



ABOUT ENISA

The European Network and Information Security Agency (ENISA) is an EU agency created to advance the functioning of the internal market. ENISA is a centre of excellence for the European Member States and European institutions in network and information security, giving advice and recommendations and acting as a switchboard for information on good practices. Moreover, the agency facilitates contacts between European institutions, the Member States, and private business and industry actors.

This work takes place in the context of ENISA's Emerging and Future Risk programme.

CONTACT DETAILS:

This report has been edited by:

e-mail: Daniele.catteddu@enisa.europa.eu and Giles.hogben@enisa.europa.eu,

Internet: <http://www.enisa.europa.eu/>

Legal notice

Notice must be taken that this publication represents the views and interpretations of the editors, unless stated otherwise. This publication should not be construed to be an action of ENISA or the ENISA bodies unless adopted pursuant to ENISA Regulation (EC) No 460/2004. This publication does not necessarily represent the state-of-the-art in cloud computing and it may be updated from time to time.

Third-party sources are quoted as appropriate. ENISA is not responsible for the content of the external sources including external websites referenced in this publication.

This publication is intended for educational and information purposes only. Neither ENISA nor any person acting on its behalf is responsible for the use that might be made of the information contained in this publication.

Reproduction is authorised provided the source is acknowledged.

© European Network and Information Security Agency (ENISA), 2009

LIST OF CONTRIBUTORS

This paper was produced by ENISA editors using input and comments from a group selected for their expertise in the subject area, including industry, academic and government experts.

The views expressed in this publication are those of the editors, unless stated otherwise, and do not necessarily reflect the opinions of the participating experts.

Alessandro Perilli	Virtualization.info (Independent Analyst)
Andrea Manieri	Ingegneria Informatica
Avner Algom	The Israeli Association of GRID Technologies
Craig Balding	Cloudsecurity.org
Dr. Guy Bunker	Bunker Associates
John Rhoton	Independent Consultant
Matt Broda	Microsoft
Mirco Rohr	Kaspersky
Ofer Biran	IBM
Pete Lindstrom	Spire Security
Dr Peter Dickman, Engineering Manager	Google Inc.
Philippe Massonet	Reservoir Project, CETIC
Raj Samani	Information Systems Security Association, UK
Simon Pascoe,	British Telecom
Srijith K. Nair, Theo Dimitrakos	The BEinGRID Project, British Telecom
Dr Simone Balboni	University of Bologna
Various	National Health Service (NHS) Technology Office, UK
Various	RSA
Various	Symantec, Symantec Hosted Services
Legal input was mainly drafted by	
Dr.Paolo Balboni	Baker & McKenzie - Tilburg University
Kieran Mccorry	Hewlett Packard
W. David Snead, P.C.	Attorney and Counselor

EXECUTIVE SUMMARY

Cloud computing is a new way of delivering computing resources, not a *new technology*. Computing services ranging from data storage and processing to software, such as email handling, are now available instantly, commitment-free and on-demand. Since we are in a time of belt-tightening, this new economic model for computing has found fertile ground and is seeing massive global investment. According to IDC's analysis, the worldwide forecast for cloud services in 2009 will be in the order of \$17.4bn (1). The estimation for 2013 amounts to \$44.2bn, with the European market ranging from €971m in 2008 to €6,005m in 2013 (2).

The key conclusion of this paper is that the cloud's economies of scale and flexibility are both a friend and a foe from a security point of view. The massive concentrations of resources and data present a more attractive target to attackers, but cloud-based defences can be more robust, scalable and cost-effective. This paper allows an informed assessment of the security risks and benefits of using cloud computing - providing security guidance for potential and existing users of cloud computing.

The security assessment is based on three use-case scenarios: 1) SME migration to cloud computing services, 2) the impact of cloud computing on service resilience, 3) cloud computing in e-Government (e.g., eHealth).

The new economic model has also driven technical change in terms of:

Scale: commoditisation and the drive towards economic efficiency have led to massive concentrations of the hardware resources required to provide services. This encourages economies of scale - for all the kinds of resources required to provide computing services.

Architecture: optimal resource use demands computing resources that are abstracted from underlying hardware. Unrelated customers who share hardware and software resources rely on logical isolation mechanisms to protect their data. Computing, content storage and processing are massively distributed. Global markets for commodities demand edge distribution networks where content is delivered and received as close to customers as possible. This tendency towards global distribution and redundancy means resources are usually managed in bulk, both physically and logically.

Given the reduced cost and flexibility it brings, a migration to cloud computing is compelling for many SMEs. However, the survey undertaken as part of this report (see [Survey - An SME Perspective on Cloud Computing](#)) confirms that major concerns for SMEs migrating to the cloud include the confidentiality of their information and liability for incidents involving the infrastructure.

Governments are also interested in the possibility of using cloud computing to reduce IT costs and increase capabilities. For example, the US government GSA (General Services Administration) now offers a portal for cloud computing services (3). Governments too, have serious hurdles to overcome - in terms of public perception of the secure processing of citizens' personal information in cloud computing infrastructures. On top of this, there are also legal and regulatory obstacles which prevent many eGovernment applications from moving to cloud. Nevertheless, both governments and SMEs face the reality that many of their employees will be using cloud-based services whether or not this is part of their official policy.

For cloud computing to reach the full potential promised by the technology, it must offer solid information security. This paper explains, based on concrete scenarios, what cloud computing means for network and information security, data protection and privacy. We look at the security benefits of cloud computing and its risks. We cover the technical, policy and legal implications. Most importantly, we make concrete recommendations on how to address the risks and maximise the benefits.

Finally, it is important to note that cloud computing can refer to several different service types, including Application/Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). The risks and benefits associated with each model will differ and so will the key considerations in contracting for this type of service. The following sections attempt to make the distinction when the risks or benefits apply differently to different cloud models.

TOP RECOMMENDATIONS

ASSURANCE FOR CLOUD CUSTOMERS

Cloud customers need assurance that providers are following sound security practices in mitigating the risks facing both the customer and the provider (e.g., DDoS attacks). They need this in order to make sound business decisions and to maintain or obtain security certifications. An early symptom of this need for assurance is that many cloud providers are swamped with requests for audits.

For this reason, we have expressed many of the report's recommendations as a standard list of questions which can be used to provide or obtain assurance.

Documents based on the check-list should provide a means for customers to:

1. assess the risk of adopting cloud services;
2. compare different cloud provider offerings;
3. obtain assurance from selected cloud providers;
4. reduce the assurance burden on cloud providers.

The security check-list covers all aspects of security requirements including legal issues, physical security, policy issues and technical issues.

LEGAL RECOMMENDATIONS

Most legal issues involved in cloud computing will currently be resolved during contract evaluation (ie, when making comparisons between different providers) or negotiations. The more common case in cloud computing will be selecting between different contracts on offer in the market (contract evaluation) as opposed to contract negotiations. However, opportunities may exist for prospective customers of cloud services to choose providers whose contracts are negotiable

Unlike traditional Internet services, standard contract clauses may deserve additional review because of the nature of cloud computing. The parties to a contract should pay particular attention to their rights and obligations related to notifications of breaches in security, data transfers, creation of derivative works, change of control, and access to data by law enforcement entities. Because the cloud can be used to outsource critical internal infrastructure, and the interruption of that infrastructure may have wide ranging effects, the parties should carefully consider whether standard limitations on liability adequately represent allocations of liability, given the parties' use of the cloud, or responsibilities for infrastructure.

Until legal precedent and regulations address security concerns specific to cloud computing, customers and cloud providers alike should look to the terms of their contract to effectively address security risks.

LEGAL RECOMMENDATIONS TO THE EUROPEAN COMMISSION

We recommend that the European Commission study or clarify the following:

- certain issues related to the Data Protection Directive and the recommendations of the Article 29 Data Protection Working Party;
- cloud providers obligation to notify their customers of data security breaches;
- how the liability exemptions for intermediaries arising from the eCommerce Directive articles 12-15 apply to cloud providers;.
- how best to support the minimum data protection standards and privacy certification schemes common across all the member States.

RESEARCH RECOMMENDATIONS

We recommend priority areas of research in order to improve the security of cloud computing technologies. The following are the categories we have considered with a few examples of specific areas from the full list:

BUILDING TRUST IN THE CLOUD

- Effects of different forms of breach reporting on security

- End-to-end data confidentiality in the cloud and beyond
- Higher assurance clouds, virtual private clouds etc

DATA PROTECTION IN LARGE SCALE CROSS-ORGANIZATIONAL SYSTEMS

- Forensics and evidence gathering mechanisms.
- Incident handling - monitoring and traceability
- International differences in relevant regulations including data protection and privacy

LARGE SCALE COMPUTER SYSTEMS ENGINEERING

- Resource isolation mechanisms - data, processing, memory, logs etc
- Interoperability between cloud providers
- Resilience of cloud computing. How can cloud improve resilience?

TOP SECURITY BENEFITS

SECURITY AND THE BENEFITS OF SCALE: put simply, all kinds of security measures are cheaper when implemented on a larger scale. Therefore the same amount of investment in security buys better protection. This includes all kinds of defensive measures such as filtering, patch management, hardening of virtual machine instances and hypervisors, etc. Other benefits of scale include: multiple locations, edge networks (content delivered or processed closer to its destination), timeliness of response, to incidents, threat management.

SECURITY AS A MARKET DIFFERENTIATOR: security is a priority concern for many cloud customers; many of them will make buying choices on the basis of the reputation for confidentiality, integrity and resilience of, and the security services offered by, a provider. This is a strong driver for cloud providers to improve security practices.

STANDARDISED INTERFACES FOR MANAGED SECURITY SERVICES: large cloud providers can offer a standardised, open interface to managed security services providers. This creates a more open and readily available market for security services.

RAPID, SMART SCALING OF RESOURCES: the ability of the cloud provider to dynamically reallocate resources for filtering, traffic shaping, authentication, encryption, etc, to defensive measures (e.g., against DDoS attacks) has obvious advantages for resilience.

AUDIT AND EVIDENCE-GATHERING: cloud computing (when using virtualisation) can provide dedicated, pay-per-use forensic images of virtual machines which are accessible without taking infrastructure off-

line, leading to less down-time for forensic analysis. It can also provide more cost-effective storage for logs allowing more comprehensive logging without compromising performance.

MORE TIMELY, EFFECTIVE AND EFFICIENT UPDATES AND DEFAULTS: default virtual machine images and software modules used by customers can be pre-hardened and updated with the latest patches and security settings according to fine-tuned processes; IaaS cloud service APIs also allow snapshots of virtual infrastructure to be taken regularly and compared with a baseline. Updates can be rolled out many times more rapidly across a homogenous platform than in traditional client-based systems that rely on the patching model.

