing it to linger on an abandoned system is not acceptable from a data protection perspective.
- Auditing and investigations may be more challenging, due to the complexity of the system (for example, the use of virtualisation technologies may make it harder to determine where data is located, and which systems may be audited without accessing another cloud user's data). Again, this was quoted repeatedly as a challenge with respect to data protection by stakeholders in the course of the study: users of cloud service providers are expected to ensure security and reliability of the service provider, in some cases by performing local audits of the data processing infrastructure. From a practical perspective, this is virtually impossible with cloud service providers, whose infrastructure can be spread geographically almost without limitation.

As noted above, none of these points is strictly unique to cloud services: traditional outsourcing models can also be confronted with these challenges. However, the fact that cloud models combine all of these elements and add a layer of complexity means that they are considered to be more problematic.

A key recurring operational challenge is increasing transparency: end users must be able to get a clear overview and guarantees of where and how their data and services will be handled, and how security risks and incidents can be detected and addressed. This need for transparency was a consistent thread running throughout the SecureCloud 2010 conference and expert workshop. Taking this into account, it is perhaps not surprising that much of the ongoing work with respect to cloud computing from an operational perspective focuses on developing methodologies to improve transparency and to obtain usable metrics for assessing security. Examples of these approaches include:

- The Cloud Security Alliance 'Security Guidance for Critical Areas of Focus in Cloud Computing', which identifies specific cloud security risks in a variety of areas (notably architecture, governance risks and operational risks) and provides recommendations for managing these.[129]

[129] Cloud Security Alliance (2009)

- The Jericho Forum Self-Assessment Scheme, which defines crucial rules in the form of 'commandments' which are mapped to acceptable/best practices that cloud users can use to assess plans.[130]
- The Shared Assessments questionnaires for vendor/service provider assessments, which allow the systematic evaluation of security, privacy and business continuity controls. It should be noted that these questionnaires are not tailored specifically to cloud computing, but are rather built around a series of general questions that any vendor/service provider should be able to provide an acceptable response to.[131]
- The ENISA Cloud Computing Risk Assessment Report, which has the additional merit of being drafted from a European perspective and therefore being geared specifically towards the legal and business context in which EU cloud stakeholders operate. The report describes the primary information security benefits and risks of cloud computing, again coupled with a set of practical recommendations.[132]

These frameworks are useful in assisting aspiring cloud users in the identification of appropriate questions to raise with their cloud service providers, and in obtaining the necessary guarantees and assurances. Especially for SMEs (who are less likely to have sufficient expertise to identify the correct questions) such frameworks can be very useful as an aid to decision making. In practice, these frameworks can operate as a form of de facto soft law, assisting cloud service users in ensuring their compliance with legal obligations and good security practices. Strictly speaking however, none of the above has a formal or official status (in the sense that compliance with these frameworks provides any assurance of compliance with legal obligations), so that their role is to enable or facilitate operational auditing.

Presentations during the SecureCloud 2010 conference indicated that cloud computing services were currently mainly used for non-critical purposes (excluding highly sensitive data and services whose continuity may not be interrupted at any point) or only internally (without entrusting data or services to third parties). This appears to be to a large extent a consequence of the perception of lesser control in cloud models: since customers lose the ability to control implementation and operational details, they are less willing to move vital data and services into the cloud for fear of being unable to handle significant technical, operational or legal challenges. This situation is likely to change in the future as cloud maturity improves, and users feel more at ease entrusting important information and services to external cloud service providers (at least up to the same point as is presently the case for non-cloud outsourcing). The case studies in the next chapter of the report will examine this observation in greater detail.

---

[130] See (as of 22 November 2010) http://www.opengroup.org/jericho/

[131] See (as of 22 November 2010) http://www.sharedassessments.org/value/

[132] ENISA (2009a)

CHAPTER 7 Case studies

7.1 Introduction

As has been illustrated in the previous chapters of this report there are a range of concerns, challenges and uncertainties relating to the security, privacy and trust issues associated with the use of cloud computing. These may be inherent in the underlying technical components or drivers of cloud computing, such as virtualisation and web services, or the legal or regulatory frameworks that surround cloud computing, such as the identification of applicable law and data protection/privacy compliance.

In any respect, the characteristics of cloud computing (shared to a certain extent with outsourcing) relating to security or privacy may be seen as both a challenge and an opportunity. For example, the degree of granularity regarding the locations of stored data in the cloud may be linked to the negotiating power of the cloud customer and may in certain cases be in conflict with the prevailing regulatory framework (for European cloud customers, the legal requirements stemming from privacy and data protection law). Equally, the wealth of data created by the logging functionality inherent in the cloud computing model presents something of a security opportunity in terms of understanding exactly what is going on in an ICT infrastructure, especially for SMEs or organisations that might not otherwise have the know-how or resources to implement such functionalities on their own. Furthermore the possibilities of quickly and seamlessly re-provisioning infrastructure also represent an opportunity in respect of availability and resilience.

7.2 Initial classification of case studies

This chapter takes a specific look at real life case studies,[133] to gain insight into how the issues discussed above are handled in practice. The data collection for these case studies was conducted between May and August 2010.

Given the role of use cases in the study, it is important to select representative examples which provide a realistic overview of cloud computing. In that respect, several classifications may be proposed.

A first important distinction lies with the control of the cloud service:

[133] The term 'case studies' is used to refer to actual instances of implementation, as distinct from 'use cases', which are generic classifications of types of implementation (eg, 'a public sector healthcare concern')

- **Private clouds** (internal clouds): the user controls every aspect of a cloud (eg, a company creates its own internal cloud service for internal use only); in other words, while cloud technologies are used, no external cloud service provider is involved, eliminating many of the security concerns defined above.
- **Public** clouds (external clouds): the cloud itself is made available to external users as a service.

The motivation behind the operation of the cloud service can be used as a second key characteristic:

- **Commercial cloud computing services**: operated and offered as for-profit ventures.
- **Public sector cloud computing services**: operated and offered for the public good; typically, in these cases a cloud service is offered by a public body and offered only to other public bodies or other entities which somehow contribute to a public service.
- **Scientific/academic/research cloud computing services**: operated and offered to advance scientific or technical progress in various disciplines eg, by pooling/sharing resources.
- **Community-driven cloud computing services**: where the cloud serves to establish, promote or interconnect more- or less-defined communities of users.

A third classification could be based on the degree of sensitivity of the usage of the cloud service, or in other words the risk it presents to the user. While much more nebulous and subjective than the distinctions above, it nonetheless presents a crucial characteristic when assessing the viability of a cloud project. One could distinguish:

- **Low risk** use cases: where data loss or service interruption would have no serious impact on the user's activities. No crucial data or services are entrusted to the cloud, and the risk of significant damage is limited. An example would be using the cloud for redundant backup of non-vital data, or for added value services which would only cause minor discomfort in case of interruption.
- **Medium risk** use cases: where data loss or service interruption would have a noticeable impact on the user's activities resulting in real damages but without being business critical. Examples include email hosting in the cloud, or relying on cloud-based human resource management systems (for businesses which do not rely entirely on email communications or on HR management for their business activities).
- **High risk** use cases: where data loss or service interruption would have a detrimental impact on the user's activities, and could result in significant and broader harm. Examples might include entrusting sensitive (eg, financial or health data) or classified national security data to the cloud.

Finally, the type of service being provided can be considered a deciding factor, with IaaS, PaaS, SaaS and HaaS being the main possibilities.

Within the scope of the present project, it would not be feasible to achieve full coverage of each of these archetypes. Rather, it seems desirable to make an optimal selection, based on the main observed uses of cloud computing. This issue will be briefly examined below.

7.3 Observations on case studies in practice

Both the literature and presentations from the SecureCloud 2010 conference illustrated several trends that may be relevant for the selection of use cases to be examined:

- Firstly, both private and public cloud initiatives were presented during this conference. For the purposes of this study, public initiatives appear to be more relevant, since security and privacy risks are much easier to identify, address and avoid in internal projects.
- The main emphasis was on commercial and public sector clouds; thus, both of these categories would need to be represented in the selected use cases. Scientific clouds and community-driven clouds would be interesting as well (especially since community-driven clouds may be more prone to privacy risks), but appear to be less common today.
- High-risk use cases are few and far between; thus it seems inevitable that most use cases fall into the low and medium risk categories.
- Finally, with respect to cloud service types, IaaS and HaaS cloud services seem less interesting for the purposes of this study, since although they may share the concerns of multi-tenancy and dynamic migration they are less fundamentally different from traditional outsourcing models than PaaS and SaaS models.

7.4 Matrix of case study typology

Table 4, below, indicates the key characteristics of the selected case studies:

Table 4. Characteristics of case studies

Case study	Cloud service model used by case study subject	Service model offered to end users	Risk	Rationale
Danish National IT Agency	IaaS	Traditional (non-cloud)	Low to Medium	This case study represents a social network and knowledge sharing platform. However, there are concerns regarding the use of personal data collected as part of this implementation
City of Los Angeles	SaaS	SaaS	Medium	Although this case study exists in a different regulatory environment to that of Europe, there are still overarching security

				and sensitivity lessons associated with the storage of criminal record data in the cloud. The LAPD was a major user community in this implementation
EU Member State eHealth cloud	IaaS	SaaS	Medium to High	This case study covers the use and processing of health related personal data in the cloud

7.5 Case Study 1: the Danish National IT and Telecom Agency

The use case owner is the Danish National IT and Telecom Agency (NITA),[134] part of the Ministry of Science, Technology and Innovation. It is a public service tasked with establishing the groundwork for the digitalisation of Denmark, and is based on five strategic objectives:

- Denmark should have an electronic communications infrastructure ranking among the best in the world.
- The Danes should obtain the necessary competencies to utilise this infrastructure optimally.
- The Danes should feel secure and confident when using the infrastructure.
- The Danes will use the infrastructure in practice because it offers valuable and useful content.
- NITA will contribute actively to the reduction of energy consumption and environmental hazards in Denmark.

As a part of its remit, NITA supports the development and deployment of IT architecture and standards in the Danish public sector. In a January 2010 policy paper on Denmark as a high-tech society, an independent committee appointed by the Ministry of Science recommended that:

- Denmark should become a leading cloud (and IT in general) nation. The government should invest heavily in cloud computing. Due to its size, the public sector would be an attractive customer for providers of cloud computing and of cloud-based services, and could help drive the resolution of challenges related to privacy, data security, expenses, etc.
- The public sector should undertake projects where cloud computing is used. The focus should be to gather experience and publish guidance on the use of cloud computing, among other things to ensure clear and attractive subscription terms and standard agreements, new paradigms for safety, etc.

[134] See (as of 22 November 2010) http://www/itst.dk/

NITA does not procure or provide IT solutions for other public sector entities; rather, its task is to support IT development in Denmark in general. The experiences of NITA are nevertheless also intended to aid private bodies.[135]

7.5.1 Description of the use case

The use case concerns the migration to the cloud of the social network Digitalisér.dk,[136] which is a community managed by NITA aiming to provide public sector institutions and suppliers (although anyone is free to join) with information about the digitalisation of the Danish public sector and to allow users to exchange information about the issue.

In other words, Digitalisér.dk is a social network, a tool for development and knowledge sharing, a forum for the digitalisation of Denmark, and both a formal central repository of information on data interchange standards and a big open digital playground. It is a creative space for everyone involved in digitising the public sector.[137]

As NITA points out in the English version of the site, 'Digitalisér.dk aims to stimulate development and adoption of digital content and business models by utilising Web 2.0 technologies and public data and digital resources. With digitaliser.dk, the Danish government has created a new model of partnership between the tech community and government which paves the way for more direct communication between the public sector, citizens, and businesses. Citizens and businesses are no longer passive recipients of public information but participate in dialogue and knowledge sharing with the public sector.'[138]

Digitalisér.dk was already operational before the migration to the cloud, using a different, more traditional technical hosting model. The migration started in November 2009 and at that time Digitalisér.dk had some 2,500 registered users. NITA decided to transfer the community to the cloud firstly in order to assess the benefits of the cloud (including in terms of innovation, efficiency, cost effectiveness and environmental impacts[139]) and secondly to reduce hosting costs. The primary goal was, however, to test the feasibility and impact of a cloud approach, and offer a practical example to other Danish IT initiatives (in the public or private sector). The expectation of cost reduction played only a secondary role.

NITA, in fact, does not own and manage any cloud infrastructure itself; this part of the platform is completely outsourced to a third party provider, Amazon. In other words Digitalisér.dk runs on a common 'off-the-shelf' cloud solution, rather than an ad hoc infrastructure.

---

[135] The information in this profile was kindly provided by M. Marc Andersen, NITA IT Security Analyst

[136] See (as of 22 November 2010) http://digitaliser.dk/

[137] For more information about Digitalisér.dk, see (both as of 22 November 2010) http://ourdata.eu/node/65 and http://www.epractice.eu/en/cases/digitaliserdk

[138] See (as of 22 November 2010) http://digitaliser.dk/resource/432461.

[139] See (as of 22 November 2010) http://www.itst.dk/filer/Publikationer/hoejhastighedskomiteen/hoejhastighedeskommiteen_del1/978-87-92572-06-9.pdf, p. 33 following.on (in Danish)

7.5.2 Chosen cloud solution

As pointed out above, Digitalisér.dk is hosted in the cloud infrastructure of Amazon, a commercial cloud provider with whom NITA has concluded a standard business agreement. The standard Amazon cloud provision contractual terms were applied (see below for further information); no unique agreements were negotiated. Due to the limited scale and risk of the application, this was seen as an acceptable solution.

The cloud resources are not shared with other Danish institutions. NITA does not operate on behalf of other governmental entities, but simply decided to migrate Digitalisér.dk to a cloud model to test the potential benefits of the cloud without the ambition to immediately apply the experience gained to other services or government departments. NITA does not provide services and/or consultancy in the field of cloud computing. Of course, nothing impedes other public and private bodies following the example of NITA and deciding, based on the results of the use case, to migrate to the cloud as well.

Regarding, more specifically, the use case, the decision to migrate Digitalisér.dk to the cloud is based more on the will to test the potential of cloud computing than on the necessity to get more scalability, as might be required for larger social networks. Taking into account, in fact, that more or less 2,500 registered users were using Digitalisér.dk at the time the migration started, efficient scalability, when needed, could also be reached in a different (traditional) hosting environment. The chosen cloud solution of course offers scalability opportunities, but the use case indicates that the increase in flexibility may make it convenient to adopt cloud computing even if scalability is not a top priority for the user.

7.5.3 Risks and challenges of cloud use

This use case can be qualified as low risk, where data loss or service interruption would have no serious impact on individuals or on society. No crucial data or services are entrusted to the cloud, and the risk of significant damage is limited.

In particular, as regards privacy issues, no sensitive data are involved, in the sense that Digitalisér.dk processes only the names and email addresses of the registered users and no other data are stored in the cloud. Given the nature and objectives of this specific social network (which focuses on professional communications rather than private liaisons) no sensitive private data are processed, and therefore no important security issues or risks of significant damages are likely to arise. Furthermore, the terms of use of Digitalisér.dk expressly prohibit users from posting messages with sensitive content or uploading sensitive data.

In any case, NITA made a risk assessment before the migration to the cloud, and no particular problems and concerns were identified at the time that would be a barrier to the migration. This risk assessment, in fact, was principally focused on technical aspects, since privacy and security concerns were not likely to play a major role after the migration. The same assessment methodologies were used as for any other type of technology project (in other words, they were general questions essential to any such project, not specifically created for cloud projects).

Technical and business challenges too have been assessed as being limited, in the sense that the migration to cloud computing was not expected to (and did not) impact to a notable

extent the operation of Digitalisér.dk (according to the key paradigm of 'business as usual' – see below for further remarks).

7.5.4 Chosen approach to address these risks

Since the risks were low, NITA decided that the contractual terms of use of the cloud services contained in the Amazon's Service Level Agreement (SLA) were acceptable (to be more precise, the contractual relations with Amazon are governed by the Amazon Web Services Customer Agreement and the Amazon EC2 Service Level Agreement).

Regarding privacy and security, the Amazon Web Services Customer Agreement states that:

> **11.6. Your Applications are Your Responsibility.** […] we specifically disclaim all liability, and you shall be solely responsible for the development, operation, and maintenance of your Application (including any Bundled Application) and for all materials that appear on or within your Application and you agree that you shall, without limitation, be solely responsible for:
>
> […]
>
> 11.6.4. ensuring that your Application accurately and adequately discloses, either through a privacy policy or otherwise, how you collect, use, store, and disclose data collected from visitors […];

> **7.2. Security.** We strive to keep Your Content secure, but cannot guarantee that we will be successful at doing so, given the nature of the Internet. Accordingly […] you acknowledge that you bear sole responsibility for adequate security, protection and backup of Your Content and Applications. We strongly encourage you, where available and appropriate, to (a) use encryption technology to protect Your Content from unauthorized access, (b) routinely archive Your Content, and (c) keep your Applications or any software that you use or run with our Services current with the latest security patches or updates. We will have no liability to you for any unauthorized access or use, corruption, deletion, destruction or loss of any of Your Content or Applications.
>
> **11.2. Applications and Content.** You represent and warrant: (i) that you are solely responsible for the development, operation, and maintenance of Your Content, including without limitation, the accuracy, security, appropriateness and completeness of Your Content and all product-related materials and descriptions; […].
>
> **11.5. Disclaimers.** […] WE AND OUR LICENSORS DO NOT WARRANT THAT THE SERVICE OFFERINGS WILL FUNCTION AS DESCRIBED, WILL BE UNINTERRUPTED OR ERROR FREE, OR FREE OF HARMFUL COMPONENTS, OR THAT THE DATA YOU STORE WITHIN THE SERVICE OFFERINGS WILL BE SECURE OR NOT OTHERWISE LOST OR DAMAGED.
>
> **11.6. Your Applications are Your Responsibility.** […]We specifically disclaim all liability, and you shall be solely responsible for the development, operation, and maintenance of your Application (including any Bundled Application) and for all materials that appear on or within your Application and you agree that you shall, without limitation, be solely responsible for:
> […]
> 11.6.4. ensuring that your Application accurately and adequately discloses, either through a privacy policy or otherwise, how you collect, use, store, and disclose data collected from visitors[…].

The above clauses show that Amazon's solutions (by default) are intended for low-risk use cases, since the cloud provider does not take any liability for the security of the content stored in the cloud. In the case of Digitalisér.dk it seemed that the contract proposed by Amazon was considered acceptable, since the benefits (reduced hosting costs) outweighed the risks and disadvantages (such as, at least de facto, absence of auditing rights and extensive guarantees).

From a more general perspective, NITA's experience is that the focus on outsourcing is pivotal but the differences between cloud hosting and traditional hosting, from the business and technical point of view, are not as big as many may believe. Of course, the content of the contracts and SLAs that apply in the case study is different from that of SLAs that would be used under the previous hosting agreement, and the use of an 'off-the-shelf' solution precluded the possibility of asking the provider to modify some aspects of the cloud service provision.

In that respect, the business relationship with a larger cloud provider may be dramatically different from that with a traditional pre-cloud hosting provider. In the case of Amazon, standard terms will be applied unless the client has sufficient bargaining power.

As regards compliance controls, the EC2 SLA states that the user is entitled to service credits if Amazon fails to provide the cloud services at the promised level:

> **Service Commitment**
>
> AWS [Amazon Web Services] will use commercially reasonable efforts to make Amazon EC2 available with an Annual Uptime Percentage […] of at least 99.95% during the Service Year. In the event Amazon EC2 does not meet the Annual Uptime Percentage commitment, you will be eligible to receive a Service Credit as described below.

The request of service credits and the Annual Uptime Percentage must be accepted and confirmed by Amazon, which then will proceed with the reimbursement of the credits.

From a technical perspective, the infrastructure used to deliver the cloud services to NITA is situated in the United States, or in other words Digitalisér.dk runs on cloud servers located in the United States.

In practical terms, the migration to cloud computing did not bring a radical change in comparison with the traditional hosting models that NITA was using before the advent of the cloud. It required only a limited amount of awareness raising with the clients, in the sense that they should understand that data and resources are stored in undefined placed, ie, in the cloud. Nonetheless, the application of technologies in cloud computing (principally the systematic use of virtualisation and hypervisors) is new, which was the key reason to test cloud computing and to see if it can fit the missions of NITA. Generally, the migration to cloud computing did not pose significant threats to Digitalisér.dk, and no new major technical risks and challenges were expected by NITA.

From a financial perspective, the expectation of a return of investment did not play a major role in the decision of NITA to migrate Digitalisér.dk to the cloud. However, the financial expectations are generally positive: NITA foresees that it will recuperate the migration costs within nine months of the migration; thereafter, cloud hosting should be marginally cheaper and thus realise small cost savings.

As regards training requirements for technical staff at NITA, little particular internal training has been required; this was partially due to the fact that much of the migration process was also outsourced. NITA provided the architectural specifications for moving to the cloud in the first place and managed the project as a whole. The online Amazon tutorials were useful to train the system administrators, who did not encounter major problems during the migration.

7.5.5 Experiences and lessons learned

The experience of switching to the cloud was largely positive. No major problems were faced, and in practice almost nothing changed in comparison with the hosting model followed before the migration. The main lesson learned is that migrating to the cloud is possible and that no major efforts are required, at least for this type of use case.

Furthermore, migrating has brought small financial benefits, as running Digitalisér.dk in a cloud computing environment is marginally cheaper than it was before. Since the reasons for NITA to migrate to the cloud related mainly to the desire to conduct a successful test and to reduce the hosting costs, the planned objectives have been reached.

More specifically, the key lessons can be formulated as follows:

- Understanding the technology: NITA increased its understanding of what cloud computing enables, but also the many parts that don't change. The approach is based on the idea that cloud computing is a form of outsourcing with more flexible technological characteristics, and in this sense the experience confirmed these expectations.
- Clearing up terminology and misconceptions: the cloud should be treated as an alternative technological approach which can be preferable to other hosting business models especially when scalability/flexibility is needed. From a business perspective, not much changed after the migration to the cloud, apart from the slightly reduced hosting costs.
- Applicability: cloud computing proved to be applicable to situations where no intense scalability is required.
- Business as usual: cloud computing does not imply dramatic changes in the business of the users, and little special training was required.

7.5.6 Remaining issues and concerns

NITA is confident that all relevant issues have been identified and addressed by the current solution, since no particular problems have been foreseen and/or appeared in reality. The system administrator and the people assigned to manage and follow Digitalisér.dk are well prepared and the migration to the cloud did not create difficulties for them.

NITA runs its business as usual and although cloud computing undoubtedly brought interesting advantages it did not revolutionise the use case. This applies also to users of Digitalisér.dk: for them, cloud computing did not change their experience of the web site, and this too was fully expected before the migration.

7.6 Case study 2: the City of Los Angeles

Los Angeles is the second largest city in the United States. The City of Los Angeles' core activities include the provision of municipal services (water, power, rubbish collection, etc.) for some 3.8 million residents. The City administration has 30,000 employees including 17,000 working for the Los Angeles Police Department (LAPD).

Their cloud computing project, termed Google Enterprise E-mail & Collaboration System (GEECS), was commissioned by the Information Technology Agency (ITA) of the City; it migrates the City's employees email and related systems to its Gov Cloud, a tailored version of the Google Apps suite designed for public sector use. The ITA provides all telecommunications regulatory services, cable television licensing and franchise enforcement, municipal cable programming and utilisation and operates the City's cable channel and related matters. It also has responsibility for planning, designing,

implementing and operating and coordinating the City's information technology systems and networks.[140]

7.6.1 Description of the use case

The City of Los Angeles wanted to achieve a number of objectives in taking the decision to use GEECS. These objectives, including cost savings, technological flexibility and availability, are detailed below:

- Firstly, to acquire a state-of-the-art technology solution to drive productivity and efficiency and enable departments to move forward and communicate in ways that they have been unable to do so now.
- Secondly, to improve information provision – so that City employees can get whatever information they needs, when they need it, on whatever device they need it on, and also so they can collaborate on all sorts of documents, changing fundamentally the way the City conducts its business.
- Thirdly, to save money without the need for extensive capital outlay. It was predicted by the General Manager of the ITA that the City would be able to achieve US$4–10m in hard savings, US$5–10m in soft savings (avoided costs) and US$10–20m in productivity enhancements (increased productivity resulting from the use of a common set of tools, such as calendaring, email and collaboration).
- Fourthly, to allow the reprioritisation of staff to other areas of technology deployment within the City that are currently understaffed.
- Finally, to improve availability: the high availability of Google Apps in comparison to the City's existing email system was noted. It was reported that the City experienced *c.*550 minutes of Post Office email failures and downtime in the twelve months between August 2008 and August 2009, compared to *c.*15 minutes of downtime experienced by Google Apps.

In internal discussions the fact that City would be the first major governmental entity to deploy a Google Apps system was also highlighted. However, reference was made to the use by the City of Orlando and Washington, DC, municipal authorities of Google products, as well by other large private sector organisations.

GEECS involves the migration of some 30,000 employees onto GovCloud. This cloud, representing a collection of Google Apps, includes a general system (for all staff) comprising email, antivirus and anti-spam, storage, archiving, e-discovery and training. Optional services per department were also provided such as an office productivity replacement, collaboration tools, file version control and storage virtual drives.

The commissioning of GEECS was highlighted in the media since Google appeared willing to enter into negotiations with the customer in order to secure the contract, specifically around a number of contractual concessions.

The migration schedule to GovCloud was relatively speedy, with non-LAPD employees migrating between January and July 2010, following a pilot of 3,000 users. This was

[140] See (as of 23 November 2010) http://cityclerk.lacity.org/cps/pdf/govtglnc.pdf

originally scheduled to be followed by full migration, by October 2010, of LAPD employees.[141] However, as of August 2010, the LAPD had only 50 users testing the GovCloud environment,[142] due to the additional protection requirements above and beyond that of the wider City.

7.6.2 Chosen cloud solution

In moving to the GEECS solution, Computer Sciences Corporation (CSC) was employed to implement, migrate and deploy Google Apps Premier Edition. Google Apps Premier Edition is a Software as a Service (SaaS) solution. This was chosen to replace an outdated deployment of Novell GroupWise used throughout the City municipal authorities for email. Once the system has been implemented, users are asked to accept, with an on-screen click, a separate usage contract with Google who run the system.

The contract is reportedly worth US$7m to CSC and Google, although, according to reports, other state entities within both the City and County of Los Angeles will be able to take advantage of the services. This 'piggybacking' permits the City to receive pricing discounts on its contract if other additional entities choose to contract with CSC/Google. The savings were identified to be either US$5 or US$10 per person for each following year's annual payment, depending on the extent of further piggybacking on the contract by other state entities. Thus, all state and local government agencies are eligible to accede to the contract and as of September 2009, the ITA had received twenty 'serious' enquiries from other California state and local government agencies.

Observers have noted that a contract with CSC as the prime contractor and Google as the subcontractor is unusual in a cloud environment. However, this structural arrangement, in which a valued-added reseller (VAR), such as CSC, builds its services on top of the services provided by another company, such as Google, represented thirteen of the fourteen offers the City received in response to its invitation to tender, meaning that this type of legal construct was difficult to avoid. In this arrangement, the City only has a service contract with CSC, whereas it has only signed a non-disclosure agreement (NDA) with Google.

Google Apps Premier Edition is a purpose-built messaging and collaboration platform. The suite of applications includes solutions for business email, information sharing and security. Common services include email, word processing, spreadsheets, presentations, calendars and instant messaging without automatic saving. Additionally, under the solution proposed by CSC, further features such as Forms, multi-purpose wiki-style web pages and Google Video are provided. In the presentation of this offer to the City, the main business benefit was proposed to be the simplification of the creation, production and dissemination of information, enabling City workers to focus less on use of the technology and more on production of information. Other characteristics that were identified as potentially valuable by the municipality include the potential for seamless interoperability and easy collaboration. Finally, it was anticipated that the short learning curve and intuitive interface (together with the public popularity of Google Apps) would

[141] Kevin Crawford, ITA, pers. comm. 16 July 2010

[142] Maggie Goodrich, Police Administrator III, Information Technology Bureau, pers. comm. 2 August 2010

offer opportunities for reducing training time and minimising disruption caused by changing over from Novell GroupWise.

Nonetheless, as official documentation relating to the deployment illustrates, the cost issue was one of the key drivers. Financial management information from the City authorities indicated that continuing to use Novell GroupWise would involve considerable further expense. For example, the currently deployed solution for email was some generations behind the latest edition of the software and therefore updates, with requisite budgetary implication, would be required.

The Google Apps SaaS solution was integrated with the City of Los Angeles Identity Management (IDM) infrastructure, based on Lightweight Directory Access Protocol (LDAP) services, to enable more secure authentication than the standard username and password provided with the public Google Apps offering.

7.6.3 Risks and challenges of cloud use

This use case may be understood as medium risk, for two main reasons. Firstly, the sheer breadth of service and potential impacts resulting from interruption (email, archiving, anti-spam) represent one key set of indicators as to its risk level. Secondly, the particular security issues associated with the replacement of the LAPD's CLETS (California Law Enforcement Telecommunications System) prompted a range of discussions in the City's law enforcement and police community after the signing of the initial contract, but prior to deployment, regarding the security aspects of sensitive information relating to criminal investigations being stored in the cloud.