BENEFITS OF RESOURCE CONCENTRATION: Although the concentration of resources undoubtedly has disadvantages for security [see Risks], it has the obvious advantage of cheaper physical perimeterisation and physical access control (per unit resource) and the easier and cheaper application of many security-related processes.

TOP SECURITY RISKS

The most important classes of cloud-specific risks identified in this paper are:

LOSS OF GOVERNANCE: in using cloud infrastructures, the client necessarily cedes control to the Cloud Provider (CP) on a number of issues which may affect security. At the same time, SLAs may not offer a commitment to provide such services on the part of the cloud provider, thus leaving a gap in security defences.

LOCK-IN: there is currently little on offer in the way of tools, procedures or standard data formats or services interfaces that could guarantee data, application and service portability. This can make it difficult for the customer to migrate from one provider to another or migrate data and services back to an in-house IT environment. This introduces a dependency on a particular CP for service provision, especially if data portability, as the most fundamental aspect, is not enabled..

ISOLATION FAILURE: multi-tenancy and shared resources are defining characteristics of cloud computing. This risk category covers the failure of mechanisms separating storage, memory, routing and even reputation between different tenants (e.g., so-called guest-hopping attacks). However it should be considered that attacks on resource isolation mechanisms (e.g., against hypervisors) are still less numerous and much more difficult for an attacker to put in practice compared to attacks on traditional OSs.

COMPLIANCE RISKS: investment in achieving certification (e.g., industry standard or regulatory requirements) may be put at risk by migration to the cloud:

- if the CP cannot provide evidence of their own compliance with the relevant requirements
- if the CP does not permit audit by the cloud customer (CC).

In certain cases, it also means that using a public cloud infrastructure implies that certain kinds of compliance cannot be achieved (e.g., PCI DSS (4)).

MANAGEMENT INTERFACE COMPROMISE: customer management interfaces of a public cloud provider are accessible through the Internet and mediate access to larger sets of resources (than traditional hosting providers) and therefore pose an increased risk, especially when combined with remote access and web browser vulnerabilities.

DATA PROTECTION: cloud computing poses several data protection risks for cloud customers and providers. In some cases, it may be difficult for the cloud customer (in its role as data controller) to effectively check the data handling practices of the cloud provider and thus to be sure that the data is handled in a lawful way. This problem is exacerbated in cases of multiple transfers of data, e.g., between federated clouds. On the other hand, some cloud providers do provide information on their data handling practices. Some also offer certification summaries on their data processing and data security activities and the data controls they have in place, e.g., SAS70 certification.

INSECURE OR INCOMPLETE DATA DELETION: when a request to delete a cloud resource is made, as with most operating systems, this may not result in true wiping of the data. Adequate or timely data deletion may also be impossible (or undesirable from a customer perspective), either because extra copies of data are stored but are not available, or because the disk to be destroyed also stores data from other clients. In the case of multiple tenancy and the reuse of hardware resources, this represents a higher risk to the customer than with dedicated hardware.

MALICIOUS INSIDER: while usually less likely, the damage which may be caused by malicious insiders is often far greater. Cloud architectures necessitate certain roles which are extremely high-risk. Examples include CP system administrators and managed security service providers.

NB: the risks listed above do not follow a specific order of criticality; they are just ten of the most important cloud computing specific risks identified during the assessment. The risks of using cloud computing should be compared to the risks of staying with traditional solutions, such as desktop-based models. To facilitate this, in the main document we have included estimates of relative risks as compared with a typical traditional environment.

Please note that it is often possible, and in some cases advisable, for the cloud customer to transfer risk to the cloud provider; *however not all risks can be transferred*: If a risk leads to the failure of a

business, serious damage to reputation or legal implications, it is hard or impossible for any other party to compensate for this damage. Ultimately, you can outsource responsibility but you can't outsource accountability.

CONTENTS

List of contributors	3
Executive summary	4
Top recommendations	5
Top security benefits	7
Top security risks	8
Contents	11
Target audience	14
Cloud computing - working definition	14
Survey of existing work	16
1. Security benefits of cloud computing	17
Security and the benefits of scale	17
Security as a market differentiator	17
Standardised interfaces for managed security services	18
Rapid, smart scaling of resources	18
Audit and evidence-gathering	18
More timely and effective and efficient updates and defaults	19
Audit and SLAs force better risk management	19
Benefits of resource concentration	19
2. Risk assessment	21
Use-case scenarios	21
Risk assessment process	21
3. Risks	23
Policy and organizational risks	25
R.1 Lock-in	25
R.2 Loss of governance	27
R.3 Compliance challenges	29
R.4 Loss of business reputation due to co-tenant activities	29

R.5	Cloud service termination or failure	30
R.6	Cloud provider acquisition	31
R.7	Supply chain failure	32
Technical risks		33
R.8	Resource exhaustion (under or over provisioning)	33
R.9	Isolation failure	34
R.10	Cloud provider malicious insider - abuse of high privilege roles	35
R.11	Management interface compromise (manipulation, availability of infrastructure)	36
R.12	Intercepting data in transit	37
R.13	Data leakage on up/download, intra-cloud	38
R.14	Insecure or ineffective deletion of data	38
R.15	Distributed denial of service (DDoS)	39
R.16	Economic denial of service (EDoS)	39
R.17	Loss of encryption keys	40
R.18	Undertaking malicious probes or scans	41
R.19	Compromise service engine	41
R.20	Conflicts between customer hardening procedures and cloud environment	42
Legal risks		43
R.21	Subpoena and e-discovery	43
R.22	Risk from changes of jurisdiction	44
R.23	Data protection risks	44
R.24	Licensing risks	45
Risks not specific to the cloud		46
R.25	Network breaks	46
R.26	Network management (ie, network congestion / mis-connection / non-optimal use)	47
R.27	Modifying network traffic	47
R.28	Privilege escalation	47
R.29	Social engineering attacks (ie, impersonation)	48
R.30	Loss or compromise of operational logs	49
R.31	Loss or compromise of security logs (manipulation of forensic investigation)	49
R.32	Backups lost, stolen	49
R.33	Unauthorized access to premises (including physical access to machines and other facilities)	50
R.34	Theft of computer equipment	50
R.35	Natural disasters	51
4.	Vulnerabilities	52
	Vulnerabilities not specific to the cloud	58
5.	Assets	60

6. Recommendations and key messages	63
Information assurance framework	63
Introduction	63
Division of liabilities	64
Division of responsibilities	64
Software as a Service	65
Platform as a Service	65
Infrastructure as a Service	66
Methodology	67
Note of caution	68
Note to governments	68
Information assurance requirements	69
Personnel security	69
Supply-chain assurance	70
Operational security	70
Identity and access management	73
Asset management	76
Data and Services Portability	76
Business Continuity Management	76
Physical security	78
Environmental controls	79
Legal requirements	80
Legal recommendations	81
Legal recommendations to the European Commission	82
Research recommendations	83
Glossary and abbreviations	86
Bibliography	91
ANNEX I – Cloud computing – Key legal issues	95
ANNEX II – SME use-case scenario	110
ANNEX III – Other use-case scenarios	118

TARGET AUDIENCE

The intended audiences of this report are:

- business leaders, of SMEs in particular, to facilitate their evaluation and mitigation of the risks associated with adopting cloud computing technologies;
- European policymakers, to aid them in deciding on research policy (to develop technologies to mitigate risks);
- European policymakers, to assist them in deciding on appropriate policy and economic incentives, legislative measures, awareness-raising initiatives, etc, vis-à-vis cloud-computing technologies;
- individuals or citizens, to enable them to evaluate the costs and benefits of using the consumer version of these applications.

CLOUD COMPUTING - WORKING DEFINITION

This is the working definition of cloud computing we are using for the purposes of this study. It is not intended as yet another definitive definition. Sources for our definition can be reviewed at (5), (6) and (54).

Cloud computing is an on-demand service model for IT provision, often based on virtualization and distributed computing technologies. Cloud computing architectures have:

- highly abstracted resources
- near instant scalability and flexibility
- near instantaneous provisioning
- shared resources (hardware, database, memory, etc)
- 'service on demand', usually with a 'pay as you go' billing system
- programmatic management (e.g., through WS API).

There are three categories of cloud computing:

- **Software as a service (SaaS):** is software offered by a third party provider, available on demand, usually via the Internet configurable remotely. Examples include online word processing and spreadsheet tools, CRM services and web content delivery services (Salesforce CRM, Google Docs, etc).

BENEFITS, RISKS AND RECOMMENDATIONS FOR INFORMATION SECURITY

- **Platform as a service (PaaS):** allows customers to develop new applications using APIs deployed and configurable remotely. The platforms offered include development tools, configuration management, and deployment platforms. Examples are Microsoft Azure, Force and Google App engine.
- **Infrastructure as service (IaaS):** provides virtual machines and other abstracted hardware and operating systems which may be controlled through a service API. Examples include Amazon EC2 and S3, Terremark Enterprise Cloud, Windows Live Skydrive and Rackspace Cloud.

Clouds may also be divided into:

- **public:** available publicly - any organisation may subscribe
- **private:** services built according to cloud computing principles, but accessible only within a private network
- **partner:** cloud services offered by a provider to a limited and well-defined number of parties.

In general, the commodity, cost, liability and assurance of clouds vary according to the following figure:




FIGURE 1: FEATURES OF PUBLIC, PARTNER AND PRIVATE CLOUDS

SURVEY OF EXISTING WORK

In compiling this report, we surveyed existing work on cloud security risks and their mitigation, including *Security Guidance for Critical Areas of Focus in Cloud Computing* (Cloud security Alliance (55)) *Cloud Cube Model: Selecting Cloud Formations for Secure Collaboration* (Jericho Forum (56)) and *Assessing the Security Risks of Cloud Computing* (Gartner (57)) in order to understand where to focus its results for maximum added value.

1. SECURITY BENEFITS OF CLOUD COMPUTING

It is hardly necessary to repeat the many rain-forests' worth of material which has been written on the economic, technical and architectural and ecological benefits of cloud computing. However, in the direct experience of the members of our expert group, as well as according to recent news from the 'real world', an examination of the security risks of cloud computing must be balanced by a review of its specific security benefits. Cloud computing has significant potential to improve security and resilience. What follows is a description of the key ways in which it can contribute.

SECURITY AND THE BENEFITS OF SCALE

Put simply, all kinds of security measures are cheaper when implemented on a larger scale. Therefore the same amount of investment in security buys better protection. This includes all kinds of defensive measures such as filtering, patch management, hardening of virtual machine instances and hypervisors, human resources and their management and vetting, hardware and software redundancy, strong authentication, efficient role-based access control and federated identity management solutions by default, which also improves the network effects of collaboration among various partners involved in defense. Other benefits of scale include:

- **Multiple locations:** most cloud providers have the economic resources to replicate content in multiple locations by default. This increases redundancy and independence from failure and provides a level of disaster recovery out-of-the-box.
- **Edge networks:** storage, processing and delivery closer to the network edge mean service reliability and quality is increased overall and local network problems are less likely to have global side-effects.
- **Improved timeliness of response:** larger **to incidents:** well-run larger-scale systems, for example due to early detection of new malware deployments, can develop more effective and efficient incident response capabilities.
- **Threat management:** cloud providers can also afford to hire specialists in dealing with specific security threats, while smaller companies can only afford a small number of generalists.

SECURITY AS A MARKET DIFFERENTIATOR

Security is a priority concern for many cloud customers [see the survey: [An SME perspective on Cloud Computing](#)] – customers will make buying choices on the basis of the reputation for confidentiality, integrity and resilience, and the security services offered by a provider, more so than in traditional environments. This is a strong driver for cloud providers to improve their security practices and compete on security.

STANDARDISED INTERFACES FOR MANAGED SECURITY SERVICES

Large cloud providers can offer a standardised, open interface to managed security services (MSS) providers offering services to all its customers. This potentially creates a more open and readily available market for security services where customers can switch providers more easily and with lower set-up costs.

RAPID, SMART SCALING OF RESOURCES

The list of cloud resources which can be rapidly scaled on demand already includes, e.g., storage, CPU time, memory, web service requests and virtual machine instances, and the level of granular control over resource consumption is increasing as technologies mature.

A cloud provider has the potential to dynamically reallocate resources for filtering, traffic shaping, encryption, etc, in order to increase support for defensive measures (e.g., against DDoS attacks) when an attack is likely or it is taking place. When this ability for dynamic resource reallocation is combined with appropriate resource optimisation methods, the cloud provider may be able to limit the effect that some attacks could have on the availability of resources that legitimately hosted services use, as well as limit the effect of increasing the use of resources by the security defence to combat such attacks. Achieving this requires however that the provider implements adequate coordination of autonomies for security defence and for resource management and optimisation.

The ability to dynamically scale defensive resources on demand has obvious advantages for resilience. Furthermore, the more all kinds of individual resources can be scaled in a granular way, without scaling all of the system resources, the cheaper it is to respond to sudden (non-malicious) peaks in demand.

AUDIT AND EVIDENCE-GATHERING

IaaS offerings support on-demand cloning of virtual machines. In the event of a suspected security breach, the customer can take an image of a live virtual machine – or virtual components thereof – for offline forensic analysis, leading to less down-time for analysis. With storage on tap, multiple clones can be created and analysis activities parallelised to reduce investigation time. This improves the ex-post analysis of security incidents and increases the probability of tracking attackers and patching weaknesses. However, it does presume the customer has access to trained forensic experts (which is not a standard cloud service as of writing).

It can also provide more cost-effective storage for logs, thus allowing more comprehensive logging without compromising performance. Pay as you go cloud storage brings transparency to your audit storage costs and makes adjusting to meet future audit log requirements easier. This makes the process of identifying security incidents as they happen more efficient (7).

MORE TIMELY AND EFFECTIVE AND EFFICIENT UPDATES AND DEFAULTS

Virtual machine images and software modules used by customers can be pre-hardened and updated with the latest patches and security settings according to fine-tuned processes; moreover, IaaS cloud service APIs also allow snapshots of virtual infrastructure to be taken regularly and compared with a baseline (e.g., to ensure software firewall rules have not changed) (8). Updates can be rolled out many times more rapidly across a homogenous platform than in traditional client-based systems that rely on the patching model. Finally in PaaS and SaaS models the applications are more likely to have been hardened to run outside the enterprise environment, which makes them likely to be more portable and robust than the equivalent enterprise software (where it exists). They are also more likely to be regularly updated and patched in a centralized fashion minimizing the window of vulnerability.

AUDIT AND SLAs FORCE BETTER RISK MANAGEMENT

The need to quantify penalties for various risk scenarios in SLAs and the possible impact of security breaches on reputation (see Security as market differentiator) motivate more rigorous internal audit and risk assessment procedures than would otherwise exist. The frequent audits imposed on CPs tend to expose risks which would not otherwise have been discovered, having therefore the same positive effect.

BENEFITS OF RESOURCE CONCENTRATION

Although the concentration of resources undoubtedly has disadvantages for security (see

Risks) it has the obvious advantage of cheaper physical perimeterisation and physical access control (per unit resource) and the easier and cheaper application of a comprehensive security policy and control over data management, patch management, incident management, and maintenance processes. The extent to which those savings are passed on to customers will obviously vary.

2. RISK ASSESSMENT

USE-CASE SCENARIOS

For the purposes of this risk assessment of cloud computing, we analyzed three use-case scenarios:

- An SME perspective on Cloud Computing
- The Impact of Cloud Computing on service resilience
- Cloud Computing and eGovernment (eHealth).

For the sake of brevity we decided to publish the complete version of the SME use-case scenario (see [ANNEX II](#)) and a summary of the resilience and eHealth scenarios (see [ANNEX III](#)).

This selection was based on the rationale that in Europe the cloud market is foreseen to have a great impact on new businesses and start-ups, as well as on the way current business models will evolve. Since EU industry is mainly composed by SMEs (99% of companies according to EU sources- (9)) it makes sense to focus on SMEs. Nevertheless, we have included several risks and recommendations which apply specifically to governments and larger enterprises.

The SME scenario is based on the results of the survey: An SME perspective on Cloud Computing (see [here](#)), and it is NOT meant to be a road map for companies considering, planning or running cloud computing projects and investments.

A medium-sized company was used as a use-case to guarantee to the assessment a high enough level of IT, legal and business complexity. The aim was to expose all possible information security risks. Some of those risks are specific to medium-sized businesses;, others are general risks that micro or small enterprises are also likely to face when migrating to a cloud computing approach.

The scenario was NOT intended to be completely realistic for any single cloud client or provider but all elements of the scenario are likely to occur in many organisations in the near future.

RISK ASSESSMENT PROCESS

The level of risk is estimated on the basis of the likelihood of an incident scenario, mapped against the estimated negative impact. The likelihood of an incident scenario is given by a threat exploiting vulnerability with a given likelihood.

The likelihood of each incident scenario and the business impact was determined in consultation with the expert group contributing to this report, drawing on their collective experience. In cases where it was judged not possible to provide a well founded estimation of the likelihood of an occurrence, the value is N/A. In many cases the estimate of likelihood depends heavily on the cloud model or architecture under consideration.

The following shows the risk level as a function of the business impact and likelihood of the incident scenario. The resulting risk is measured on a scale of 0 to 8 that can be evaluated against risk acceptance criteria. This risk scale could also be mapped to a simple overall risk rating:

- Low risk: 0-2
- Medium Risk: 3-5
- High Risk: 6-8

		Likelihood of incident scenario	Very Low (Very Unlikely)	Low (Unlikely)	Medium (Possible)	High (Likely)	Very High (Frequent)
Business Impact	Very Low	0	1	2	3	4	
	Low	1	2	3	4	5	
	Medium	2	3	4	5	6	
	High	3	4	5	6	7	
	Very High	4	5	6	7	8	

We have based the estimation of risk levels on ISO/IEC 27005:2008 (10).


3. RISKS

The following points should be noted in relation to the descriptions of risk below:

- Risk should always be understood in relation to overall business opportunity and appetite for risk – sometimes risk is compensated by opportunity.
- Cloud services are not only about convenient storage, accessible by multiple devices, but include important benefits such as more convenient communication and instant multi-point collaboration. Therefore, a comparative analysis needs to compare not only the risks of storing data in different places (on premises v the cloud) but also the risks when on premises-data stored on premises – e.g. a spreadsheet - is emailed to other persons for their contributions, against the security issues of a spreadsheet stored in the cloud and open to collaboration between those persons. Therefore, the risks of using cloud computing should be compared to the risks of staying with traditional solutions, such as desktop-based models.
- The level of risk will in many cases vary significantly with the type of cloud architecture being considered.
- It is possible for the cloud customer to transfer risk to the cloud provider and the risks should be considered against the cost benefit received from the services. However *not all risks can be transferred*: if a risk leads to the failure of a business, serious damage to reputation or legal implications, it is hard or impossible for any other party to compensate for this damage.
- The risk analysis in this paper applies to cloud technology. It does not apply to any specific cloud computing offering or company. This paper is not meant to replace a project-specific organisational risk assessment.
- The level of risks is expressed from the perspective of the cloud customer. Where the cloud provider point of view is considered, this is explicitly stated.

The following table shows the distribution of the risk probabilities and impacts:

PROBABILITY



IMPACT

FIGURE 2: RISK DISTRIBUTION

The risks identified in the assessment are classified into three categories:

- policy and organizational
- technical
- legal.

Each risk is presented in tables which include:

- probability level
- impact level
- reference to vulnerabilities
- reference to the affected assets
- level of risk.

Furthermore, *where meaningful*, we have added a comparative probability and impact cell to compare cloud computing risks and risks in standard IT approaches. We have not included a comparative risk since it is assumed that all the risks selected are higher.

POLICY AND ORGANIZATIONAL RISKS

R.1 LOCK-IN

Probability	HIGH	Comparative: Higher
Impact	MEDIUM	Comparative: Equal
Vulnerabilities	V13. Lack of standard technologies and solutions V46. Poor provider selection V47. Lack of supplier redundancy V31. Lack of completeness and transparency in terms of use	
Affected assets	A1. Company reputation A5. Personal sensitive data A6. Personal data A7. Personal data - critical A9. Service delivery – real time services A10. Service delivery	
Risk	HIGH	

There is currently little on offer in the way of tools, procedures or standard data formats or services interfaces that could guarantee data and service portability (although some initiatives do exist, e.g., see. (58)). This makes it extremely difficult for a customer to migrate from one provider to another, or to migrate data and services to or from an in-house IT environment. Furthermore, cloud providers may have an incentive to prevent (directly or indirectly) the portability of their customers services and data.

This potential dependency for service provision on a particular CP, depending on the CP's commitments, may lead to a catastrophic business failure should the cloud provider go bankrupt (see

R.5) and the content and application migration path to another provider is too costly (financially or time-wise) or insufficient warning is given (no early warning).