Since the LAPD has access to sensitive electronic data, any email system that it uses will have to be secure enough to allow the safe transfer of this information. As the sensitive data are held in accordance with Department of Justice (DoJ) requirements, the email system will also have to adhere to DoJ rules. These requirements, to be enforced by the LAPD, are that the data must be: encrypted to a known standard certified by NIST; held in an auditable system; and that background checks need to be conducted on those who have access to the data. These requirements were not in the original contract and there were initially objections from Google. This was due to the level of alterations required (Google does not generally use NIST-certified encryption, but rather obfuscation of data) and the new procedures insisted upon (the requirement to do separate background checks on Google staff for each user organisation of the GovCloud product). However, these conditions are now accepted by Google, and elements that do not currently exist, such as the auditing tools for the LAPD, have been placed on the road map for future development. The original contract between the City and CSC is in the process of being rewritten to reflect these alterations.[143]

The alterations to reflect the requirements of the LAPD and the DoJ will build on the current version of GovCloud, 1.5, delivered on 22 June 2010, which holds sharded and encrypted data in one or more of three separate storage devices housed at existing US data centers.

[143] As at 3 August 2010

Other important challenges raised by the decision-makers include concerns relating to availability (eg, the consequences of a loss in Internet connectivity in key working areas), the knock-on effects of the deployment on other services (eg, whether the use of an Internet-based service will result in increased latency for other network applications), and the fact that disaster recovery is now entrusted to Google.

7.6.4 Chosen approach to address these risks

The associated contracts, SLAs and other agreements between the City of Los Angeles and both CSC and Google appear to be the main evidence of the approach to address the risks. So far, as the use case is in its preliminary phases, details of operational aspects and how the providers actually meet the requirements set out have not come to light. However, there is enough detail in the terms of the contracts to enable discussion on what legal, technical and other measures are present to provide for privacy, trust and security.

CSC (as the primary contractor to the City of Los Angeles) is required by the contract to establish a security programme to ensure the confidentiality and availability of protected information. The contract obliges CSC to comply with all federal, state, local and international law covering privacy and security. The City has a right to conduct a security audit and CSC has the responsibility to implement any security safeguards identified by the City in its audit. These audits can be on an ongoing basis, from time to time and without notice.

Additionally, CSC is required to undertake an SAS 70 or equivalent compliant audit of Google's information security programme and release the results of this audit to the City authorities upon request. However, as Google will be storing, processing and transmitting the data, audits from CSC may offer limited added value, depending on their scope (for example, whether CSC were permitted to check Google's infrastructure as well).

CSC is contracted to indemnify the City under certain circumstances, but there are caveats and limitations of liability (LOL). For example, there is no indemnification for breaches of privacy and security laws. Additionally, breaches outside warranty are subject to LOL.

It was stated in official documentation provided for the case study that Google has agreed in its contract with the City to provisions regarding a higher level of audit, transparency and security than normally offered as part of its cloud services. In practice, Google can perform a full audit in a two-week timeframe, and the City can itself audit individual emails spontaneously. However, the ability for the City to fully audit email communications is currently still under development.[144] Despite these abilities, the arrangement lacks details on incident response procedures, such as forensic investigations into data exposure.

Google is not contracted to comply with specific privacy and security laws, just federal and state breach notices. Ironically, if the City ends up violating relevant privacy laws, this may be viewed as a breach of contract on the part of the City.

There are also security provisions in the Google contract in relation to its 'facilities' but these apply only in respect of the contract between Google and CSC, rather than between

[144] As at 3 August 2010

the City of Los Angeles and Google. These provisions require that data from the City of Los Angeles is kept in facilities which adhere to 'reasonable security standards', similar to those in which Google stores and processes its own data.

The detail of this section of the agreement seems to indicate that Google regards its current systems to be 'compliant with industry standards'. Nonetheless, a purely legal interpretation of this form of compliance compared to reasonable security may be viewed as insufficient.

The SLAs contain confidentiality and privacy obligations, notably that CSC cannot disclose information to any third party without the City's written consent. 'Information' as defined by the NDA includes: 'technical information', 'business information', 'security' or 'rights'. The NDA makes it clear that the City retains ownership of all the information.[145]

7.6.5 Confidentiality provisions

During the consideration and selection of Google Apps, security and confidentiality concerns were raised by the LAPD and City Attorney. According to official documentation, these concerns related to sensitive information, subpoenas and DoJ requirements. The Californian Department of Justice and the LAPD worked closely with the ITA during the design and preparation phases to ensure that confidentiality concerns were met.

In this situation the City retains ownership of the data, but the storage, processing and transmission of data has been entrusted to Google. If Google violates security and privacy laws, the City as 'data controller' is still liable. Thus, having a contract framed to prevent breaches, rather than to allow subsequent recovery or for the providing of financial compensation, may be best.

The CSC contract with the City allows for unlimited liability with respect to privacy and security breaches, and additionally responsibility for certain acts, errors and omissions by its sub-contractor Google.

Google's arrangement by comparison, contains broader confidentiality provisions which include customer data as confidential information (which by extension covers personal information stored, processed or transmitted through Google's services), as well as a stipulation that Google will refrain from disclosing confidential information. Google is required to protect such information 'with the same standard of care it uses to protect its own confidential information'. This distinction between an obligation not to disclose and an obligation to protect may suggest that Google is more comfortable with broader contractual security obligations than CSC. However, the arrangement does contain LOL and caveats for non-indemnification covering acts such as violation of privacy and security laws. Another point not defined in the arrangement is consideration of whether liability exists in respect of mandated disclosure (as might be required by law enforcement agencies).

[145] In the context of the regulatory framework of EU privacy and data protection laws, a practical comparison would make the City the 'data controller'

Both the CSC contract and Google arrangement contain provisions regarding compliance with relevant laws which would appear to be broad enough to include relevant privacy and data security laws.

Furthermore, technical measures have been deployed to address security, namely that other directory authentication procedures will be used to gain access to services, rather than the standard username and password used on Google's free public cloud offering. Additionally, City users will only be able to access Google services on City-run networks, provisioned mobile devices or via VPN, identical to how security was currently provided. To access data on GovCloud, a location key, an encryption key and a domain key are all required, meaning that one cannot access others' data. Rules regarding password rotation and strength will also continue to be enforced. City of Los Angeles email addresses are not publicly available from the Gmail directory and the City will continue to use its own email addresses (@LACity.org).

From a legal perspective the contract between Google and the City includes the following contract terms:

- Unlimited damages for a data breach
- Provisions allowing audits
- Guarantees that the data remain in the contiguous forty-eight states and that this is auditable
- Penalties if Google services are unavailable for any longer than five minutes per month
- Unlimited damages if the NDA is breached by Google
- Requirements for encryption and breaking data into pieces
- Bar on Google viewing data without permission from the City.

7.6.6 Experiences and lessons learned

The implementation of GovCloud has progressed successfully. According to the City, the cloud solution provides great versatility and swift updates. Certain workarounds have been used, but these are not detrimental. The VAR arrangement has worked relatively well, coupling Google and its interest in moving into the public sector and CSC's previous experience. Because of its experience in the public sector, CSC was able to accelerate the transition to Google's cloud-based services.

In the contract between CSC and the City, it was reported that there was insufficient due diligence undertaken by the City during the piloting stage. Problems arose during the migration regarding the incomplete transfer of contact and calendar information. Additionally, timelines were an issue as there was a miscommunication of deadlines and requirements between the City and CSC.

Regarding the interaction of the City with Google, since this was Google's first move into the public sector there were minor points of contention. However, the City was generally happy with the interaction.

The LAPD believed that the adoption of GovCloud was complicated by the internal arrangements within the department and by interactions with Google as a technical sub-

contractor.[146] Given that the project was led by the City's ITA without LAPD being present (despite the fact that they represented 40 percent of the final users), important topics were missed in the original contract, leading to contractual rewrites. Additionally, the chains of communication were perceived to be too long, with the LAPD having to liaise via CSC to talk with Google.

The LAPD focused on the protection of sensitive data, meaning that user interface issues were not as prominent as they should have been. Thus when interface issues were identified they had to be resolved late in the process, when ideally they could have been addressed considerably earlier.

Internally within the City, there has been some reluctance to migrate from Novell GroupWise to Google Apps. However, following initial opposition[147] Google has been extremely responsive in including additional features in their solution to better match those of GroupWise.[148] Of 100 feature requests, Google had implemented 70 by mid-July 2010.

As a matter of policy, City users were given access to both GovCloud and GroupWise for a week after the migration. This lead to synchronisation complications, and other errors, which users blamed incorrectly on GovCloud. With the benefit of hindsight, the City would have preferred not to allow this simultaneous access. Additionally, the City emphasised the need for better communication and better training to keep its employees informed of changes.

The initially anticipated financial savings are hard to quantify, particularly with respect to additional features of GEECS. The GroupWise system is actually cheaper than a package containing just Google's email solution. However, once additional Google Apps are available to the LAPD and other departments, savings are expected to be realised via productivity enhancements.

7.6.7 Remaining issues and concerns

There are still a number of remaining feature requests, and further adjustments are expected to be identified with the migration of LAPD employees in late 2010. However, the City anticipates most of these will be addressed.

The LAPD is concerned about the level of trust that is being placed in Google. Google are to be contractually obliged to provide the required auditing tools, but these are still in development. They are expected in early 2011,[149] but given the sometimes strained relationship between Google and the LAPD, the latter will feel more comfortable when these tools are in existence and use.

Concern also exists within the City regarding a lack of employee buy-in to the migration to a cloud solution. However, the City expects cloud solutions in general to develop in the future and that its own solution will correspondingly improve. Improvements are expected

[146] Maggie Goodrich, Police Administrator III, Information Technology Bureau, pers. comm. 2 August 2010

[147] Ibid.

[148] Kevin Crawford, ITA, pers. comm. 16 July 2010

[149] Maggie Goodrich, Police Administrator III, Information Technology Bureau, pers. comm. 2 August 2010

to include increased storage space, additional products, and better collaboration, all of which will address some of the remaining employee concerns.

7.7 Case study 3: EU eHealth provider

The use case analysed here has been set up by a consortium of seven private entities in a specific EU Member State.[150] While the entities provide their own technical expertise with respect to cloud services and health care, Amazon delivers the physical cloud capability.

Other partners of the consortium include a regional medical council that provides patients' data and medical knowledge, entities providing pharmaceutical and medical information, and a social network web site, which has the main goal of distributing information among users.

7.7.1 Description of the use case

The use case involves the processing of patients' data in the cloud, thus it is a cloud experiment in the area of eHealth. Its purpose is to provide health data on specific patients in the chosen region, via a cloud infrastructure, to doctors of the same region. Doctors can retrieve the data directly and provide patients with treatment and medicaments.

The use case in question is run on a non-commercial basis: the main goal of the consortium is not to make profit but rather to test the possibilities of cloud computing in the area of eHealth and to assess if the use case can be extended to a national level. An important point is that the project was created *ex novo* in a cloud environment and was not running before under a different hosting model.

The data moved to the cloud are the following:

- Physicians' information (including address and contact details).
- Patients' health records: these do not relate to specific illnesses but rather are generic records including all relevant health-related information.
- General health information: this information is not related to any specific patient.
- Details and locations of medical centres and pharmacies: this information is generally useful for users and can be retrieved through a link with Google Maps (the same applies to the abovementioned physicians' information).

The tool, whose goal is that of sharing data and information about patients of the concerned region, is accessible only by doctors of that region. It is also accessible by patients, so that the list of users includes physicians and patients. Health insurance companies are currently excluded, but they may be included in the future. The system is still at an experimental level but nevertheless it is operational in the concerned region.

Regarding the main objectives to be achieved by the use case, the general (and ambitious) aim of the project is to get valuable information for physicians and other health care professionals from personal health information records (based on patient data) and a

[150] For the purposes of this study and at the request of the participants, this case study has been anonymised

general health information database (based on medical knowledge), and for consumers to get information on physicians, their locations and addresses (including a route map) in a given country, accessible on any device from any location.

The solution allows interaction between patients and doctors in the sense that:

- Doctors can retrieve information about patients from the cloud, so that every physician can obtain detailed information about the patient even if he/she never seen the concerned patient before.
- Patients can obtain information about names and locations of doctors in the whole region (and potentially throughout the whole country).

7.7.2 Chosen cloud solution

The use case falls into the category of public and non-commercial cloud implementations, since the goal is not that of making profit but rather to propose an open solution for doctors and patients in a specific region. The cloud solution adopted is public in the sense that it uses the infrastructure of Amazon and not an ad hoc cloud.

The consortium, in fact, decided to adopt the cloud solutions offered by Amazon mainly because of their low hosting costs.

7.7.3 Risks and challenges of cloud use

The use case presents important risks and challenges, especially in the area of privacy and security threats. While data loss or service interruption has no direct impact on the user's health or safety, crucial (patients') data are entrusted to the cloud, meaning that there are substantial privacy challenges to be managed. The risk of significant damage to the quality of health care is limited, since no data perceived as of critical importance from the operational or business perspective are uploaded in the cloud, and above all because in case of technical problems or failures, doctors can retrieve patients' data using other means. The solution does not aim to substitute traditional health records but rather to complement them, so that physicians have alternative and more efficient ways of retrieving patients' data and information.

The consortium was and is fully aware of the privacy issues that the use case concerns and that they need to be carefully managed in order to avoid major incidents. Although the failure of service would not cause much damage, no major technical risks are involved, and no lock-ins are expected, the consortium felt that the data were very crucial with lot of privacy involved. For this reason, they performed a systematic risk assessment before starting the project. See below for further information about the solutions adopted.

7.7.4 Chosen approach to address these risks

The main security and privacy issues have been addressed by a careful contractual architecture and through technical means. From the legal perspective, the consortium realised that the standard Amazon SLA was not sufficient to mitigate the privacy risks. They managed to negotiate with Amazon the content of an ad hoc SLA, which is very similar to the standard SLA and users' terms of Amazon but differs in one sensitive clause.

Amazon agreed to store all the data uploaded by users in the framework of the use case only in servers inside the European Union, with the further commitment not to migrate the data to cloud components outside the European Union itself. Three issues should be

highlighted: (i) this contractual arrangement has been signed by Amazon and by the leading partner on behalf of the whole consortium; (ii) the modified SLA and users' terms are intended to avoid privacy risks and reduce compliance costs for the consortium, which is now not obliged to obtain authorisations from the national privacy authority concerned for the export of data outside the European Union; and (iii) the other standard clauses of Amazon's SLA and users' terms basically remain applicable.

It should also be pointed out that no specific maintenance agreements have been signed and that the consortium has auditing rights towards Amazon, which can thus be audited by the consortium itself. Generally speaking, this auditing right is not common, but it is necessary due to the nature of the sensitive data involved. External legal auditors are involved in the use case as well.

From a technical point of view, guidelines such as incident handling mechanisms, logging obligations, and so on, have been established. Regarding incident handling mechanisms, a complete backup of all data is regularly performed but no other tools have been foreseen. There are no dedicated servers to be used in case of incident. The consortium believes that Amazon is reliable enough for the hosting of the patients' data, and as pointed out above the loss of such data would not impede the treatment of patients by doctors.

As for the timing of the use case, it started in December 2009 and after 2 weeks of trial period, the experiment proper began. The use case was revealed to be stable during the trial period; presently it is still in the experimental phase before full and definitive implementation. It is foreseen that a business continuity plan will be drafted and implemented after this experimental phase is complete.

From an operational perspective, no migration planning was needed, since the data stored in the cloud were not previously hosted under a different hosting model. Only limited staff training was required, since many of the people involved in the consortium had previously gained experience with cloud computing and therefore were already familiar with the technology. They are aware of the risks and challenges of cloud computing. Training is rather more required at the doctors' side, and it was (and still is) necessary to provide physicians with detailed information about the operating modalities of cloud computing.

Finally, in brief, the primary advantage of this project is that the data of patients from a particular region of a European Member State are made available via a cloud provider in Europe to be retrieved by doctors in the same region. Patients too can obtain information on relevant physicians. Although the use case does not fall in the area of e-government (no public institutions are involved in the project), no revenues are foreseen, no fees will be charged, and the use case is operated on a non-commercial basis.

7.7.5 Experiences and lessons learned

The use case brought (and continues to bring) valuable insights to the consortium. In particular, the following issues need to be highlighted:

- Although the case study envisions an integrated platform linking hospitals and healthcare centres across the region (and potentially within a country) allowing secure access to shared information with patients' records instantly available, the support from the health centres and physicians was not complete as some were afraid to entrust patients' data to the cloud due to

privacy concerns. In other words, even if the benefits arising from the use case are notable for the users (especially for doctors), some health professionals felt that the risks involved were very substantial. Nonetheless, some 75–80 percent of the physicians of the region concerned did participate in the project and shared their patient data in the cloud.

- The project seems to be viable and interesting so that a further implementation at national level would be desirable.
- However, such further implementation would be more difficult to realise, with greater risks of confusion when a huge amount of data are migrated to the cloud. Although hosting costs and possibilities are not big concerns, since cloud computing is potentially unlimited and relatively cheap, the migration of large amounts of sensitive data is not a trivial task.
- The use case does not pose significant problems and difficulties from the technical point of view; major issues, however, are expected to arise if the consortium decided to change cloud provider and to migrate all data to the new supplier, due to the need to extract and relocate substantial amounts of highly sensitive data.

Generally speaking, the use case brought benefits in the sense that doctors could more easily obtain detailed and updated information about patients of the region even if they did not know a patient personally and had not treated them before.

7.7.6 Remaining issues and concerns

The implementation of the use case required a high level of care considering the sensitive nature of the data involved. Due to this, the consortium adopted a phased approach to meet the needs of various stakeholders (in particular doctors) and embraced best practice recommendations. They have also engaged users in the process and in joint application development.

However, although 75–80 percent of the physicians in the concerned region did collaborate with the project, reluctance to provide patients' data may be an issue if the use case should be further developed and extended to the national (or even international) level. Obtaining the consent of the doctors involved in the experiment to provide patients' data was not an easy task.

In general the project has been well received and future plans of further nationwide implementation are under discussion. There are, of course, still issues to be carefully assessed, such as respect of doctors' professional secrecy. In this sense it may be recommended that the local or national physicians' associations are involved and provide the consortium with guidelines, recommendations and useful tips about how to avoid the processing and undue disclosure of information covered by professional secrecy.

Security is also a major topic, since it is absolutely essential that the cloud solution envisaged eliminates security threats as much as possible. An open cloud may not be the best solution, but in any case the experiment, which for the moment runs on a commercial open cloud infrastructure, demonstrated that cloud computing may bring benefits in the field of eHealth.

7.8 Implications

This section presents an analysis of the primary risks and challenges, along with the observed strategies to address these, arising from the practical implementation of the case studies described. The use of third-party cloud computing services can be considered as a form of outsourcing with more flexible technological characteristics: therefore a key question is if and how the existing policy and regulatory framework is capable of handling this increased flexibility. Is it necessary to rewrite the regulatory rule book in order to deal with cloud computing? And what dynamics or aspects (if any) are different enough to warrant specific regulatory intervention? Some of the emergent issues arising from the case studies are presented below.

7.8.1 Cloud computing is still experimental

Whichever way cloud computing is defined, from a pragmatic perspective cloud competing deployments have been around for only 2–3 years. In this space of time, many organisations implementing cloud computing have barely got to grips with the realisation of commercial benefits, let alone grappling with the complex issues of security. With that in mind, based both on the interviews with stakeholders and the case studies above, it seems that organisations are growing increasingly keen to experiment with cloud-based models to try to take advantage of the supposed benefits, whilst relegating compliance and security issues to a secondary role (or even not considering them at all).

As will be explored further in this report, each of the case studies revealed specific challenges or weaknesses which were not identified until a relatively late stage in the project, or not at all, such as:

- The lack of specific provisions with respect to personal data exports in the Danish use case, based on the consideration that a social network was estimated to be a low-risk application.
- The regulations with respect to confidential information established by the Department of Justice in the Los Angeles use case.
- The issue of patient consent with health care information processing in the eHealth use case.

Especially in an era of economic crisis, companies may be inclined to choose the presumed economic benefits of cloud solutions and leave the more ambiguous compliance and security issues for a later stage. This obviously creates risks that need to be managed at some point.

7.8.2 Moving to the cloud makes it even more imperative to be fully aware of your own risk tolerance and security level

Cloud providers (much like any other type of service provider) offer standard terms with generally limited guarantees in respect of privacy and security. This is not a problem provided that the users of cloud services have the diligence and negotiating power to identify their own risks and responsibilities, and to seek appropriate solutions when standard terms are inadequate. For those cloud customers that have more unique or complex security arrangements and are aware of them (as in the eHealth and LAPD cases), they can negotiate tailored SLAs which permit a greater degree of assurance.

One consequence is that cloud customers have to be more aware of their security needs before embarking upon a cloud deployment, lest they accidentally outsource risks that they were insufficiently aware of to the cloud service provider (who may decline any liability for these risks under its standard terms). So, paradoxically, the benefits gained in terms of reduced capex from moving to the cloud might be lost if it is sunk into understanding the internal security landscape and risk tolerance.

On the other hand, it should also be recognised that this risk is not entirely new or unique to the cloud: any type of outsourcing contract involving sensitive data or services requires due diligence on the side of the outsourcing party, along with a framework to ensure that liabilities and guarantees are correctly managed. The main new element in the cloud is the greater degree of flexibility and complexity with respect to the data and services, which requires appropriate provisions (eg, with respect to applicable law, competent courts, location where data/services will be hosted, and SLAs).

These issues are not unique to cloud services, but have become much more apparent due to the global scale on which cloud service providers operate, the smaller inherent transparency of their infrastructure, and the perception of greater operational ease of use which often belies the complicated legal and trust requirements behind a project (even if it is cheap and simple to move into the cloud, that does not necessarily mean that compliance, security and trustworthiness will be equally cheap and simple). On this point, cloud projects require a greater degree of diligence, at least until users have more experience with this type of model.

7.8.3 Can cost drivers override security?

Another challenge consistent with take up of new technology or computing paradigms by customers is the extent to which cost concerns can override security. This may be seen in the City of Los Angeles case study where despite the LAPD representing the majority end user community of the case study, six months after initial deployment only 50 LAPD users were converted to the system. This was due to ongoing concerns regarding security. In current harsh financial climate, it can be seen that this might evolve to become even more of a pressing concern. In the case of the LAPD the cost drivers were significant with expected gains in the tens of millions of dollars; nonetheless, the security issues associated with the replacement of the LAPD's CLETS prompted a range of discussions in the city's law enforcement and police community after the initial contract was signed but before deployment began to this class of users. Indeed the auditing tools for the LAPD system are now being placed on the roadmap for future development.

7.8.4 Integration with existing security measures

The City of Los Angeles demonstrated with its use of Lightweight Directory Access Protocol (LDAP) authentication that integration with existing IDM infrastructures is possible in a cloud computing deployment. Indeed, it may be that this is the only way to retain assurance of certain aspects of security or privacy controls, given the challenges for cloud computing outlined in the literature assessment. The lesson, therefore, is that whilst contracting out to the cloud may improve the overall security position of an organisation (since cloud computing can also have positive aspects for security), this is not a foregone conclusion (except in cases where no real security measures existed beforehand, which may well be the case for smaller organisations, including SMEs). However, especially in cases

where more sensitive information is entrusted to the cloud, any inherent security characteristics of the desired cloud would have to be carefully measured against security needs in order to assess if further steps would be necessary. It is not always possible to go from a no security situation to a more secure situation simply by turning to cloud services; rather, some initial security and trust mechanisms will typically need to be in place in the cloud customer organisation prior to any deployment, especially in medium or higher risk application fields.

7.8.5 Untoward dependencies can be a problem

One of the operational challenges cited by the City of Los Angeles case study regards Internet connectivity and the availability challenges presented by the reliance on a service so heavily dependent upon Internet connectivity. Although many approaches to mitigating this risk focus on considering availability at the end point (eg, via mechanisms to keep some kind of offline last-known good representation of the data), what is not considered is how the resilience requirements evolve in a cloud computing deployment and how it may require potential cloud users to devote more consideration to their own resilience (eg, whether they have dual home Internet connections with (physically and logically) separate service providers).

7.8.6 Granularity *is* possible (but only if you have the right cloud provider and sufficient negotiating power)

Cloud computing services are often likened to utility services 'on tap': they offer a homogenous service which can be elastically increased or decreased as needed in a relatively easy fashion that requires limited technical expertise and which is ultimately cost effective to the user. However, this metaphor also suggests a weakness, namely that of offering a standardised product with limited customisability, which may not match cleanly with the end users' needs. This begs the question: can a cloud service user realistically obtain customised service levels that ensure compliance with technical, legal and operational requirements?

Both the City of Los Angeles and the EU Member State eHealth case studies show that a degree of customisation of 'off-the-shelf' cloud solutions is in fact possible. For instance, in both of these cases questions of data location were dealt with in a slightly more granular fashion to address specific needs. In the case of GovCloud, Google has asserted that data will be held in three separate storage devices located within existing US data centres and that this is auditable. In the case of the EU Member State eHealth case study, the contractor asserted not to place sensitive patient data outside of the EU, in order not to run afoul of personal data export rules. Nonetheless the ability to achieve this is clearly dependent upon the intersection between the cloud customer and the market: if the main cloud providers are based outside of Europe then they might be less inclined to meet contractual requirements of EU-based companies who stipulate, for example, that all personal data must be kept within the EU.

The same observation applies in relation to other crucial issues to be governed through contractual terms, such as applicable law, competent courts, auditing rights and possibilities, etc. Again, this issue is not unique to cloud computing, as it manifests itself in precisely the same way as with traditional outsourcing projects: both parties will have their standard terms, preferences and prerequisites, and the final result will typically depend on

negotiating power. In that respect, it is worrisome that the leading cloud service providers (including those identified in the two EU use cases) are established outside the EU, as this implies that EU law will typically not apply unless users of cloud services manage to negotiate a favourable outcome on this point. This is also relevant because it implies that EU regulatory intervention with respect to cloud service providers is unlikely to have a real market impact: key service providers are simply not subject to EU law. Rather, effective policymaking should target the users of cloud solutions, who may well be European and subject to EU requirements, as (again) the two EU case studies show.

CHAPTER 8 Gap analysis

This chapter presents an overview of gaps in various aspects of current European policy approaches relevant to cloud computing. The available approaches will be characterised, including legislation, standards, guidelines, support to implementation and R&D. This gap analysis is designed to indicate where (from the literature, case studies, interviews and expert workshop) gaps, problems, challenges or inadequacies exist in current European policy approaches to addressing the previously identified security, privacy and trust issues of cloud computing.

Three distinct policy domains of relevance can be identified: legislation, implementation and ICT research. With respect to identifying gaps in legislation, some may be clear; there may be wording of laws that is regarded by stakeholders as outdated or requirements imposed upon the subject of such regulation which may not be viable or an appropriate way of achieving outcomes proposed in the legislation.

The multi-stakeholder necessity of successfully addressing such issues means that some of the gaps are not necessarily in specific legislation per se, but rather reflect areas where industry and Member States could improve implementation. Nonetheless, there exist certain 'levers' which may be pulled by policymakers to improve this, most notably the use of means to encourage the private sector to improve transparency and security. Some of these levers include the dissemination of additional interpretative guidance to support efforts to improve the implementation of practices to achieve positive security, privacy and trust outcomes. Other examples include standards (the most well known being the NIST standards and the ISO2700x suite of Information Security Management System best practices) which in this context have a dual use. Firstly, by capturing and codifying a canon of internationally recognised practices regarding the management of information security risks, such 'standards' can evolve into useful guidance derived from front line 'operational' practice. Secondly, it is hoped that the process of third party certification against a strict standard will create a security market allowing customers to differentiate supply based on their processes relating to the management of privacy, security and trust.

In terms of ICT research, this report has emphasised that technology neutrality can have benefits in the operational, legal and economic management of security, privacy or trust risks. Nonetheless, there may be specific challenges or concerns with respect to technology (outlined earlier in this report in Chapter 4) so public policy intervention could be aimed at stimulating scientific and technical research and development to improve or address these concerns. This report also emphasises policy research as a domain aimed at providing the evidence base for the improvement of compliance or implementation of policy.

Each policy domain listed in the table below is assigned a number to facilitate the identification and correlation of recommendations to the specific gap in question. Addressing the gaps will require action by stakeholders who may not necessarily be the owner of the issue.