The acquisition of the cloud provider (R.6) can also have a similar effect, since it increases the likelihood of sudden changes in provider policy and non-binding agreements such as terms of use (ToU).

It is important to understand that the extent and nature of lock-in varies according to the cloud type:

SaaS Lock-in

- Customer data is typically stored in a custom database schema designed by the SaaS provider. Most SaaS providers offer API calls to read (and thereby 'export') data records. However, if the provider does not offer a readymade data 'export' routine, the customer will need to develop a program to extract their data and write it to file ready for import to another provider. It should be noted that there are few formal agreements on the structure of business records (e.g., a customer record at one SaaS provider may have different fields than at another provider), although there are common underlying file formats for the export and import of data, e.g., XML. The new provider can normally help with this work at a negotiated cost. However, if the data is to be brought back in-house, the customer will need to write import routines that take care of any required data mapping unless the CP offers such a routine. As customers will evaluate this aspect before making important migration decisions, it is in the long-term business interest of CPs to make data portability as easy, complete and cost-effective as possible.
- Application lock-in is the most obvious form of lock-in (although it is not specific to cloud services). SaaS providers typically develop a custom application tailored to the needs of their target market. SaaS customers with a large user-base can incur very high switching costs when migrating to another SaaS provider as the end-user experience is impacted (e.g., re-training is necessary). Where the customer has developed programs to interact with the providers API directly (e.g., for integration with other applications), these will also need to be re-written to take into account the new provider's API.

PaaS Lock-in

PaaS lock-in occurs at both the API layer (ie, platform specific API calls) and at the component level. For example, the PaaS provider may offer a highly efficient back-end data store. Not only must the customer develop code using the custom APIs offered by the provider, but they must also code data access routines in a way that is compatible with the back-end data store. This code will not necessarily be portable across PaaS providers, even if a seemingly compatible API is offered, as the data access model may be different (e.g., relational v hashing).

- PaaS lock-in at the API layer happens as different providers offer different APIs.
- PaaS lock-in happens at the runtime layer as ‘standard’ runtimes are often heavily customised to operate safely in a cloud environment. For example, a Java runtime may have ‘dangerous’ calls removed or modified for security reasons. The onus is on the customers' developers to understand and take into account these differences.
- PaaS also suffers from data lock-in, in the same way as in SaaS, but in this case the onus is completely on the customer to create compatible export routines.

IaaS-Lock-in

IaaS lock-in varies depending on the specific infrastructure services consumed. For example, a customer using cloud storage will not be impacted by non-compatible virtual machine formats.

- IaaS computing providers typically offer hypervisor based virtual machines. Software and VM metadata is bundled together for portability – typically just within the provider’s cloud. Migrating between providers is non-trivial until open standards, such as OVF (11), are adopted.
- IaaS storage provider offerings vary from simplistic key/value based data stores to policy enhanced file based stores. Feature sets can vary significantly, hence so do storage semantics. However application level dependence on specific policy features (e.g., access controls) may limit the customer’s choice of provider.
- Data lock-in is the obvious concern with IaaS storage services. As cloud customers push more data to cloud storage, data lock-in increases unless the CP provides for data portability.

Common to all providers is the possibility of a ‘run on the banks’ scenario for a cloud provider. For this scenario, suppose there is a crisis of confidence in the cloud provider’s financial position, and therefore a mass exit and withdrawal of content on a first come, first served basis. Then, in a situation where a provider limits the amount of ‘content’ (data and application code) which can be ‘withdrawn’ in a given timeframe, some customers will never be able to retrieve their data and applications.

R.2 LOSS OF GOVERNANCE

Probability	VERY HIGH	Comparative: Higher
Impact	VERY HIGH (depends on organization) (IaaS VERY HIGH, SaaS Low)	Comparative: Equal
Vulnerabilities	V34. Unclear roles and responsibilities V35. Poor enforcement of role definitions V21. Synchronizing responsibilities or contractual obligations external to cloud	

	<ul style="list-style-type: none"> V23. SLA clauses with conflicting promises to different stakeholders V25. Audit or certification not available to customers V22. Cross-cloud applications creating hidden dependency V13. Lack of standard technologies and solutions V29. Storage of data in multiple jurisdictions and lack of transparency about THIS V14. No source escrow agreement V16. No control on vulnerability assessment process V26. Certification schemes not adapted to cloud infrastructures V30. Lack of information on jurisdictions V31. Lack of completeness and transparency in terms of use V44. Unclear asset ownership
Affected assets	<ul style="list-style-type: none"> A1. Company reputation A2. Customer trust A3. Employee loyalty and experience A5. Personal sensitive data A6. Personal data A7. Personal data - critical A9. Service delivery – real time services A10. Service delivery
Risk	HIGH

In using cloud infrastructures, the client necessarily cedes control to the CP on a number of issues which may affect security. For example ToUs may prohibit port scans, vulnerability assessment and penetration testing. Moreover, there may be conflicts between customer hardening procedures and the cloud environment (see R 20). On the other hand, SLAs may not offer a commitment to provide such services on the part of the cloud provider, thus leaving a gap in security defenses. Moreover the cloud provider may outsource or sub-contract services to third-parties (unknown providers) which may not offer the same guarantees (such as to provide the service in a lawful way) as issued by the cloud provider. Or the control of the cloud provider changes, so the terms and conditions of their services may also change.

The loss of governance and control could have a potentially severe impact on the organization's strategy and therefore on the capacity to meet its mission and goals. The loss of control and governance could lead to the impossibility of complying with the security requirements, a lack of confidentiality, integrity and availability of data, and a deterioration of performance and quality of service, not to mention the introduction of compliance challenges (see R.3).

R.3 COMPLIANCE CHALLENGES

Probability	VERY HIGH – depends on PCI, SOX	Comparative: Higher
Impact	HIGH	Comparative: Equal
Vulnerabilities	V25. Audit or certification not available to customers V13. Lack of standard technologies and solutions, V29. Storage of data in multiple jurisdictions and lack of transparency about THIS V26. Certification schemes not adapted to cloud infrastructures V30. Lack of information on jurisdictions V31. Lack of completeness and transparency in terms of use	
Affected assets	A20. Certification	
Risk	HIGH	

Certain organisations migrating to the cloud have made considerable investments in achieving certification either for competitive advantage or to meet industry standards or regulatory requirements (e.g., PCI DSS). This investment may be put at risk by a migration to the cloud:

- if the CP cannot provide evidence of their own compliance to the relevant requirements;
- if the CP does not permit audit by the CC.

In certain cases, it also means that using a public cloud infrastructure implies that certain kinds of compliance cannot be achieved and hence cloud hosted services cannot be used for services that need them. For example, EC2 says customers would be hard-pressed to achieve PCI compliance on their platform. So EC2 hosted services cannot be used to handle credit card transactions.

R.4 LOSS OF BUSINESS REPUTATION DUE TO CO-TENANT ACTIVITIES

Probability	LOW
Impact	HIGH
Vulnerabilities	V6. Lack of resource isolation V7. Lack of reputational isolation

	V5. HYPERVISOR VULNERABILITIES
Affected assets	A1. Company reputation A5. Personal sensitive data A6. Personal data A7. Personal data - critical A9. Service delivery – real time services A10. Service delivery
Risk	MEDIUM

Resource sharing means that malicious activities carried out by one tenant may affect the reputation of another tenant. For example, spamming, port scanning or the serving of malicious content from cloud infrastructure can lead to:

- a range of IP addresses being blocked, including the attacker and other innocent tenants of an infrastructure;
- confiscation of resources due to neighbour activities (neighbour subpoenaed).

The impact can be deterioration in service delivery and data loss, as well as problems for the organization's reputation.

R.5 CLOUD SERVICE TERMINATION OR FAILURE

Probability	N/A	
Impact	VERY HIGH	Comparative: Higher
Vulnerabilities	V46. Poor provider selection V47. Lack of supplier redundancy V31. Lack of completeness and transparency in terms of use	
Affected assets	A1. Company reputation A2. Customer trust A3. Employee loyalty and experience A9. Service delivery – real time services A10. Service delivery	
Risk	MEDIUM	

As in any new IT market, competitive pressure, an inadequate business strategy, lack of financial support, etc, could lead some providers to go out of business or at least to force them to restructure their service portfolio offering. In other words, it is possible that in the short or medium term some cloud computing services could be terminated.

The impact of this threat for the cloud customer is easily understandable, since it could lead to a loss or deterioration of service delivery performance, and quality of service, as well as a loss of investment.

Furthermore, failures in the services outsourced to the CP may have a significant impact on the cloud customer’s ability to meet its duties and obligations to its own customers. The customer of the cloud provider may thus be exposed to contractual and tortious liability to its customers based on its provider’s negligence. Failures by the cloud provider may also result in liability by the customer to its employees.

R.6 CLOUD PROVIDER ACQUISITION

Probability	N/A	
Impact	MEDIUM	Comparative: Higher
Vulnerabilities	V31. Lack of completeness and transparency in terms of use	
Affected assets	A1. Company reputation A2. Customer trust A3. Employee loyalty and experience A4. Intellectual property A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A9. Service delivery – real time services A10. Service delivery	
Risk	MEDIUM	

Acquisition of the cloud provider could increase the likelihood of a strategic shift and may put non-binding agreements at risk (e.g., software interfaces, security investments, non-contractual security controls). This could make it impossible to comply with the security requirements. The final impact could be damaging for crucial assets such as: the organization’s reputation, customer or patient trust, and employee loyalty and experience.

R.7 SUPPLY CHAIN FAILURE

Probability	LOW	Comparative: Higher
Impact	MEDIUM	Comparative: Higher
Vulnerabilities	V31. Lack of completeness and transparency in terms of use V22. Cross-cloud applications creating hidden dependency V46. Poor provider selection V47. Lack of supplier redundancy	
Affected assets	A1. Company reputation A2. Customer trust A5. Personal sensitive data A6. Personal data A7. Personal data - critical A9. Service delivery – real time services A10. Service delivery	
Risk	LOW	

A cloud computing provider can outsource certain specialised tasks of its ‘production’ chain to third parties. In such a situation the level of security of the cloud provider may depend on the level of security of each one of the links and the level of dependency of the cloud provider on the third party. Any interruption or corruption in the chain or a lack of coordination of responsibilities between all the parties involved can lead to: unavailability of services, loss of data confidentiality, integrity and availability, economic and reputational losses due to failure to meet customer demand, violation of SLA, cascading service failure, etc.