Table 5. Gap Analysis

Domain	Gap or issue	
	No.	**Description**
Legislation		
Data Protection Directive 95/46/EC	1	*1.1* The definition of data controller/processor and applicability of these terms on cloud models *1.2* Control over personal data *1.3* Location as a criterion for determining applicable law, in combination with differences between national laws *1.4* Rules to support accountability may not be effective or optimal
e-Commerce Directive 2000/31/EC	2	*2.1* Effectiveness of notice and takedown regimes in international context *2.2* Liability rules for intermediary services providers; applicability in the cloud *2.3* Transparency rules *2.4* Applicable law is again linked to physical location *2.5* Exceptions to determining applicable law (consumer protection, intellectual property rights)
BRUSSELS I Council Regulation (EC) 2001/44/EC	3	*3.1* Only applicable in the EU/EEA, and contains exceptions (including some consumer contracts). Only applies when no other choice has been made (rare for cloud models), which can be problematic for cloud users with limited negotiation power
Unfair Commercial Practices Directive 2005/29/EC	4	*4.1* Consumer protection rules: Member States may choose to add more specific rules at the national level (eg, with respect to information to be provided or language to be used); additional diverging requirements may apply depending on the sector/context of the cloud service (eg, information society services, health care services, financial services, etc.)
Electronic Communications Privacy Directive 2002/58/EC, as amended by the Citizen's Rights Directive 2009/136/EC	5	*5.1* Breach notification requirement only applies to Communication Service Providers not Information Society Service Providers *5.2* Communications secrecy obligations may be difficult or impossible to meet in a cloud environment
Data Retention Directive 2006/24/EC	6	*6.1* Data retention rules may not apply to cloud-based service providers
Framework Decision on Attacks against Information Systems 2005/222/JHA	7	*7.1* Attacks emanating from the cloud *7.2* Liability for harmful and illegal content *7.3* Identification of applicable law *7.4* International collaboration between law enforcement bodies
Implementation and governance		
Cloud users (public and private sector)	8	*8.1* Government security risk management programmes may not align with cloud computing (eg, certification or accreditation requirements lead to inefficiencies and inconsistencies in meeting security requirements)

		8.2 Best practices for security management *8.3* Absence of or reluctance to participate in awareness raising programmes for citizens or consumers *8.4* Low accountability in respect of how use of cloud computing could help to improve trust *8.5* Lack of transparency *8.6* Poor management of continuity and resilience
Cloud service providers	9	*9.1* Absence of or reluctance to participate in awareness raising programmes for customers *9.2* Lack of transparency *9.3* Unaccountable or poorly accountable use of data
Risk controls	10	*10.1* Physical access controls *10.2* Application development and maintenance *10.3* Vulnerability management *10.4* Monitoring *10.5* Identification and authentication *10.6* Access control *10.7* Management of encryption and key infrastructure *10.8* Continuity and incident management *10.9* People (security team) *10.10* People (business function) *10.11* Security control testing *10.12* Accredited components *10.13* Residual data *10.14* Asset management
Guidelines	11	*11.1* No official guidance on application of European data protection law in the cloud *11.2* Little or no industry incident identification/ response guidance *11.3* Little or no auditing guidance *11.4* Little or no guidance on forensics *11.5* Lack of consideration of European specificities in currently available industry guidance
Research		
ICT research	12	*12.1* Little exploitation of research on improving encryption technologies to facilitate data processing in the cloud *12.2* Low development of technological interfaces to facilitate measurement of security levels *12.3* Ineffective security monitoring technologies *12.4* Little research on user risk perception and behaviour *12.5* Current technologies for identity management may not be viable in the cloud *12.6* No robust and broadly accepted technology standards supporting interoperability
Policy research	13	*13.1* Absence of research into best practice for metrics: how to understand and assess the maturity of a cloud service provider in respect of how they treat security, privacy and trust risks *13.2* Absence of research into best practice and guidance for privacy in cloud computing (eg, privacy by design, data minimisation)

The table above lists gaps identified from the literature, case studies, expert workshop and interactions with stakeholders during the course of this study.

The evidence presented throughout this study illustrates a number of challenges, both general and specific in respect of a broad range of European legislation that may support

the achievement of high levels of security, privacy or trust. One of the key themes common to many of these legal frameworks is the principle of jurisdiction, which is normally determined by reference to the provider or customer, or in the case of data protection, the location of the data. However, since location-indifference is one of the key traits of cloud computing (and is arguably an important factor enabling users of cloud computing to derive efficiency benefits), this constitutes a serious challenge for such regimes. This challenge is present across many European legal frameworks but there are also specific gaps or inadequacies, detailed below:

1. The Data Protection Directive is regarded as one of the pre-eminent examples of legal regimes seen as untenable in a cloud computing context. The legal framework for privacy and personal data protection established by the Data Protection Directive is regarded as the main means to provide for the protection of the personal data of European citizens. Its clauses relating to defining types of entities such as 'data controller' and 'data processor' were seen as incompatible in respect of cloud computing deployments, as was the capability of the tools foreseen in the Directive to permit data subjects to exercise 'control' over their personal data. Furthermore, the means envisaged in the Data Protection Directive to ensure data controllers can be held accountable for how they adhere to various provisions in respect of data protection were also seen as being challenging in a cloud computing context due to the global nature of cloud service provision and use.
2. The e-Commerce Directive establishes a legal regime to improve trust in the use of e-commerce: however, because of the global and opaque nature of cloud service provision, such rules may not work as well where the information society service provider uses a cloud service whose location is difficult to determine. In addition, the transparency rules in the e-Commerce Directive may be difficult to enforce in a cloud service environment.
3. The framework for contractual law established by the BRUSSELS I regime presents a challenge in the cloud context. BRUSSELS I is effective in supporting trust and confidence by giving each party to a contract assurance regarding what terms apply. However, the evidence illustrates that this could constitute a challenge particularly for cloud users with limited negotiating power (such as SMEs) since it is the default applicable framework in the absence of a joint agreement between the parties as to applicable law may be overly cumbersome and complicated to apply in a cloud computing context.
4. Specific requirements included in national obligations of the Unfair Commercial Practices Directive (such as in respect of the sector or context of the cloud service) may unnecessarily complicate the task of compliance in a cloud environment by exposing even more contradictory obligations based on diverging national implementations. Furthermore, it remains to be seen whether the use of behavioural advertising as a key element of 'free' cloud services should fall in the framework of the Unfair Commercial Practices Directive in countering misleading or aggressive advertising.
5. The breach notification rules present in recent revisions to the European Telecommunications Regulatory package (aimed at improving security by inciting providers of public electronic communications networks to avoid

data losses and subsequent notification of breaches) may also not be effective in achieving beneficial outcomes for security in the cloud. This is because cloud computing providers do not easily fall into the definition of an 'information society service provider' or a 'provider of public e-communications network'. The current breach notification regime only applies to the latter. Since a carefully and precisely drafted breach notification framework might be effective in contributing towards improved levels of security) the fact that this regime only covers providers of e-communication networks constitutes a challenge.

6. Similarly, the existing scope of the rules in the Data Retention Directive (specifying periods of retention for personal data on the grounds of necessity for law enforcement purposes) to communications providers may not be effective in achieving the desired security or criminal justice outcomes where traffic and logging data may be captured by cloud service providers.
7. The legal regime created by the Framework Decision on Attacks against Information Systems effectively requires Member States to implement common legal definitions of computer and network misuse and minimum sanctions. This may be undermined in a cloud computing environment where it is even more difficult to identify attacks emanating from the cloud, establish liability for harmful and illegal content (since there may be opaque chains of providers in a cloud computing context), identify applicable law (for the reasons regarding location indifference specified above), and establish international collaboration. This is because a cloud computing service may be dynamically and simultaneously composed of providers in a number of different jurisdictions.

Turning to instruments to support the implementation and governance of cloud computing and the management of associated security, privacy and trust challenges, the evidence gathered during the course of this project points to the following challenges:

8. There are a number of unique challenges common to those deploying or seeking to deploy cloud computing solutions. For example, the ability of security risk management programmes and frameworks (such as the UK government's Information Assurance Governance Framework or the BSi's 100-3 ITGrundschutz framework) to remain effective in providing a high level of security and contributing to the protection of personal data may be challenged in a cloud computing environment. This is largely due to the incompatibility of the administrative processes and governance mechanisms of such frameworks with the rapidity and scalability of cloud service provision. For example, the US FedRAMP (Federal Risk and Authorization Management Program) has been established to overcome this precise challenge: FedRAMP provides a standardised means of assessing and authorising cloud services and products, permitting an 'authorise once use many' approach. This is so that potential US government users can commission cloud services under a consistent risk management framework. Similarly, there may be inadequacies in established best practices for security management exposed by cloud computing. For example, establishing a

culture of information security within public administrations, or good data housekeeping practices could be undermined in a cloud context which offers the potential for limitless and indefinite storage of personal data (where it is not clear whether such measures are effective). As has been shown, other deficiencies exist with respect to security risk training and awareness raising campaigns aimed at creating positive behaviours and understanding toward security risks. Current means to establish transparency in respect of 'classical IT' systems (for example by audit regimes) might be similarly ineffective in a cloud computing environment due to the complexity of meeting audit requirements: for example, if an audit regime recommends physical site visits to data centres in order to provide assurance that risks to security, privacy or trust are being managed appropriately then, in a cloud computing context, it might not be economically possible to undertake such visits. Finally, current guidance or practice on managing continuity and resilience may not be tenable in a cloud environment. This applies on two levels: firstly in respect of resilience (given that an Internet connection is the single point of failure of many public cloud deployment models), but also regarding lock-in and the ability to extract and take data elsewhere should a provider no longer continue to be a viable business.

9. Turning to cloud providers, a number of aspects regarding the implementation of measures to provide for security, privacy or trust may not be effective in a cloud computing context. The most significant of these is regarding transparency and accountability. Because of the complexity of cloud service offerings, the levels of security and protection of personal data are not always transparent. Furthermore, the rapidity, scalability and accessibility of cloud computing also means that it is difficult for service providers to be held accountable (both in a general sense of governance but also in the more precise context of legal liability) for their implementation of security measures and how they go about meeting their responsibilities regarding privacy and trust. SLAs might be the main instrument to provide for accountability in such a scenario. However, smaller organisations using public cloud services might not have the bargaining power to negotiate specifically tailored SLAs. It might be more difficult to hold the cloud service provider accountable for a certain level of security or levels of protection of personal data. Other aspects relevant to accountability include provision of the mechanisms by which providers might be audited. Finally, cloud service providers are (understandably) reluctant to provide information or alert users to possible security issues. This can be for the reason that everyday consumers do not use security as a market differentiator and therefore it is inefficient for providers to invest in awareness raising measures.
10. At a more detailed level, as has been noted already, the current canon of best practice regarding security controls may not be effective in the governance of information security risks in a cloud context. Such guidance is often articulated in good practice guidelines such as the ISO2700x suite and the NIST standard. However, the guidance was developed prior to the concepts of outsourcing, off-shoring, virtualisation, IT service and multi-tenancy that

are key characteristics of cloud computing. The common factor across a number of these risk controls, such as physical access, vulnerability management, monitoring, identification and authentication, is the fact that (as with outsourcing) the effective implementation of these controls requires shared responsibility between the cloud service provider and the cloud user. The cloud user might not even be the ultimate end user, which may complicate the effective implementation of controls like 'creating a culture of information security'. To take a further example, in respect of the control of vulnerability management, in a cloud context the effective implementation of this control requires negotiation and shared responsibility between the cloud service provider and the cloud user to determine who will patch what machines, the patching and update tempo, and so on. Similarly, in respect of failure of components, services and so on, in a cloud context it would be difficult to establish thresholds that would trigger management escalation since the cloud service provider would have responsibility for the management of these components. Other risk controls are also affected. Providing for continuity and incident management in a public (or even hybrid) cloud environment is complex since the type of responses in an incident response plan might have cascading effects upon other tenants. Implementing controls related to residual data (data remanence) also represent a concern. Providing for a high level of confidence that all copies or shards of data have been deleted or moved is problematic in a cloud context since it is difficult to verify that no remnants of data remain physically remain.

11. Challenges also exist regarding guidance from regulators or developed by industry groups about certain aspects of security or privacy in the cloud. This includes absence of guidance on the expectations of data protection and privacy regulators regarding how organisations can meet their obligations in a cloud computing context. There is little guidance on how to conduct incident response or effective auditing in the cloud as well as how to go about forensics (eg, how to maintain evidential integrity of data in the cloud). Finally, although the ENISA cloud security report of late 2009 was the exception, much of the guidance is orientated towards either a US or generic context which either omits or does not adequately cover specifics of addressing security, privacy and trust issues from a European perspective.

Finally, this study identifies some specific technological challenges or areas where little or no R&D is being conducted. Such challenges are in respect of either scientific or technical domains or policy research:

12. A number of gaps in R&D research to overcome vulnerabilities or deficiencies in the underlying technological components of cloud computing have been identified in the literature review, interviews and expert workshop. These include further exploiting and building upon emergent cutting-edge research to develop methods to process data without having to decrypt; research into application programming interfaces (APIs) permitting the logging of events relevant for the measurement of security levels; research into

how to present security data to the cloud customer (eg, via real-time dashboards or security 'cockpits'); and technical research into user perceptions and behaviours in respect of privacy in the clouds (in order to inform privacy by design approaches). Current technical means to provide for identity management may not be viable in a cloud context but as yet there is little research into how to improve IDM. This is a challenge that will become even more acute in the public sector context given policy objectives surrounding the use of IDM across European public administrations. Finally, there are challenges relating to the extent of research being conducted around interoperability and the creation of technical standards which would facilitate the easy movement of data between cloud providers.

13. Turning to policy research (which would inform the creation of different types of guidance noted above), the main challenges are in respect of best practice for metrics – notably the investigation into designing appropriate independent third-party frameworks for cloud users to 'measure' the relative levels of security or privacy protection in the cloud market place. Similarly, there is little independent policy research into privacy guidance, particularly how cloud users and providers can meet their obligations under different privacy regimes to enable the effective use of the cloud whilst upholding European values.

The next chapter covers the recommendations as to how these identified gaps may be addressed by various forms of policy intervention or action by different types of stakeholder, including European institutions, standards bodies, Member States, industry, and end users.

CHAPTER 9 Solving the challenges: recommendations and actions

9.1 Introduction

Addressing the security and privacy challenges of cloud computing is a complex undertaking since it requires a combination of technological solutions and legal approaches that is capable of addressing operational realities and concerns. The extent of the technical, legal and operational challenges has been presented in the preceding chapters, and it is clear that an appropriate policy framework will be needed to address these challenges in a way that reconciles business drivers with public interests, while avoiding the pitfall of imposing burdens that put European businesses at a competitive disadvantage. However, the novelty of the subject has led to a situation where no such comprehensive public policy framework for framing these issues exists yet. Still, it has been possible to identify some initial attempts and useful inputs.

At an OECD-sponsored workshop in 2009, one of the key observations was that national governments have been assigned an important role.[151] The workshop concluded that the public sector should work together in fostering the development of standards and protocols, including by using their IT procurement power to push companies towards these common standards. The overall strategic goal is to avoid interoperability barriers and cases of technological lock-ins that may undermine the development of new cloud-based services. This is one of the ways in which governments could drive the resolution of technical challenges in a way that would also beneficially impact cloud use in the private sector.

However, from a legal perspective, the development of appropriate regulations should remain a major concern for governments. These need to be 'future proof' in the sense that they need to be able to take into consideration the development of new technological solutions.[152] This issue was flagged and extensively discussed during this study's own workshop, where much of the attention was devoted to the difficulty of complying with existing regulations in a cloud environment. European data protection regulations were discussed as the primary example of this issue, with the implication that they no longer fitted present-day economic realities or citizen interests. Specifically, the controller-

[151] OECD (2009)

[152] Ibid.

processor model was signalled as problematic, as it would lead to cloud service providers being classified as data controllers much more frequently than traditional outsourcing service providers, meaning that they would have to shoulder additional responsibilities and therefore be at a competitive disadvantage. Similarly, the importance that European data protection law accords to geographic location as a relevant criterion for determining data protection obligations (including applicable law and thus administrative formalities) was perceived as increasingly anachronistic and inefficient in an environment where national regulations should instead be converging. Thus, there was a strong call for improving international harmonisation as a key tool for facilitating compliance and thereby improving the effectiveness of data protection approaches. These elements should be taken into account in the ongoing review of European data protection regulation.