An important example here is where a critical dependency exists on a third party single-sign-on or identity management service. In this case, an interruption of the third party service or of the CP’s connection to the service or a weakness in their security procedures may compromise the availability or confidentiality of a cloud customer or indeed the entire cloud offering.

In general, a lack of transparency in the contract can be a problem for the whole system. If a provider does not declare which core IT services are outsourced - it is not realistic that providers should list the contractors since these may change frequently - the customer is not in a position to properly evaluate the risk he is facing. This lack of transparency could decrease the level of trust in the provider.

TECHNICAL RISKS

R.8 RESOURCE EXHAUSTION (UNDER OR OVER PROVISIONING)

Probability	A. Inability to provide additional capacity to a customer: MEDIUM	Comparative: N/A
	B. Inability to provide current agreed capacity level: LOW	Comparative: Higher
Impact	A. Inability to provide additional capacity to a customer: LOW/MEDIUM (e.g., at Christmas)	Comparative: N/A
	B. Inability to provide current agreed capacity level: HIGH	Comparative: Same
Vulnerabilities	V15. Inaccurate modelling of resource usage V27. Inadequate resource provisioning and investments in infrastructure V28. No policies for resource capping V47. Lack of supplier redundancy	
Affected assets	A1. Company reputation A2. Customer trust A10. Service delivery A11. Access control / authentication / authorization (root/admin v others)	
Risk	MEDIUM	

Cloud services are on-demand services [see Cloud computing - working definition]. Therefore there is a level of calculated risk in allocating all the resources of a cloud service, because resources are allocated according to statistical projections. Inaccurate modelling of resources usage - common resources allocation algorithms are vulnerable to distortions of fairness - or inadequate resource provisioning and inadequate investments in infrastructure can lead, from the CP perspective, to:

- Service unavailability: failure in certain highly specific application scenarios which use a particular resource very intensively (ie, CPU/Memory intensive number crunching or simulation (eg. forecasting stock prices;

- Access control compromised: in some cases it may be possible to force a system to 'fail open' in the event of resource exhaustion. [Ref: CWE-400: Uncontrolled Resource Consumption - Resource Exhaustion (12)];
- Economic and reputational losses: due to failure to meet customer demand.
- The opposite consequences of inaccurate estimation of resource needs could lead to:
- Infrastructure oversize: excessive provisioning leading to economic losses and loss of profitability.

From the cloud customer perspective, a poor provider selection and lack of supplier redundancy could lead to:

- Service unavailability: failure in the delivery (or degrading performance) of services both in real time and not in real time;
- Access control system compromised: put the confidentiality and Integrity of data at risk;
- Economic and reputational losses: due to failure to meet customer demand, violation of SLA, cascading service failure, etc.

Note: this risk could be also a consequence of a DDoS attack (see R. 15) and of misbehaving applications due to poor application compartmentalization in some cloud providers' systems.

R.9 ISOLATION FAILURE

Probability	LOW (Private Cloud) MEDIUM (Public Cloud)	Comparative: Higher
Impact	VERY HIGH	Comparative: Higher
Vulnerabilities	V5. Hypervisor vulnerabilities V6. Lack of resource isolation V7. Lack of reputational isolation V17. Possibility that internal (cloud) network probing will occur V18. Possibility that co-residence checks will be performed	
Affected assets	A1. Company reputation A2. Customer trust A5. Personal sensitive data A6. Personal data A7. Personal data - critical A9. Service delivery – real time services A10. Service delivery	
Risk	HIGH	

Multi-tenancy and shared resources are two of the defining characteristics of cloud computing environments. Computing capacity, storage, and network are shared between multiple users. This class of risks includes the failure of mechanisms separating storage, memory, routing, and even reputation between different tenants of the shared infrastructure (e.g., so-called guest-hopping attacks, SQL injection attacks exposing multiple customers’ data stored in the same table, and side channel attacks).

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

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

R.10 CLOUD PROVIDER MALICIOUS INSIDER - ABUSE OF HIGH PRIVILEGE ROLES

Probability	MEDIUM (Lower than traditional)	Comparative: Lower
Impact	VERY HIGH (Higher than traditional)	Comparative: Higher (aggregate) Comparative: Same (for a single customer)
Vulnerabilities	V34. Unclear roles and responsibilities V35. Poor enforcement of role definitions V36. Need-to-know principle not applied V1. AAA vulnerabilities V39. System or OS vulnerabilities V37. Inadequate physical security procedures V10. Impossibility of processing data in encrypted form V48. Application vulnerabilities or poor patch management	
Affected assets	A1. Company reputation A2. Customer trust A3. Employee loyalty and experience A4. Intellectual property A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A9. Service delivery – real time services A10. Service delivery	
Risk	HIGH	

The malicious activities of an insider could potentially have an impact on: the confidentiality, integrity and availability of all kind of data, IP, all kind of services and therefore indirectly on the organization's reputation, customer trust and the experiences of employees. This can be considered especially important in the case of cloud computing due to the fact that cloud architectures necessitate certain roles which are extremely high-risk. Examples of such roles include CP system administrators and auditors and managed security service providers dealing with intrusion detection reports and incident response. As cloud use increases, employees of cloud providers increasingly become targets for criminal gangs (as has been witnessed in the financial services industry with call centre workers (13), (14)).

R.11 MANAGEMENT INTERFACE COMPROMISE (MANIPULATION, AVAILABILITY OF INFRASTRUCTURE)

Probability	MEDIUM	Comparative: Higher
Impact	VERY HIGH	Comparative: Higher
Vulnerabilities	V1. AAA vulnerabilities V4. Remote access to management interface V38. Misconfiguration V39. System or OS vulnerabilities V48. Application vulnerabilities or poor patch management	
Affected assets	A1. Company reputation A2. Customer trust A5. Personal sensitive data A6. Personal data A7. Personal data - critical A9. Service delivery – real time services A10. Service delivery A14. Cloud service management interface	
Risk	MEDIUM	

The customer management interfaces of public cloud providers are Internet accessible and mediate access to larger sets of resources (than traditional hosting providers) and therefore pose an increased risk especially when combined with remote access and web browser vulnerabilities. This includes customer interfaces controlling a number of virtual machines and, most importantly, CP interfaces controlling the operation of the overall cloud system. Of course, this risk may be mitigated by more investment in security by providers.

R.12 INTERCEPTING DATA IN TRANSIT

Probability	MEDIUM	Comparative: Higher (for a given piece of data)
Impact	HIGH	Comparative: Same
Vulnerabilities	V1. AAA vulnerabilities V8. Communication encryption vulnerabilities V9. Lack of or weak encryption of archives and data in transit V17. Possibility that internal (cloud) network probing will occur V18. Possibility that co-residence checks will be performed V31. Lack of completeness and transparency in terms of use	
Affected assets	A1. Company reputation A2. Customer trust A4. Intellectual property A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A23. Backup or archive data	
Risk	MEDIUM	

Cloud computing, being a distributed architecture, implies more data in transit than traditional infrastructures. For example, data must be transferred in order to synchronise multiple distributed machine images, images distributed across multiple physical machines, between cloud infrastructure and remote web clients, etc. Furthermore, most use of data-centre hosting is implemented using a secure VPN-like connection environment, a practice not always followed in the cloud context.

Sniffing, spoofing, man-in-the-middle attacks, side channel and replay attacks should be considered as possible threat sources.

Moreover, in some cases the CP does not offer a confidentiality or non-disclosure clause or these clauses are not sufficient to guarantee respect for the protection of the customer’s secret information and ‘know-how’ that will circulate in the ‘cloud’.

R.13 DATA LEAKAGE ON UP/DOWNLOAD, INTRA-CLOUD

Probability	MEDIUM (N/A)
Impact	HIGH
Vulnerabilities	V1. AAA vulnerabilities V8. Communication encryption vulnerabilities V17. Possibility that internal (cloud) network probing will occur V18. Possibility that co-residence checks will be performed V10. Impossibility of processing data in encrypted form V48. Application vulnerabilities or poor patch management
Affected assets	A1. Company reputation A2. Customer trust A3. Employee loyalty and experience A4. Intellectual property A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A12. Credentials A13. User directory (data) A14. Cloud service management interface
Risk	MEDIUM

This is the same as the previous risk, but applies to the transfer of data between the cloud provider and the cloud customer.

R.14 INSECURE OR INEFFECTIVE DELETION OF DATA

Probability	MEDIUM	Comparative: Higher
Impact	Very HIGH	Comparative: Higher
Vulnerabilities	V20. Sensitive media sanitization	
Affected assets	A5. Personal sensitive data A6. Personal data A7. Personal data - critical	

BENEFITS, RISKS AND RECOMMENDATIONS FOR INFORMATION SECURITY

	A12. Credentials
Risk	MEDIUM

Whenever a provider is changed, resources are scaled down, physical hardware is reallocated, etc, data may be available beyond the lifetime specified in the security policy. It may be impossible to carry out the procedures specified by the security policy, since full data deletion is only possible by destroying a disk which also stores data from other clients. When a request to delete a cloud resource is made, this may not result in true wiping of the data (as with most operating systems). Where true data wiping is required, special procedures must be followed and this may not be supported by the standard API (or at all).

If effective encryption is used then the level of risk may be considered to be lower.

R.15 DISTRIBUTED DENIAL OF SERVICE (DDoS)

Probability	Customer: MEDIUM	Comparative: Lower
	Provider: LOW	Comparative: N/A
Impact	Customer: HIGH	Comparative: Higher
	Provider: VERY HIGH	Comparative: Lower
Vulnerabilities	V38. Misconfiguration V39. System or OS vulnerabilities V53. Inadequate or misconfigured filtering resources	
Affected assets	A1. Company reputation A2. Customer trust A9. Service delivery – real time services A10. Service delivery A14. Cloud service management interface A16. Network (connections, etc)	
Risk	MEDIUM	

R.16 ECONOMIC DENIAL OF SERVICE (EDOS)

Probability	LOW
Impact	HIGH

Vulnerabilities	V1. AAA vulnerabilities V2. User provisioning vulnerabilities V3. User de-provisioning vulnerabilities V4. Remote access to management interface V28. No policies for resource capping
Affected assets	A1. Company reputation A2. Customer trust A9. Service delivery – real time services A10. Service delivery
Risk	MEDIUM

There are several different scenarios in which a cloud customer's resources may be used by other parties in a malicious way that has an economic impact:

- Identity theft: an attacker uses an account and uses the customer's resources for his own gain or in order to damage the customer economically.
- The CC has not set effective limits on the use of paid resources and experiences unexpected loads on these resources through no malicious actions.
- An attacker uses a public channel to use up the customer's metered resources - for example, where the customer pays per HTTP request, a DDoS attack can have this effect.