Not all of the effort should be exclusively shouldered by the public sector, however, as some of the legal, technical and operational challenges identified in the course of this study could also be addressed by or in collaboration with private industry. Security and privacy were discussed extensively during the aforementioned OECD workshop, and have been classified as being 'difficult' topics in light of the global and cross-jurisdictional nature of cloud computing. Aside from standards and protocols, the workshop strongly argued for a joint effort from policymakers and technologists to work together with industry in scoping the future landscape in order to identify innovative technical and policy solutions.

A similar multi-stakeholder approach has also been espoused by the International Telecommunication Union (ITU). However, the approach differs from the perspective taken by the OECD as the ITU indicates that the focus should be in devising appropriate technical standards that build upon similar activities undertaken in the context of grid computing or other utility computing environments. This means that all stakeholders are to be involved, starting from industry and government institutions. More importantly, it indirectly affirms that substantial work has been completed and the results of these efforts should provide the basis for addressing the public policy issues associated with cloud computing.[153]

It is important to emphasise that the multi-stakeholder approach has been espoused at the European level as well, including recently by ENISA in preparing its ENISA Cloud Computing Risk Assessment Report, addressing the security aspects of cloud computing. Its work has been primarily focused on devising technological and operational guidelines for addressing security and privacy concerns, which could be useful to assist stakeholders in ensuring their compliance with applicable rules and in adopting appropriate strategies to mitigate security and privacy threats in a cloud environment.[154] Similar reports and guideline documents have been prepared by other actors as well, including the aforementioned Cloud Security Alliance's Security Guidance for Critical Areas of Focus in Cloud Computing, and the Jericho Forum Self-Assessment Scheme. Thus, some of the building blocks for addressing cloud security and privacy challenges are already in place.

[153] ITU (2009)

[154] ENISA (2009)

Nonetheless, a more comprehensive approach is needed at the European level, where the interests and objectives of all stakeholders are reflected, and through which the gaps and challenges identified in the literature review and case studies can be addressed effectively. To this end, a series of recommendations are presented below.

9.2 Recommendations

Given the body of evidence presented in the literature and case studies in the preceding chapters, as well as the study's expert workshop, it is clear that most of the challenges in deploying or using cloud services revolve around three crucial issues:

- **Compliance**: ensuring that a cloud deployment meets the requirements imposed by the applicable normative framework, including general legislation, sector specific rules, and contractual obligations. The challenges in complying with data protection rules are a key example of this.
- **Accountability**: ensuring that security or privacy breaches in a cloud deployment are correctly addressed, including through appropriate compensation mechanisms towards any victims.
- **Transparency**: ensuring that the operation of a cloud deployment is sufficiently clear to all stakeholders, including service providers and users, both professional businesses and private consumers, and that its operation in practice can be assessed where necessary (including the identification of incidents). This can be witnessed, for example, in the difficulty of determining who/where a cloud service provider is, and where his responsibilities/liabilities end.

It is important to recognise that these three issues do not impact all stakeholders equally. Compliance will be mainly of interest to professional service providers and users, who are reasonably held to a higher standard than private consumers. Accountability is important for users (both professional and private) of cloud services, as it impacts their ability to rely on the cloud. Finally, transparency is important to all stakeholders, as it is the only basis on which they can build their trust in the cloud. Given the importance of these three pillars, the recommendations presented here should be able to provide useful improvements to each of these areas.

9.2.1 Compliance

A frequently heard complaint in interactions with stakeholders is the difficulty of ensuring that cloud usage complies with all applicable laws. Much of the problem stems from differences in national laws, and from the fact that some of these laws (with data protection being the most commonly quoted example) do not apply easily to the flexibility of the cloud model.

This problem can be addressed through the following recommendations:

- **In the longer run, a greater emphasis should be placed on international harmonisation**. A situation where applicable law is determined by

geographical criteria, and where applicable laws differ substantially, is only going to lead to either clear economic inefficiency (ie, having to assess compliance on a case by case basis), or large-scale non-compliance (ie, ignoring details that cannot be fitted into the existing model). Thus, a closer approximation of national laws is needed, initially at the European level, but eventually also at the global level. This applies to data protection law, but also to other areas such as consumer protection and international private law.

- **In the shorter run, clarification (eg, in the form of a policy paper) is needed on what cloud providers/users are expected to do.** Specifically with respect to data protection, guidelines at the European level were frequently requested explaining (i) when cloud providers are considered to be data controllers versus data processors, and (ii) how they can meet their obligations under European data protection law in a harmonised way. It is important to stress that these recommendations should be grounded in economic reality; an approach that merely requires that any cloud provider adapts its model to match any Member State regulations (which differ substantially and may even conflict) is unrealistic, especially given the dominance of non-European cloud service providers at this time. It may be advisable to link such clarifications to existing normative documents, such as the aforementioned ENISA work, and to sample policies/contracts, to give cloud service providers and users a more tangible goal to achieve.

9.2.2 Accountability

It is crucial to be able to ensure that security/privacy incidents can be appropriately investigated, and that suitable consequences (including compensation for victims) can be tied to such incidents. This is especially appropriate when considering the needs of private consumers, who have the least negotiating power (if any at all) when dealing with cloud service providers, and who will often have little practical means at their disposal to take legal action when such incidents occur.

This problem can be addressed through the following recommendations:

- **Cloud service providers should be encouraged to provide easier means for citizens to exercise their rights through cloud solutions**. Citizens should be able to easily verify which data relating to them are being held in the cloud, and should be able to export and/or delete such data if required, using automated tools. This would apply most clearly to citizens as direct users of cloud services. These rights are already contained to some extent in data protection regulations, and cloud service providers are in a unique position to enhance trust by making them more easily usable. That is not to say that addressing this will be easy since there is still further room for improvement in understanding behaviour and psychology of decision-making. Nonetheless in this way, consumers may have the possibility of taking their data elsewhere in case of incidents, which is an economically viable incentive for improving security/privacy.
- **Effective means for collective action against security/privacy breaches should be sought in future consumer and data protection policy**. If consumers are expected to take action in isolation, incidents are unlikely to

be sanctioned effectively. Collective actions (either in the form of a class action model or through representation by consumer and data protection bodies) are a viable option for improving accountability.

- **Research into standardised SLAs and liability provisions could lead to greater accountability**. SLAs and liability provisions are currently presented largely as a standardised approach by each service provider; however, they are usually not very comparable, making it harder for cloud users to determine their rights and possibilities of recourse. Research into standardisation of these provisions could help improve this situation to some extent.
- **Cloud users should be made aware of possible risks and appropriate behaviours**. It is important not to lay all responsibility at the feet of cloud service providers. Users too must be made aware of inappropriate risks and how to improve their own privacy and security. Usage of appropriate tools to manage this risk (including, for example, secure identification mechanisms) should be further encouraged.

9.2.3 Transparency

One of the major challenges related to cloud computing was the lack of transparency of cloud services: for any type of user (both professional and private), it is hard to determine how cloud services operate, and when incidents occur.

This problem can be addressed through the following recommendations:

- **As noted above, automated tools to exercise citizen rights can provide a first useful improvement.** Cloud service providers should be encouraged to provide such tools to improve transparency towards consumers.
- **Further research is recommended in the areas of security and performance metrics**. For both metrics and logging, several guidelines and methodologies are being developed, but there does not appear to be much consensus yet on best practices and common approaches to applying them. European guidance in this area could be useful to assist cloud users in assessing the effectiveness of cloud providers.
- **Incident response guidelines should be developed, followed by response disclosure obligations if appropriate and necessary**. A limited breach notification obligation already exists in respect of electronic communications in European law, but no similar obligation exists relating to information society service providers such as cloud service providers. This means that incidents may go by unnoticed, thus depriving cloud users of the possibility to take quick mitigating action. Cloud service providers should be at a minimum encouraged to inform their customers of any incidents which may impact them negatively, especially when there is still a possibility for these users to take mitigating actions. If ineffective, an obligation to inform cloud users in this way (preferably not targeted to cloud service providers specifically, but rather to information society service providers in general) may be considered.

9.2.4 Governance

By governance we mean top down measures that Europe can collectively take to support positive governance of cloud computing to make sure that security, privacy and trust is not undermined. This includes ensuring that the policy objectives and actions of the European Commission are well aligned with ongoing stakeholder activities, including by actively participating in the establishment and promotion of standards and best practices, and in interactions with cloud service providers.

- **Supporting ongoing promotion and implementation of best** practices tailored to effective management of operational risks arising from cloud deployments, for example the ENISA cloud guidelines but also supporting the evolution of current standards and best practices to better take account of risks arising from cloud deployments.
- **Acting as a leading customer** by using the e-Commission initiative in the Digital Agenda to roll out cloud computing deployments for a Commission service in order to explore the security, privacy and trust implications. An output of this would be a risk management framework for the use of cloud computing services such as is currently under consultation in the US FedRAMP programme.

9.3 Specific actions

Table 6, below, maps out and specifically articulates these recommendations against the gaps identified in Chapter 8 and according to the appropriate stakeholder best placed to address the gap.

Table 6. Concrete recommendations per stakeholder

Gap	Stakeholder	Action	Rationale
1–7	The European Commission, the Article 29 Data Protection Working Party and Member States	Work towards a greater emphasis on international harmonisation of relevant legal and normative frameworks via: • Further efforts to improve consistent application of relevant legal frameworks across the Member States through the conduct of implementation, monitoring and evaluations of relevant EU legislation • Continue to support broad international dialogue fostering harmonisation of relevant legal frameworks (especially regarding privacy) between, eg, the OECD, APEC and the US	Excessive differences (and even contradictions) between applicable national laws may, in a cloud context, lead to economic inefficiency and large scale non compliance
8, 10	The European Commission	Consider making use of cloud services itself, relying on a mixture of private and public cloud systems, and sharing its experiences and best practices with respect to security and privacy with other stakeholders. This could be done by identifying and selecting an appropriate pilot e-Commission service to be delivered via cloud services (eg, launching an European statistical data portal such as www.data.gov in the US)	Lead initiatives can show the viability of cloud solutions to the market, stimulate the European market for cloud services, and improve the awareness and viability of European security and privacy standards throughout the cloud community. The Digital Agenda also notes the possibilities of initiatives in the context of the e-Commission
1, 8, 11	Regulators in the Member States and cloud service providers	Develop suitable awareness raising mechanisms to help users to become aware of their own privacy and security risks by, for example: • Draft, prepare and issue guidance for cloud users (both organizations and individual consumers) on the benefits, risks and consequences of the storage and use of personal data in the cloud • Draft, prepare and issue guidance for cloud providers on how they should inform cloud users (especially consumers) of their rights in an accessible and understandable manner	Only by setting policies to establish joint responsibility will the risks be effectively and efficiently managed
8.4, 8.5	The European Commission and regulators in the Member States	Consider introducing/supporting means for collective action against privacy/security breaches in future consumer and data protection policy. This could be achieved via:	If consumers are expected to take action in isolation, incidents are unlikely to be sanctioned effectively. Collective actions (either in the form of

		• The incorporation of such rules in the proposed new European legal framework on privacy and data protection • The dissemination of best practice and common approaches amongst European privacy and data protection regulators in such fora as the London Initiative	a class action model or through representation by consumer and data protection bodies) are a viable option for improving accountability
11.1	The Article 29 Data Protection Working Party	Prepare guidelines (eg, in the form of a policy paper) on expectations of cloud users and obligations imposed on providers in the context of data protection • Prepare and release an Opinion on aspects of meeting data protection obligations in a cloud context relevant for data controllers and data processors	There exists a lack of clarity regarding when/how cloud providers and users should meet their obligations (eg, at what point cloud providers/users are considered to be data controllers, and how they are expected to comply with a multitude of national data protection laws)
11.2, 11.4	ENISA	Prepare incident response guidelines and response disclosure obligations, if appropriate and necessary • In conjunction with stakeholders such as cloud service providers and practitioners, prepare and release guidance covering how to respond to incidents targeted towards the three main categories of stakeholders (cloud providers, cloud users, and law enforcement). The cross-border nature of cloud computing and effective identification of competent authorities should take priority in this guidance document • Consider the preparation of guidance helping law enforcement and cloud providers and users understand the complexities of forensics in a cloud computing environment	Especially in an international cloud context, cloud stakeholders need to be able to respond effectively to incidents, including by communicating quickly and efficiently with competent bodies and possible victims if this can help mitigate/manage possible damages
1.4, 11.1	The European Commission	Explore if and how cloud users should be made aware of incidents in the cloud that could negatively impact them through consideration of the creation and inclusion of appropriate generalised and breach notification rules in the future legal European framework regarding privacy and data protection	Although a limited breach notification law exists with respect to electronic communications, no such equivalent exists for cloud service providers who do not qualify as electronic communications service providers
10	The European Committee for Standardization (CEN) and ENISA	Actively participate in the establishment and promotion of security standards and best practices, and should interact with cloud service providers to foster a culture of security and privacy awareness • Consider the extent to which current globally accepted standards in	In the shorter term, non-regulatory approaches and direct interaction with the cloud industry may be a more effective means to improve security and privacy in the cloud at the international level, without creating disincentives and competitive

		the realm of security and privacy (such as the ISO2700x Information Security Management System or NIST 800-series standards) might be updated to be more suitable for a cloud environment through stakeholder engagement (eg, via a conference or workshop) • Explore the necessity of new cloud security standards, keeping in mind their potential future use for the purpose of obtaining security certification (eg, as a guideline to obtain ISO/IEC 2700x compliance accreditation) • Explore and support the development of internationally acceptable standards for interoperability (eg, through dialogue with ISO and industry groups such as the Cloud Security Alliance and OpenCloud)	disadvantages for European cloud providers
9	Industry: cloud service providers	Work to make it easier for cloud users and citizens to exercise their rights when using cloud solutions • Be encouraged to participate fully in interoperability and standardisation processes and initiatives to permit users to more easily move their data between providers • Investigate, develop and implement simple, transparent, automated and effective solutions to support the exercise of privacy and data protection rights in a cloud service, permitting data to be easily accessed, erased, moved and rectified, for example	Clear means for users to migrate their data to other providers will act as a incentive for cloud service providers to increase their levels of security, thereby creating economic drivers for security and privacy. Furthermore, providing simple tools to exercise one's rights as a data subject can be a useful first improvement in improving transparency towards consumers, which can be instrumental for improving trust
9, 11, 13.1	Academia, the European Commission and ENISA	Investigate how policy and practice diverge regarding demand-side factors concerning security, privacy and trust • Conduct policy research into best practice and common approaches to applying security metrics in the cloud • Investigate the creation of a generic maturity model to independently ex-ante measure and assess cloud security provision	No consensus yet exists on best practice and common approaches to applying security metrics
13.1	Academia, the European Commission and ENISA	Undertake policy and legislative research into standardised SLAs/liability provisions for cloud computing Undertake policy research to identify common and useful liability provisions and templates for SMEs that serve to uphold a high level of security and privacy for cloud users and European citizens	The isolated and largely non-negotiable approach currently taken by cloud providers is not very accessible nor beneficial to end users (especially SMEs and consumers, who have little negotiating power), resulting in difficulties for cloud users to determine or exercise their rights and possibilities

			of recourse
12	Academia	Attract European research & development funding (through the scientific European 7th Framework Programme) to: • Conduct research to establish effective locations and types of security detection mechanisms (sensors) in cloud architectures to support rapid detection of emergent threats from the cloud • Investigate and pilot secure virtualised architectures and trust domains for cloud computing environments (virtual machines running on specially configured hypervisors according to a set of common security domain of controls) • Investigate tools and instruments to support shared responsibility of security controls contained in best practice and guidance such as the ISO2700x suite (eg, configuration of roles and responsibilities for execution of patch management activities between the cloud provider and cloud user) • Design testbeds and pilots for middleware, software interfaces and APIs for Security Event and Incident Management (SEIM) to permit interrogation and collation of all events that might be of interest in measuring security (eg, CPU failure, hardware failure, etc.) as well as establishing alert thresholds. Further explore the design of SEIM 'dashboards' for cloud computing across different use cases (SMEs, consumers) • Explore and investigate appropriate means for interoperability of both data exchange (where possible) and enforcement of security, privacy and business policies attached to data across distributed cloud environments. • Research approaches to implementing interoperable abstraction of security, privacy and trust policies to enable data 'policy stickiness' in untrustworthy cloud environments • Research technical measures to support transparency in security for the SME/private end user of cloud (eg, using SEIM to provide for accessible assurance and compliance checking when the cloud user moves across different trust boundaries in the cloud) • Research frameworks and common specifications for handling identity access and identity assurance in the cloud (eg, mapping of	Greater awareness and transparency of security risks and incidents is needed, and international standardization work is a good inroad to ensuring that these are promoted at the international level, especially for higher risk applications where audits using such standards are likely to be used

		identities across cloud service providers, designing suitable Trusted Third Party mechanisms to establish federated identity management for cloud computing and dealing with the identity lifecycle in the cloud) • Further research into applicability of technical means to provide for increased levels of data security across trust domains (eg, automated data expiry mechanisms and secure movement or deletion of data) • Chart flows of security meta information in cloud computing environments to enable dynamic, verifiable and attestable trustmarks and certificates	