EDoS destroys economic resources; the worst case scenario would be the bankruptcy of the customer or a serious economic impact. NOTE: the general asset MONEY is not mentioned in the list.

R.17 LOSS OF ENCRYPTION KEYS

Probability	LOW	Comparative: N/A
Impact	HIGH	Comparative: Higher
Vulnerabilities	V11. Poor key management procedures V12. Key generation: low entropy for random number generation	
Affected assets	A4. Intellectual property A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A12. Credentials	
Risk	MEDIUM	

This includes disclosure of secret keys (SSL, file encryption, customer private keys, etc) or passwords to malicious parties, the loss or corruption of those keys, or their unauthorised use for authentication and non-repudiation (digital signature).

R.18 UNDERTAKING MALICIOUS PROBES OR SCANS

Probability	MEDIUM	Comparative: Lower
Impact	MEDIUM	Comparative: Lower
Vulnerabilities	V17. Possibility that internal (cloud) network probing will occur V18. Possibility that co-residence checks will be performed	
Affected assets	A1. Company reputation A2. Customer trust A9. Service delivery – real time services A10. Service delivery	
Risk	MEDIUM	

Malicious probes or scanning, as well as network mapping, are indirect threats to the assets being considered. They can be used to collect information in the context of a hacking attempt. A possible impact could be a loss of confidentiality, integrity and availability of service and data.

R.19 COMPROMISE SERVICE ENGINE

Probability	LOW
Impact	VERY HIGH
Vulnerabilities	V5. Hypervisor vulnerabilities V6. Lack of resource isolation
Affected assets	A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A9. Service delivery – real time services A10. Service delivery
Risk	MEDIUM

Each cloud architecture relies on a highly specialized platform, the service engine that sits above the physical hardware resources and manages customer resources at different levels of abstraction. For example, in IaaS clouds this software component can be the hypervisor. The service engine is developed and supported by cloud platform vendors and the open source community in some cases. It can be further customized by the cloud computing providers.

Like any other software layer, the service engine code can have vulnerabilities and is prone to attacks or unexpected failure. An attacker can compromise the service engine by hacking it from inside a virtual machine (IaaS clouds), the runtime environment (PaaS clouds), the application pool (SaaS clouds), or through its APIs.

Hacking the service engine may be useful to escape the isolation between different customer environments (jailbreak) and gain access to the data contained inside them, to monitor and modify the information inside them in a transparent way (without direct interaction with the application inside the customer environment), or to reduce the resources assigned to them, causing a denial of service.

R.20 CONFLICTS BETWEEN CUSTOMER HARDENING PROCEDURES AND CLOUD ENVIRONMENT

Probability	LOW
Impact	MEDIUM
Vulnerabilities	V31. Lack of completeness and transparency in terms of use V23. SLA clauses with conflicting promises to different stakeholders V34. Unclear roles and responsibilities
Affected assets	A4. Intellectual property A5. Personal sensitive data A6. Personal data A7. Personal data - critical
Risk	LOW

Cloud providers must set out a clear segregation of responsibilities that articulates the minimum actions customers must undertake. The failure of customers to properly secure their environments may pose a vulnerability to the cloud platform if the cloud provider has not taken the necessary steps to provide isolation. Cloud providers should further articulate their isolation mechanisms and provide best practice guidelines to assist customers to secure their resources.

Customers must realize and assume their responsibility as failure to do so would place their data and resources at further risk. In some cases cloud customers have inappropriately assumed that the cloud

provider was responsible for, and was conducting, all activities required to ensure security of their data. This assumption by the customer, and/or a lack of clear articulation by the cloud provider, placed unnecessary risk on the customer’s data. It is imperative that cloud customers identify their responsibilities and comply with them.

Cloud providers, by their very nature, are tasked with providing a multi-tenant environment, whether this is via virtualization on a server or the common network shared by the customers. The co-location of many customers inevitably causes conflict for the cloud provider as customers’ communication security requirements are likely to be divergent from each other.

Take, for example, the case of two customers on a shared traditional network infrastructure. If one customer wishes the network firewall to block all traffic except for SSH, but another customer is running a web server farm and requires passage of HTTP and HTTPS, who wins? This same type of issue is raised by customers who have competing and conflicting compliance requirements. This type of challenge only worsens as the number of tenants and the disparity of their requirements increase. Therefore, cloud providers must be in a position to deal with these challenges by way of technology, policy and transparency (where appropriate).

LEGAL RISKS

R.21 SUBPOENA AND E-DISCOVERY

Probability	HIGH
Impact	MEDIUM
Vulnerabilities	V6. Lack of resource isolation V29. Storage of data in multiple jurisdictions and lack of transparency about THIS V30 Lack of information on jurisdictions
Affected assets	A1. Company reputation A2. Customer trust A5. Personal sensitive data A6. Personal data A7 Personal data - critical A9. Service delivery – real time services A10. Service delivery
Risk	HIGH

In the event of the confiscation of physical hardware as a result of subpoena by law-enforcement agencies or civil suits (15), the centralisation of storage as well as shared tenancy of physical hardware means many more clients are at risk of the disclosure of their data to unwanted parties (16), (17), (18).

At the same time, it may become impossible for the agency of a single nation to confiscate 'a cloud' given pending advances around long distance hypervisor migration.

R.22 RISK FROM CHANGES OF JURISDICTION

Probability	VERY HIGH
Impact	HIGH
Vulnerabilities	V30. Lack of information on jurisdictions V29. Storage of data in multiple jurisdictions and lack of transparency about THIS
Affected assets	A1. Company reputation A2. Customer trust A5. Personal sensitive data A6. Personal data A7. Personal data - critical A9. Service delivery – real time services A10. Service delivery
Risk	HIGH

Customer data may be held in multiple jurisdictions, some of which may be high risk. If data centres are located in high-risk countries, e.g., those lacking the rule of law and having an unpredictable legal framework and enforcement, autocratic police states, states that do not respect international agreements, etc., sites... could be raided by local authorities and data or systems subject to enforced disclosure or seizure. Note that we are not implying here that all subpoena law-enforcement measures are unacceptable, merely that some may be so and that some legitimate seizures of hardware (which appear to be rare) may affect more customers than the targets of a law-enforcement action depending on how the data is stored (19), (20).

R.23 DATA PROTECTION RISKS

Probability	HIGH
--------------------	------

Impact	HIGH
Vulnerabilities	V30. Lack of information on jurisdictions V29. Storage of data in multiple jurisdictions and lack of transparency about THIS
Affected assets	A1. Company reputation A2. Customer trust A5. Personal sensitive data A6. Personal data A7. Personal data - critical A9. Service delivery – real time services A10. Service delivery
Risk	HIGH

Cloud computing poses several data protection risks for cloud customers and providers.

- It can be difficult for the cloud customer (in its role of data controller) to effectively check the data processing that the cloud provider carries out, and thus be sure that the data is handled in a lawful way. It has to be clear that the cloud customer will be the main person responsible for the processing of personal data, even when such processing is carried out by the cloud provider in its role of external processor. Failure to comply with data protection law may lead to administrative, civil and also criminal sanctions, which vary from country to country, for the data controller.. This problem is exacerbated in the case of multiple transfers of data e.g., between federated clouds. On the other hand, some cloud providers do provide information on the data processing that they carry out. Some also offer certification summaries of their data processing and data security activities and the data controls they have in place, e.g., SAS70 certification providers.
- There may be data security breaches which are not notified to the controller by the cloud provider.
- The cloud customer may lose control of the data processed by the cloud provider. This issue is increased in the case of multiple transfers of data (e.g., between federated cloud providers).
- The cloud provider may receive data that have not been lawfully collected by its customer (the controller).

R.24 LICENSING RISKS

Probability	MEDIUM	Comparative: Higher
Impact	MEDIUM	Comparative: Higher

Vulnerabilities	V31. Lack of completeness and transparency in terms of use
Affected assets	A1. Company reputation A9. Service delivery – real time services A20. Certification
Risk	MEDIUM

Licensing conditions, such as per-seat agreements, and online licensing checks may become unworkable in a cloud environment. For example, if software is charged on a per instance basis every time a new machine is instantiated then the cloud customer's licensing costs may increase exponentially even though they are using the same number of machine instances for the same duration. In the case of PaaS and IaaS, there is the possibility for creating original work in the cloud (new applications, software etc...). As with all intellectual property, if not protected by the appropriate contractual clauses (see ANNEX I – Cloud computing – Key legal issues , Intellectual Property), this original work may be at risk.

RISKS NOT SPECIFIC TO THE CLOUD

In the course of our risk analysis, we identified the following threats which are not specific to cloud computing, but should nevertheless be considered carefully when assessing the risk of a typical cloud-based system.

R.25 NETWORK BREAKS

Probability	LOW	Comparative: Same
Impact	VERY HIGH	Comparative: Higher
Vulnerabilities	V38. Misconfiguration V39. System or OS vulnerabilities V6. Lack of resource isolation V41. Lack of, or a poor and untested, business continuity and disaster recovery plan	
Affected assets	A9. Service delivery – real time services A10. Service delivery	
Risk	MEDIUM	

One of highest risks! Potentially thousands of customers are affected at the same time.