CHAPTER 10 Conclusions

This research study has shown that the emergent landscape for cloud computing is characterised by a number of challenges for security, privacy and trust. In many respects these challenges are not new (some exhibiting similarities with outsourcing), but cloud computing models brings them into sharper focus. These challenges exist in relation to the use of specific technologies in the cloud context as well as existent or latent vulnerabilities in the technological building blocks comprising different cloud computing deployments. There are also important questions in respect of the pertinent legal and regulatory domains as applied to cloud computing, most notably relating to legal obligations stemming from location of (personal) data in the cloud, accountability, transparency, consent, security and definition of responsibilities of those using and processing personal data. From an operational perspective, some of the common tools available to organisations to help manage security, privacy and trust challenges arising in cloud computing deployments may also be less viable or effective. This problem arises from the complexity of establishing and monitoring risk controls across the technological, organisational and legal boundary between the cloud service provider and the cloud user. The case studies illustrated that cloud computing deployments remain relatively new, that consideration of security and privacy may be an afterthought or treated separately and that it is difficult (but not impossible) to negotiate specific security terms with cloud service providers.

The study has identified a number of challenges for current EU policy relating to, for example, the legal and regulatory framework regarding privacy and data protection; electronic communications privacy law; cyber-crime law; copyright but also horizontal legal issues including establishing jurisdiction and handling disputes. There are also opportunities for enhancement of the way public and private sectors implement measures to provide for privacy, security and trust in cloud computing environments. These most notably exist in respect of implementing risk control frameworks, best practice guidance and the availability of guidance from regulators to cloud service providers and users on meeting their legal obligations. Other issues include the challenge posed by cloud computing to establishing appropriate accountability, transparency, governance and compliance. Finally, there are opportunities for R&D in the areas of security monitoring, identity management and interoperability of privacy and security policies in un-trusted cloud contexts.

As with many other public policy questions in respect of security and privacy, successfully addressing these challenges should not be considered a static goal but rather a constant process. Furthermore, this process will necessarily involve the active, responsible and

engaged participation of all stakeholders: governments; public policymakers at both European and national levels; industry; citizens and consumers. Only such a multi-stakeholder approach will ensure that the economic benefits of cloud computing can be secured whilst providing for a high level of security and privacy which will contribute to increased trust in cloud computing.

REFERENCES

List of references

Aguilar Melchor, Carlos, Philippe Gaborit & Javier Herranz, *Additively Homomorphic Encryption with d-Operand Multiplications*, Cryptology ePrint Archive, 2008/378.

Amazon Web Services, 'Cost Comparison Example for High-Performance Computing', December 2009. As of 25 November 2010 available at: http://aws.amazon.com/economics/

Amazon Web Services, 'Overview of Security Processes', August 2010. As of November 2010: http://awsmedia.s3.amazonaws.com/pdf/AWS_Security_Whitepaper.pdf

Armbrust, Michael et al., 'Above the Clouds: A Berkeley View of Cloud Computing', University of California at Berkeley, Technical Report No. UCB/EECS-2009-28, 10 February 2009. As of 25 November 2010: http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf

Armstrong, Mark, 'Competition in Two-sided Markets', *The RAND Journal of Economics*, 2006/37(3):668–691.

Article 29 Data Protection Working Party, 'Working document on determining the international application of EU data protection law to personal data processing on the Internet by non-EU based web sites', WP56, 2002. As of 25 November 2010: http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2002/wp56_en.pdf

Article 29 Data Protection Working Party & Working Party on Police and Justice, 'The Future of Privacy: Joint contribution to the Consultation of the European Commission on the legal framework for the fundamental right to protection of personal data', WP168, 2009. As of 25 November 2010: http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2009/wp168_en.pdf

Beco, Stefano, 'Resources and Services Virtualization without Barriers: Cloud Computing and RESERVOIR project', Powerpoint presentation, 2008. As of 25 November 2010: https://www.research.ibm.com/haifa/projects/systech/reservoir/present/RESERVOIR_CSFI08.ppt

benchai7, 'Cloudburst – A Weaponized Attack on the Cloud', posting on Security Vibes, 5 June 2009. As of 25 November 2010: http://securityvibes.com/community/en/blog/2009/06/05/cloudburst--a-weaponised-attack-on-the-cloud

Birman, Ken, Gregory Chockler & Robbert van Renesse, 'Toward a Cloud Computing Research Agenda', *ACM SIGACT News*, 2009/40(2).

Bradshaw, Simon, Christopher Millard & Ian Walden, 'Contracts for Clouds: Comparison and Analysis of the Terms and Conditions of Cloud Computing Services', Queen Mary School of Law Legal Studies Research Paper No. 63, 2010. As of 25 November 2010: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1662374

Brookings Institution, *Saving Money Through Cloud* Computing, Washington DC, 2010.

Buyya, Rajkumar, 'Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility', *Future Generation Computer Systems*, 2008/25(6).

Cameron, Kim, 'Identity software and services roadmap,' Microsoft Professional Developers' Conference, 2008.

Cameron, Kim, 'Software + Services Identity Roadmap', Microsoft Professional Developers' Conference, 2009.

Casassa Mont, Marco, Siani Pearson & Pete Bramhall, 'Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services. In: *IEEE Workshop on Data and Expert Systems Applications*, Washington, DC: IEEE Computer Society Press, pp. 377–382, 2003.

Centre for the Protection of National Infrastructure, *Cloud Security Risks*, London, 2010.

Chow, Richard et al. 'Controlling data in the cloud: outsourcing computation without outsourcing control'. In *Proceedings of the 2009 ACM Workshop on Cloud Computing Security*, pp.85–90. New York: ACM, 2009. As of 25 November 2010: http://doi.acm.org/10.1145/1655008.1655020

Ciurana, Eugene, *Developing with Google App Engine*, New York: Springer-Verlag, 2008.

Cloud Computing Use Case Discussion Group 'Cloud Computing Use Cases White Paper' version 4.0 As of 22 November 2010: http://www.scribd.com/doc/18172802/Cloud-Computing-Use-Cases-Whitepaper

Cloud Security Alliance, 'Security Guidance for Critical Areas of Focus in Cloud Computing v2.1', 2009. As of 21 November 2010: http://www.cloudsecurityalliance.org/csaguide.pdf

Couillard, David A., 'Defogging the Cloud: Applying Fourth Amendment Principles to Evolving Privacy Expectations in Cloud Computing', *Minnesota Law* Review, June 2009.

Cramer, R., I. Damgård & J. B. Nielsen, 'Multiparty Computations from Threshold Homomorphic Encryption'. In *Proceedings of the international Conference on the theory and Application of Cryptographic Techniques: Advances in Cryptology (May 06–10, 2001)*, ed. B. Pfitzmann, Lecture Notes In Computer Science, vol. 2045, pp.280–299. London: Springer-Verlag, London, 2002.

Creese, Sadie, Mike Auty, Michael Goldsmith & Paul Hopkins, *Inadequacies of Current Risk Controls for the Cloud*, forthcoming (to appear in the proceedings of the CloudCom 2010 conference).

Damgård, I., M. Geisler & M. Kroigard, 'Homomorphic encryption and secure comparison', *Int. J. Appl. Cryptol.* 2008/1(1):22–31. As of 25 November 2010: http://dx.doi.org/10.1504/IJACT.2008.017048

Damgård, I., M. Geisler & M. Kroigard, 'A correction to "efficient and secure comparison for on-line auctions"', *Int. J. Appl. Cryptol.* 2009/1(4):323–324. As of November 2010: http://dx.doi.org/10.1504/IJACT.2009.028031

D'Andria, F. et al., 'An Enhanced Strategy for SLA Management in the Business Context of New Mobile Dynamic VO'. In *Exploiting the Knowledge Economy: Issues, Applications, Case Studies*, eds. Paul Cunningham and Miriam Cunningham, Amsterdam: IOS Press, 2006. As of 25 November 2010: http://www.it-tude.com/slamanagement_businesscontext.html

Dushin, Fred & Eric Newcomer, 'Handling Multiple Credentials in a Heterogeneous SOA Environment', *IEEE Security and Privacy* 2007/5(5):80–82.

European Network and Information Security Agency (ENISA), 'Position Paper: Web 2.0 Security and Privacy', December 2008.

European Network and Information Security Agency (ENISA), *Cloud Computing: Benefits, risks and recommendations for information security*, 2009a. As of 25 November 2010: http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport

European Network and Information Security Agency (ENISA) *Cloud Computing: Information Assurance Framework*, 2009b. As of 25 November 2010: http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-information-assurance-framework/at_download/fullReport

Etro, Federico, 'The Economic Impact of Cloud Computing on Business Creation, Employment and Output in Europe, *Review of Business and Economics*, 2009/2:180–208.

European Commission, *A Digital Agenda for Europe*, Brussels, 2010a. As of 20 November 2010: http://ec.europa.eu/information_society/digital-agenda/documents/digital-agenda-communication-en.pdf

European Commission, Directorate General for Information Society and Media, *Expert Group Report Meeting on Cloud Computing*, Brussels, 2010b.

Fontaine, Caroline & Fabien Galand, 'A survey of homomorphic encryption for nonspecialists', *EURASIP J. Inf. Secur.* 2007:1–15. As of 25 November 2010: http://dx.doi.org/10.1155/2007/13801

Garfinkel, Tal & Mendel Rosenblum, 'When Virtual is Harder than Real: Security Challenges in Virtual Machine Based Computing Environments'. In *Proceedings of the 10th Workshop on Hot Topics in Operating Systems (HotOS-X)*, 2005. As of 25 November 2010: http://www.stanford.edu/~talg/papers/HOTOS05/virtual-harder-hotos05.pdf

Gartner, *Building Blocks for Trusted, Secure Hypervisors* (G00147722), 2007.

Gartner, *Assessing the Security Risks of Cloud Computing* Gartner, 2008a

Gartner, *Location Matters*, 2008b.

Gartner, *Cloud Computing Will Cement the Mainstream Role of SOA* (G00167246), 2009a.

Gartner, *Top Five Cloud Computing Adoption Inhibitors*, 2009b.

Gentry, Craig, 'Fully homomorphic encryption using ideal lattices'. In *Proceedings of the 41st Annual ACM Symposium on theory of Computing*, pp.169–178. New York: ACM, 2009. As of November 2010: http://doi.acm.org/10.1145/1536414.1536440

Gervais, Daniel, 'The Tangled Web of UGC: Making Copyright Sense of User-Generated Content', *Vanderbilt Journal of Entertainment and Technology Law*, Summer, 2009.

Hasselmeyer, Peer et al., 'Negotiating SLAs with Dynamic Pricing Policies', Service Oriented Computing: a look at the Inside (SOC@Inside'07), 2007. As of 25 November 2010: http://www.nextgrid.org/download/publications/Negotiating_SLAs_with_Dynamic_Pricing_Policies.pdf

Hewitt, Carl, 'ORGs for Scalable, Robust, Privacy-Friendly Client Cloud Computing', *IEEE Internet Computing*, 2008/12(5):96–99.

Hogben, Giles, 'Privacy, Security and Identity in the Cloud', presentation, ENISA, 2009. As of 25 November 2010:
http://www.enisa.europa.eu/act/it/oar/cloud-computing/Cloud_Identity_Hogben.pdf/at_download/file

ISF, *The Standard of Good Practice for Information Security*, Information Security Forum, 2005

ITU, 'Distributed Computing: Utilities, Grids & Clouds', ITU-T Technology Watch Report 9, 2009. As of 25 November 2010 available at:
http://www.itu.org

Kerschbaum, F. & R. J. Deitos, 'Security against the business partner'. In *Proceedings of the 2008 ACM Workshop on Secure Web Services*, pp.1–10, New York: ACM, 2008. As of 25 November 2010:
http://doi.acm.org/10.1145/1456492.1456493

King, S.T. et al., 'SubVirt: Implementing malware with virtual Machines', *Proc. 2006 IEEE Symp. Security and Privacy (SP 06),* pp.314–327, IEEE CS Press, 2006. As of 25 November 2010:
http://www.eecs.umich.edu/virtual/papers/king06.pdf

Kirkpatrick, Marshall, 'Amazon releases early info on S3 storage use', TechCrunch, 12 July 2006. As of 20 November 2010:
http://www.techcrunch.com/tag/s3

Kiyanclar, Nadir, 'A Survey of Virtualization Techniques Focusing on Secure On-Demand Cluster Computing', University of Illinois at Urbana-Champaign, National Center for Supercomputing Applications, 2 November 2005.

Laplante, P., J. Zhang & J. Voas. 'What's in a name? Distinguishing between saas and soa'. *IT Professional* 2008/10(3):46–50.

Lee, Edward, 'Warming up to user generated content', *University of Illinois Law Review*, 2008.

Luo, Y., S.S. Cheung, & S. Ye, 'Anonymous biometric access control based on homomorphic encryption'. In *Proceedings of the 2009 IEEE international Conference on Multimedia and Expo*, pp.1046–1049, Institute of Electrical and Electronics Engineers, 2009.

Mather, Tim, Subra Kumaraswarmy & Shahed Latif, *Cloud Security and Privacy*, Sebastopol, CA: O'Reilly, 2009.

Mell, Peter and Tim Grance, 'The NIST Definition of Cloud Computing', National Institute of Standards and Technology, 2009a. As of 25 November 2010: http://www.*csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc*

Mell, Peter & Tim Grance, 'Effectively and Securely Using the Cloud Computing Paradigm'. Presentation, NIST Information Technology Division, 2009b.

Mills, Laurin H., 'Legal Issues Associated with Cloud Computing', presentation, May 2009. As of 25 November 2010: http://www.secureit.com/resources/Cloud%20Computing%20Mills%20Nixon%20Peabody%205-09.pdf

Molina, A. D., M. Salajegheh, & K. Fu, 'HICCUPS: health information collaborative collection using privacy and security'. In *Proceedings of the First ACM Workshop on Security and Privacy in Medical and Home-Care Systems*, pp.21–30. New York: ACM, 2009. As of 25 November 2010: http://doi.acm.org/10.1145/1655084.1655089

Monahan, Brian & Mike Yearworth, 'Meaningful Security SLAs', HP Laboratories Technical Report HPL-2005-218R1, 2008.

Murley, Diane, 'Technology for Everyone…: Law Libraries in the Cloud', *Law Library Journal*, Spring 2009.

Murty, James, *Programming Amazon Web Services: S3, EC2, SQS, FPS, and SimpleDB*, Sebastopol, CA: O'Reilly, 2008.

Nature 'web matters', 2000. As of 25 November 2010: http://www.nature.com/nature/webmatters/grid/grid.html

Nelson, Michael R., 'The Cloud, the Crowd, and Public Policy', *Issues in Science and Technology*, 2009/Summer:71–76. As of 25 November 2010: http://cct.georgetown.edu/Nelson%20Cloud%20Article.pdf

NIST, 'Recommended Security Controls for Information Systems', Special Publication 800-31-1, 2005, National Institute of Standards and Technology.

Nurmi et al., 'The Eucalyptus Open-Source Cloud Computing System', *J. Phys.: Conf. Ser.* 2009/180:012051. As of 25 November 2010: http://open.eucalyptus.com/documents/ccgrid2009.pdf

OECD, 'Briefing Paper for the ICCP Technology Foresight Forum: Cloud Computing and Public Policy', DSTI/ICCP(2009)17, 2009.

O'Malley, Owen, 'Hadoop World: Security and API Compatibility', presentation, 2009. As of 25 November 2010:

http://www.cloudera.com/blog/2009/12/08/hadoop-world-security-and-api-compatibility/

Ormandy, T., *An Empirical Study into the Security Exposure to Hosts of Hostile Virtualized Environments*', Google, Inc, 2007. As of 25 November 2010: http://taviso.decsystem.org/virtsec.pdf

Ostrovsky, R & William E. Skeith III, 'Algebraic Lower Bounds for Computing on Encrypted Data', Cryptology ePrint Archive, 2007/064.

Pearson, Siani, Yun Shen & Miranda Mowbray, *A Privacy Manager for Cloud Computing*, HP Laboratories HPL-2009-156, CloudCom 2009, Beijing, Springer LNCS, 2009.

Perez, R., L. Van Doorn & R. Sailer, 'Virtualisation and Hardware Based Security', *IEEE Security & Privacy*, 2008/6(5):24–31.

Picker, Randal C., *Competition and Privacy in Web 2.0 and the Cloud*, University of Chicago, Olin Working Paper No. 414, 2008. As of 25 November 2010: http://ssrn.com/abstract=1151985

Prabhakaran, M. & M. Rosulek, 'Homomorphic Encryption with CCA Security'. *Lecture Notes In Computer Science*, 2008/5126:667–678. As of 25 November 2010: http://dx.doi.org/10.1007/978-3-540-70583-3_54

Ramasamy, H.V. & M. Schunter, 'Architecting Dependable Systems Using Virtualization', IBM Zurich Research Laboratory, Ruschlikon, Switzerland, 2007.

Rayport, Jeffrey F. & Andrew Heyward, 'Envisioning the Cloud: The Next Computing Paradigm', *Marketspace Point of View*, 20 March 2009.

Ristenpart, Thomas, Eran Tromer, Hovav Shacham & Stefan Savage, 'Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds'. In *Proceedings of CCS 2009*, pp.199–212. New York: ACM Press, 2009.

Rochet, Jean-Charles & Jean Tirole, 'Platform Competition in Two-sided Markets', *Journal of the European Economic Association*, 2003/1(4):990–1029.

Rochwerger, B. et al., The RESERVOIR Model and Architecture for Open Federated Cloud Computing. IBM Journal of Research & Development, 2009/53(4).

Rosenblum, Mendel, & Tal Garfinkel, 'Virtual Machine Monitors: Current Technology and Future Trends', *IEEE Computer*, May 2005.

Sailer, R. et al., 'sHype: Secure Hypervisor approach towards trusted Virtualisation', *IBM Technical Research Report*, February 2005.

Sang, Y. and H. Shen, 'Efficient and secure protocols for privacy-preserving set operations'. *ACM Trans. Inf. Syst. Secur.* 2009/13(1):1–35. As of 25 November 2010: http://doi.acm.org/10.1145/1609956.1609965

Schwartz, Paul M., 'From Victorian Secrets to Cyberspace Shaming', *University of Chicago Law Review*, Summer 2009.

Scolnik, Alexander, 'Protections for electronic communications: the Stored Communications Act and the Fourth Amendment', *Fordham Law Review*, October 2009.

Scope Alliance, 'Virtualization: State of the Art', Version 1.0, 3 April 2008. As of 25 November 2010: http://www.scope-alliance.org/sites/default/files/documents/SCOPE-Virtualization-StateofTheArt-Version-1.0.pdf

SensePost, 'Clobbering the Cloud', SensePost Blog, DEF CON 17 Conference, August 2009. As of 25 November 2010: http://www.sensepost.com/blog/3706.html

Sharif, Monirul et al. Secure In-VM Monitoring Using Hardware Virtualization. In *Proceedings of the 16th ACM conference on Computer and communications security CCS 09*, 477. New York: ACM, 2009. As of 25 November 2010: http://research.microsoft.com/en-us/um/people/wdcui/papers/sim-ccs09.pdf

Simmendinger, Christian, 'Whitepaper: Support for Client-Server based License Management Schemes in the Grid', IT-Tude.com, 2008. As of 25 November 2010: http://www.it-tude.com/fileadmin/gridipedia/lm_cluster/LM_WHITEPAPER_V0.1-1.pdf

Smart, N.P. & F. Vercauteren, 'Fully Homomorphic Encryption with Relatively Small Key and Ciphertext Sizes', Cryptology ePrint Archive, 2009/571.

Vaidya, J. et al., 'Privacy-preserving decision trees over vertically partitioned data,. *ACM Trans. Knowl. Discov. Data* 2008/2(3):1–27. As of 25 November 2010: http://doi.acm.org/10.1145/1409620.1409624

Van Cleeff, André, Wolter Pieters & Roel Wieringa, 'Security Implications of Virtualization: A Literature Study', cse, 2009/3:353–358, International Conference on Computational Science and Engineering, 2009.

Vaquero, L.M. et al. 'A Break in the Clouds: Toward a Cloud Definition', *ACM SIGCOMM Computer Communication Review*, 2009/39(1):50–55.

VMware Inc, *Technical White Paper*. Palo Alto, 1999.

Vogels, Werner, 'Eventually Consistent: Building reliable distributed systems at a worldwide scale demands trade-offs between consistency and availability', All Things Distributed, December 2008. As of November 2010: http://www.allthingsdistributed.com/2007/12/eventually_consistent.html

Vouk, M.A., 'Cloud computing: Issues, research and implementations', *Journal of Computing and Information Technology*, 2008/16(4):235–246.

Werbach, Kevin, 'The Centripetal Network: How the Internet Holds Itself Together, and the Forces Tearing It Apart', *UC Davis Law Review*, December 2008.

Wojtczuk, Rafal, Joanna Rutkowska & Alexander Tereshkin, 'Another Way to Circumvent Intel® Trusted Execution Technology', Invisible Things Lab, 2009. As of 25 November 2010: http://invisiblethingslab.com/resources/misc09/Another%20TXT%20Attack.pdf

World Economic Forum, *Exploring the Future of Cloud Computing: Riding the Next Wave of Technology Driven Transformation* (WEF 2010). As of 22 November: http://www3.weforum.org/docs/WEF_ITTC_FutureCloudComputing_Report_2010.pdf

Yokoo, M. & K. Suzuki, 'Secure multi-agent dynamic programming based on homomorphic encryption and its application to combinatorial auctions'. In *Proceedings of the First international Joint Conference on Autonomous Agents and Multiagent Systems: Part 1 (Bologna, Italy, July 15–19, 2002)*, pp.112–119. New York, NY, 112-119. DOI=http://doi.acm.org/10.1145/544741.544770: ACM, 2002. As of 25 November 2010: http://doi.acm.org/10.1145/544741.544770

Young, P., 'The Evolution of Conventions', *Econometrica* 1993/61(1):57–84.

APPENDICES

Appendix 1: List of interviewees

The following persons were interviewed in the course of the SecureCloud 2010 conference:

- Wendy Goucher, Idrach
- Randy Barr, Qualys
- Marcus Lasance, Verizon
- Tim Grance, NIST
- Gorka Sadowski, LogLogic
- Bernd Grobauer and Thomas Schreck, Siemens

Appendix 2: Cloud computing case study framework

Profile of the use case owner

A summary description of the use case owner (ie, the company, public service or entity adopting a cloud service) should be provided, including its core activities.

Include contact information of whomever provided the information.

Description of the use case

Describe what data/service was moved to the cloud.

What were the main objectives to be achieved with a cloud computing implementation?

Chosen cloud solution

Describe the chosen solution. Indicate the type of cloud service, and explain why it falls into this category (if not obvious):

- Private or public clouds
- Commercial / public sector / scientific / community-driven cloud
- IaaS / PaaS / SaaS / HaaS.

Risks and challenges of cloud use

Is the use case:

- A low-risk use case, where data loss or service interruption has no serious impact on the user's activities. No crucial data or services are entrusted to the cloud, and the risk of significant damage is limited.
- A medium-risk use case, where data loss or service interruption could have a noticeable impact on the user's activities, and could result in real damages without being business critical.
- A high-risk use case, where data loss or service interruption would have a detrimental impact on the user's activities, and could result in significant damages.

Explain the risks involved, and why they were assessed as low/medium/high.

Was there a systematic risk assessment prior to the switch? What were the main identified problems and concerns?

Chosen approach to address these risks

How were the main security/privacy and trust issues ultimately addressed:

- From a legal perspective, eg, via:
  - Specific contracts, including service level agreements, maintenance agreements, auditing rights, etc.
  - Compliance controls
  - External legal audits

- From a technical perspective, eg, via:
  - Establishing technical guidelines such as incident handling protocols, logging obligations, security metrics, etc.
  - Testing periods and/or trial phases during implementation
  - External technical audits

- From an operational perspective, eg, via:
  - Migration planning, including staff training
  - Risk Assessment, including lock-in, incident management planning, etc.
  - Implementing contingency and/or continuity plans
  - External operational audits

Experiences and lessons learned

What was your experience in switching to a cloud service? What were the main lessons learned?

Having switched, have the anticipated benefits materialised? What have been the main advantages and disadvantages?

Remaining issues and concerns

Are you confident that all relevant issues have been identified and addressed by your current solution?

Please identify any areas in which you are less confident at this point, or where you see room for improvement.

Appendix 3: Workshop Agenda

Background

The European Commission has asked a study team led by RAND Europe with time.lex and International Digital Lab (IDL) of the University of Warwick to explore the security, privacy and trust aspects of cloud computing. This study has undertaken a review of the literature and a number of real-life case studies to identify how challenges in respect of the privacy, security and trust issues were overcome in various implementations of cloud computing.

Purpose of the workshop

The overall purpose of this workshop is to discuss the findings and test the identified gaps in respect of EU policy regarding Cloud Computing, building on the approaches taken and challenges encountered by real life case studies. Following the workshop, this prioritised list of gaps will then be used to identify and propose policy actions and concrete recommendations for the European, national government and industry stakeholders. This will be integrated into the final report of the study, which will be made available to participants in late 2010. The workshop will very much be an interactive event, where, following an introduction to the study and discussion of the key issues involved in this, participants will be asked to engage in an interactive discussion regarding the various broad issues identified. This will allow the participants to learn from the study findings, but also to inject their own experiences from implementations of cloud computing.

The event will be conducted under the Chatham House rule (non-attribution) and a sandwich lunch will be provided.

Why you should attend

By participating you will:

- **gain access to unique research findings** – this study is one of the major EU level efforts to analyse the security, privacy and trust implications of cloud computing;
- **obtain data** to support your own business case for the use of cloud computing and to improve awareness of risks and opportunities;
- **have an opportunity to network** and strengthen your informal relationships with peers, solution providers, practitioners and European experts from both the public and private sector;

- **have an opportunity to participate** in the preparatory activities prior to the discussion and formulation of EU policy in this domain.

Agenda

Times and speakers may be subject to change

Topic	Speaker	Time
Registration		09.30
Introduction and rationale for the study	**Gustav Kalbe** Deputy Head of Unit F.5, DG Information Society & Media	10.00
Introduction to the project and agenda: Background to the study, explanation of findings so far and purpose of the day	**Neil Robinson**, Study Team Leader, RAND Europe	10.20
Summary and Discussion on identified technical, legal and policy issues during the study	• **Jonathan Cave,** RAND Europe (tbc) • **Hans Graux**, time.lex • **Prof Sadie Creese,** University of Warwick	10.30
Morning Coffee		11.00
Session 1: Reality check: Interactive session to discuss the operational relevance and impact of the issues identified above	Interactive discussion involving all participants moderated by **Dr Lorenzo Valeri**, Scientific Manager, LUISS School of Government, Rome	11.30
Lunch		13.00
Models of deployment	Panel session	14.00
Session 2: Security, privacy and trust aspects across public and private sectors: Public, private and research perspectives	Debate: • Regulator – **Iain Bourne**, Information Commissioner's Office • Government – **Leon-Paul Rouw**, Ministry of Interior and Kingdom Relations, Netherlands (tbc) • Industry - **Marco Casassa Mont,** HP Labs	
Afternoon tea		15.30
Session 3: Policy options & objectives going forward: Discussion and validation of possible policy recommendations to address the implications	Interactive discussion involving all participants moderated by **Dr Lorenzo Valeri**, Scientific Manager, LUISS School of Government, Rome	16.00
Wrap up & Concluding remarks	**Neil Robinson**, RAND Europe & **Valeria Bricola**, Project Officer Unit F.5, DG Information Society and Media	16.30
Close		16.45