R.26 NETWORK MANAGEMENT (IE, NETWORK CONGESTION / MIS-CONNECTION / NON-OPTIMAL USE)

Probability	MEDIUM	Comparative: Same
Impact	VERY HIGH	Comparative: Higher
Vulnerabilities	V38. Misconfiguration V39. System or OS vulnerabilities V6. Lack of resource isolation V41. Lack of, or a poor and untested, business continuity and disaster recovery PLAN	
Affected assets	A1. Company reputation A2. Customer trust A3. Employee loyalty and experience A9. Service delivery – real time services A10. Service delivery A16 Network (connections, etc)	
Risk	HIGH	

R.27 MODIFYING NETWORK TRAFFIC

Probability	LOW	
Impact	HIGH	
Vulnerabilities	V2. User provisioning vulnerabilities V3. User de-provisioning vulnerabilities V8. Communication encryption vulnerabilities V16. No control on vulnerability assessment process	
Affected assets	A1. Company reputation A2. Customer trust A5. Personal sensitive data A6. Personal data A7. Personal data - critical A9. Service delivery – real time services A10. Service delivery	
Risk	MEDIUM	

R.28 PRIVILEGE ESCALATION

Probability	LOW	Comparative: Lower
Impact	HIGH	Comparative: Higher (for cloud provider)
Vulnerabilities	V1. AAA vulnerabilities	

	<p>V2. User provisioning vulnerabilities V3. User de-provisioning vulnerabilities V5. Hypervisor vulnerabilities V34. Unclear roles and responsibilities V35. Poor enforcement of role definitions V36. Need-to-know principle not applied V38. Misconfiguration</p>
Affected assets	<p>A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A11. Access control / authentication / authorization (root/admin v others) A13. User directory (data)</p>
Risk	MEDIUM

R.29 SOCIAL ENGINEERING ATTACKS (IE, IMPERSONATION)

Probability	MEDIUM	Comparative: Same
Impact	HIGH	Comparative: Higher
Vulnerabilities	<p>V32. Lack of security awareness V2. User provisioning vulnerabilities V6. Lack of resource isolation V8. Communication encryption vulnerabilities V37. Inadequate physical security procedures</p>	
Affected assets	<p>A1. Company reputation A2. Customer trust A3. Employee loyalty and experience A4. Intellectual property A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A11. Access control / authentication / authorization (root/admin v others) A12. Credentials</p>	
Risk	MEDIUM	

R.30 LOSS OR COMPROMISE OF OPERATIONAL LOGS

Probability	LOW	Comparative: Lower
Impact	MEDIUM	Comparative: Same (for customer)
Vulnerabilities	V52. Lack of policy or poor procedures for logs collection and retention V1. AAA vulnerabilities V2. User provisioning vulnerabilities V3. User de-provisioning vulnerabilities V19. Lack of forensic readiness V39. System or OS vulnerabilities	
Affected assets	A21. Operational logs (customer and cloud provider)	
Risk	LOW	

R.31 LOSS OR COMPROMISE OF SECURITY LOGS (MANIPULATION OF FORENSIC INVESTIGATION)

Probability	LOW	Comparative: Lower
Impact	MEDIUM	Comparative: Same (for customer)
Vulnerabilities	V52. Lack of policy or poor procedures for logs collection and retention V1. AAA vulnerabilities V2. User provisioning vulnerabilities V3. User de-provisioning vulnerabilities V19. Lack of forensic readiness V39. System or OS vulnerabilities	
Affected assets	A22. Security logs	
Risk	LOW	

R.32 BACKUPS LOST, STOLEN

Probability	LOW	Comparative: Lower
Impact	HIGH	Comparative: Same (for customer)
Vulnerabilities	V37. Inadequate physical security procedures Error! Reference source not found. V1. AAA vulnerabilities V2. User provisioning vulnerabilities V3. User de-provisioning vulnerabilities	
Affected assets	A1. Company reputation A2. Customer trust A5. Personal sensitive data	

	A6. Personal data A7. Personal data - critical A8. HR data A9. Service delivery – real time services A10. Service delivery A23. Backup or archive data
Risk	MEDIUM

R.33 UNAUTHORIZED ACCESS TO PREMISES (INCLUDING PHYSICAL ACCESS TO MACHINES AND OTHER FACILITIES)

Probability	VERY LOW	Comparative: Lower
Impact	HIGH (to have a very high impact it should a target attack (pointing to a specific machine, etc) otherwise the impact should be high.	Comparative: Higher
Vulnerabilities	V37. Inadequate physical security procedures	
Affected assets	A1. Company reputation A2. Customer trust A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A23. Backup or archive data	
Risk	LOW	

Since cloud providers concentrate resources in large data centres, and although the physical perimeter controls are likely to be stronger, the impact of a breach of those controls is higher.

R.34 THEFT OF COMPUTER EQUIPMENT

Probability	VERY LOW	Comparative: Lower
Impact	HIGH	Comparative: Higher
Vulnerabilities	V37. Inadequate physical security procedures	
Affected assets	A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A17. Physical hardware	

BENEFITS, RISKS AND RECOMMENDATIONS FOR INFORMATION SECURITY

Risk	LOW
-------------	-----

R.35 NATURAL DISASTERS

Probability	VERY LOW	Comparative: Lower
Impact	HIGH	Comparative: Higher
Vulnerabilities	V41. Lack of, or a poor and untested, business continuity and disaster recovery plan	
Affected assets	A1. Company reputation A2. Customer trust A5. Personal sensitive data A6. Personal data A7. Personal data - critical A8. HR data A9. Service delivery – real time services A10. Service delivery A23. Backup or archive data	
Risk	LOW	

Generally speaking, the risk from natural disasters is lower compared to traditional infrastructures because cloud providers offer multiple redundant sites and network paths by default.

4. VULNERABILITIES

The following list of vulnerabilities it is not exhaustive but is, however, detailed enough for the purposes of our analysis. It contains both cloud-specific and general information security vulnerabilities.

V1. AAA VULNERABILITIES

A poor system for authentication, authorization and accounting, could facilitate unauthorized access to resources, privileges escalation, impossibility of tracking the misuse of resources and security incidents in general, etc, through:

- insecure storage of cloud access credentials by customer;
- insufficient roles available;
- credentials stored on a transitory machine.

Furthermore, the cloud makes password based authentication attacks (trend of fraudster using a Trojan to steal corporate passwords) much more impactful since corporate applications are now exposed on the Internet. Therefore password-based authentication will become insufficient and a need for stronger or two-factor authentication for accessing cloud resources will be necessary.

V2. USER PROVISIONING VULNERABILITIES

- Customer cannot control provisioning process.
- Identity of customer is not adequately verified at registration.
- Delays in synchronisation between cloud system components (time wise and of profile content) happen.
- Multiple, unsynchronised copies of identity data are made.
- Credentials are vulnerable to interception and replay.

V3. USER DE-PROVISIONING VULNERABILITIES

De-provisioned credentials are still valid due to time delays in roll-out of revocation.

V4. REMOTE ACCESS TO MANAGEMENT INTERFACE

Theoretically, this allows vulnerabilities in end-point machines to compromise the cloud infrastructure (single customer or CP) through, for example, weak authentication of responses and requests.

V5. HYPERVISOR VULNERABILITIES

Hypervisor-layer attacks are very attractive: the hypervisor in fact fully controls the physical resources and the VMs running on top of it, so any vulnerability in this layer is extremely critical. Exploiting the hypervisor potentially means exploiting every VM. The first proof of concept of a layer-below attack against a hypervisor was given by King et al in the paper (21), where the authors introduce the concept of a virtual machine-based Rootkit. By then a few vulnerabilities had been identified in the most popular hypervisors (e.g., (22) and (23)) which can be exploited without administrator access rights at this time, but none of their results had been un-patched at the time of writing.

A typical scenario enabled by exploiting a hypervisor's vulnerability is the so called 'guest to host escape', an example of which is 'Cloudburst', a VMware vulnerability recently discovered and documented in reference (24). Another scenario is 'VM hopping': in which an attacker hacks a VM using some standard method and then – exploiting some hypervisor vulnerability – takes control of other VMs running on the same hypervisor. For more information, see an *Empirical Study into the Security Exposure to Hosts of Hostile Virtualized Environments*, see (25).

V6. LACK OF RESOURCE ISOLATION

Resource use by one customer can affect resource use by another customer.

IaaS cloud computing infrastructures mostly rely on architectural designs where physical resources are shared by multiple virtual machines and therefore multiple customers.

Vulnerabilities in the hypervisor security model may lead to unauthorized access to these shared resources. For example, virtual machines of Customer 1 and Customer 2 have their virtual hard drives saved in the same shared LUN (Logical Unit Number) inside a SAN. Customer 2 may be able to map the virtual hard drive of Customer 1 to its virtual machine and see and use the data inside it.

Hypervisors used in IaaS clouds offer rich APIs that the cloud provider uses to develop a proprietary management, provisioning and reporting interface that is exposed to its customers. Vulnerabilities in the hypervisor security model or in the 'management interfaces' may lead to unauthorized access to customer information. At the same time a vulnerability at this level may allow an attacker to manipulate the assets inside the cloud facility, provoking denial of service (e.g., shut down of running virtual machines), data leakage (e.g., the copying and transfer outside the cloud of virtual machines), data compromise (e.g., replacement of virtual machines with modified copies), or direct financial damage (e.g., replication and launch of many copies of the virtual machines).

Moreover lack of controls on cloud cartography and co-residence and the cross side channel vulnerabilities (see (26)) can pose serious risks to resources isolation. For example, if resource usage is not independent between Customer 1 and Customer 2, Customer 1 can map Customer 2's resources. This can be done, for example, by using controlled loading of Customer 2's resources while measuring changes in Customer 1's own patterns of resource availability.

Finally, the lack of tools to enforce a term of service (ToS) or a more specific service level agreement (SLA), such as quality of service (QoS) or distributed resource scheduling (DRS) products, could allow a customer to monopolize the use of the cloud facility, impacting the other customers with denial of service or poor performance.

V7. LACK OF REPUTATIONAL ISOLATION

Activities from one customer impact on the reputation of another customer.

V8. COMMUNICATION ENCRYPTION VULNERABILITIES

These vulnerabilities concern the possibility of reading data in transit via, for example, MITM attacks, poor authentication, acceptance of self-signed certificates, etc.

V9. LACK OF OR WEAK ENCRYPTION OF ARCHIVES AND DATA IN TRANSIT

Failure to encrypt data in transit, data held in archives and databases, un-mounted virtual machine images, forensic images and data, sensitive logs and other data at rest puts the data at risk. Of course the costs of implementing key management [V11] and processing costs must be taking account and set against the business risk introduced.

V10. IMPOSSIBILITY OF PROCESSING DATA IN ENCRYPTED FORM

Encrypting data at rest is not difficult, but despite recent advances in homomorphic encryption (27), there is little prospect of any commercial system being able to maintain this encryption during processing. In one article, Craig Gentry estimates that performing a web search with encrypted keywords -- a perfectly reasonable simple application of this algorithm -- would increase the amount of computing time by about a trillion (28). This means that for a long time to come, cloud customers doing anything other than storing data in the cloud must trust the cloud provider.

V11. POOR KEY MANAGEMENT PROCEDURES

Cloud computing infrastructures require the management and storage of many different kinds of keys; examples include session keys to protect data in transit (e.g., SSL keys), file encryption keys, key pairs identifying cloud providers, key pairs identifying customers, authorisation tokens and revocation certificates (29). Because virtual machines do not have a fixed hardware infrastructure and cloud based content tends to be geographically distributed, it is more difficult to apply standard controls, such as hardware security module (HSM) storage, to keys on cloud infrastructures. For example:

- HSMs are by necessity strongly physically protected (from theft, eavesdrop and tampering). This makes it very difficult for them to be distributed in the multiple locations used in cloud

architectures (i.e., geographically distributed and highly replicated). Key management standards such as PKCS#10 and associated standards such as PKCS#11 (30) do not provide standardised wrappers for interfacing with distributed systems.

- Key management interfaces which are accessible via the public Internet (even if indirectly) are more vulnerable, as security is reduced in the communication channel between the user and the cloud key storage and the mutual remote authentication mechanisms used.
- New virtual machines needing to authenticate themselves must be instantiated with some form of secret. The distribution of such secrets may present problems of scalability. The rapid scaling of certification authorities issuing key pairs is easily achieved if resources are determined in advance, but dynamic, unplanned scaling of hierarchical trust authorities is difficult to achieve because of the resource overhead in creating new authorities (registration or certification, in authenticating new components and distributing new credentials, etc).
- Revocation of keys within a distributed architecture is also expensive. Effective revocation essentially implies that applications check the status of the key (certificate usually) according to a known time constraint which determines the window of risk. Although distributed mechanisms exist for achieving this (see, e.g., (31) and (32)) it is challenging to ensure that different parts of the cloud receive an equivalent level of service so that they are not associated with different levels of risk. Centralised solutions such as OCSP are expensive and do not necessarily reduce the risk unless the CA and the CRL are tightly bound.

V12. KEY GENERATION: LOW ENTROPY FOR RANDOM NUMBER GENERATION

The combination of standard system images, virtualisation technologies and a lack of input devices means that systems have much less entropy than physical RNGs; see *Cloud Computing Security* (33). This means that an attacker on one virtual machine may be able to guess encryption keys generated on other virtual machines because the sources of entropy used to generate random numbers might be similar. This is not a difficult problem to solve, but if it is not considered during system design, it can have important consequences.

V13. LACK OF STANDARD TECHNOLOGIES AND SOLUTIONS

A lack of standards means that data may be 'locked-in' to a provider. This is a big risk should the provider cease operation. This may inhibit the use of managed security services and external security technologies such as FIM.

V14. NO SOURCE ESCROW AGREEMENT

Lack of source escrow means that if a PaaS or SaaS provider goes into bankruptcy, its customers are not protected.

V15. INACCURATE MODELLING OF RESOURCE USAGE

Cloud services are particularly vulnerable to resource exhaustion because they are provisioned statistically. Although many providers allow customers to reserve resources in advance, resource provisioning algorithms can fail due to:

- inaccurate modelling of resource usage, which can lead to overbooking or over-provisioning (in turn, leading to wasted resources on the part of the cloud provider). Well-known resource allocation algorithms are Token Bucket (34), Fair Queuing (35) and Class Based Queuing (36). These are vulnerable to distortions of fairness; for an example, see (37).
- failure of resource allocation algorithms due to extraordinary events (e.g., outlying news events for content delivery).
- failure of resource allocation algorithms using job or packet classification because resources are poorly classified.
- failures in overall resource provisioning (as opposed to temporary overloads).

V16. NO CONTROL ON VULNERABILITY ASSESSMENT PROCESS

Restrictions on port scanning and vulnerability testing are an important vulnerability which, combined with a ToU which places responsibility on the customer for securing elements of the infrastructure, is a serious security problem.

V17. POSSIBILITY THAT INTERNAL (CLOUD) NETWORK PROBING WILL OCCUR

Cloud customers can perform port scans and other tests on other customers within the internal network.

V18. POSSIBILITY THAT CO-RESIDENCE CHECKS WILL BE PERFORMED

Side-channel attacks exploiting a lack of resource isolation allow attackers to determine which resources are shared by which customers.

V19. LACK OF FORENSIC READINESS

While the cloud has the potential to improve forensic readiness, many providers do not provide appropriate services and terms of use to enable this. For example, SaaS providers will typically not provide access to the IP logs of clients accessing content. IaaS providers may not provide forensic services such as recent VM and disk images.

V20. SENSITIVE MEDIA SANITIZATION

Shared tenancy of physical storage resources means that sensitive data may leak because data destruction policies applicable at the end of a lifecycle may either be impossible to implement because, for example, media cannot be physically destroyed because a disk is still being used by another tenant or it cannot be located, or no procedure is in place.

V21. SYNCHRONIZING RESPONSIBILITIES OR CONTRACTUAL OBLIGATIONS EXTERNAL TO CLOUD

Cloud customers are often unaware of the responsibilities assigned to them within the terms of service. There is a tendency towards a misplaced attribution of responsibility for activities such as archive encryption to the cloud provider even when it is clearly stated in the terms of the contract between the two parties that no such responsibility has been undertaken.

V22. CROSS-CLOUD APPLICATIONS CREATING HIDDEN DEPENDENCY

Hidden dependencies exist in the services supply chain (intra- and extra-cloud dependencies) and the cloud provider architecture does not support continued operation from the cloud when the third parties involved, subcontractors or the customer company, have been separated from the service provider and vice versa.

V23. SLA CLAUSES WITH CONFLICTING PROMISES TO DIFFERENT STAKEHOLDERS

SLA clauses may also be in conflict with promises made by other clauses or clauses from other providers.

V24. SLA CLAUSES CONTAINING EXCESSIVE BUSINESS RISK

SLAs may carry too much business risk for a provider, given the actual risk of technical failures. From the customer point of view, SLAs may contain clauses which turn out to be detrimental - for example, in the area of intellectual property, an SLA might specify that the CP has the rights to any content stored on the cloud infrastructure.

V25. AUDIT OR CERTIFICATION NOT AVAILABLE TO CUSTOMERS

The CP cannot provide any assurance to the customer via audit certification.

For instance, some CP are using open source hypervisors or customised versions of them (e.g., Xen (38)) which have not reached any Common Criteria (39) certification, which is a fundamental requirement for some organizations (e.g., US government agencies).

Please note that we are not saying that there is a direct correlation between certification and vulnerability level (since we do not have enough information about the profile protection and security target of certified products).

V26. CERTIFICATION SCHEMES NOT ADAPTED TO CLOUD INFRASTRUCTURES

There are not any cloud-specific control, which means that security vulnerabilities are likely to be missed.

V27. INADEQUATE RESOURCE PROVISIONING AND INVESTMENTS IN INFRASTRUCTURE

Infrastructure investments take time. If predictive models fail, the cloud provider service can fail for a long period.

V28. NO POLICIES FOR RESOURCE CAPPING

If there is not a flexible and configurable way for the customer and/or the cloud provider to set limits on resources, this can be problematic when resource use is unpredictable.

V29. STORAGE OF DATA IN MULTIPLE JURISDICTIONS AND LACK OF TRANSPARENCY ABOUT THIS

Mirroring data for delivery by edge networks and redundant storage without real-time information available to the customer of where data is stored introduces a level of vulnerability. Companies may unknowingly violate regulations, especially if clear information is not provided about the jurisdiction of storage.

V30. LACK OF INFORMATION ON JURISDICTIONS

Data may be stored and/or processed in high risk jurisdictions where it is vulnerable to confiscation by forced entry. If this information is not available to the cloud customer, they cannot take steps to avoid it.

V31. LACK OF COMPLETENESS AND TRANSPARENCY IN TERMS OF USE**VULNERABILITIES NOT SPECIFIC TO THE CLOUD**

In the course of our risk analysis, we identified the following vulnerabilities which are not specific to cloud computing but which should, nevertheless, be considered carefully when assessing a typical cloud-based system.

V32. LACK OF SECURITY AWARENESS

Cloud customers are not aware of the risks they could face when migrating into the cloud, particularly those risks that are generated from cloud specific threats, ie, loss of control, vendor lock-in, exhausted CP resources, etc. This lack of awareness could also affect the cloud provider who may not be aware of the actions that should be taken to mitigate these risks.

V33. LACK OF VETTING PROCESSES

Since there may be very high privilege roles within cloud providers, due to the scale involved, the lack or inadequate vetting of the risk profile of staff with such roles is an important vulnerability.

V34. UNCLEAR ROLES AND RESPONSIBILITIES

These vulnerabilities regard the inadequate attribution of roles and responsibilities in the cloud provider organization.

V35. POOR ENFORCEMENT OF ROLE DEFINITIONS

Within the cloud provider, a failure to segregate roles may lead to excessively privileged roles which can make extremely large systems vulnerable. For example, no single person should be given access privileges to the entire cloud.

V36. NEED-TO-KNOW PRINCIPLE NOT APPLIED

This is a special case of a vulnerability regarding roles and responsibilities. Parties should not be given unnecessary access to data. If they are then this constitutes an unnecessary risk.

V37. INADEQUATE PHYSICAL SECURITY PROCEDURES

These can include:

- lack of physical perimeter controls (smart card authentication at entry);
- lack of electromagnetic shielding for critical assets vulnerable to eavesdropping.

V38. MISCONFIGURATION

This class of vulnerabilities include: inadequate application of security baseline and hardening procedures, human error and untrained administrator.

V39. SYSTEM OR OS VULNERABILITIES

V40. UNTRUSTED SOFTWARE

V41. LACK OF, OR A POOR AND UNTESTED, BUSINESS CONTINUITY AND DISASTER RECOVERY PLAN

V42. LACK OF, OR INCOMPLETE OR INACCURATE, ASSET INVENTORY

V43. LACK OF, OR POOR OR INADEQUATE, ASSET CLASSIFICATION

V44. UNCLEAR ASSET OWNERSHIP

V45. POOR IDENTIFICATION OF PROJECT REQUIREMENTS

These include a lack of consideration of security and legal compliance requirements, no systems and applications user involvement, unclear or inadequate business requirements, etc.

V46. POOR PROVIDER SELECTION**V47. LACK OF SUPPLIER REDUNDANCY****V48. APPLICATION VULNERABILITIES OR POOR PATCH MANAGEMENT**

This class of vulnerabilities include: bugs in the application code, conflicting patching procedures between provider and customer, application of untested patches, vulnerabilities in browsers, etc.

V49. RESOURCE CONSUMPTION VULNERABILITIES**V50. BREACH OF NDA BY PROVIDER****V51. LIABILITY FROM DATA LOSS (CP)****V52. LACK OF POLICY OR POOR PROCEDURES FOR LOGS COLLECTION AND RETENTION****V53. INADEQUATE OR MISCONFIGURED FILTERING RESOURCES****5. ASSETS**

Asset	Description or reference to above described elements	Owner [<i>actors or organisations involved</i>]	Perceived Value [<i>Scale: VERY LOW – LOW – MEDIUM – HIGH – VERY HIGH</i>]
A1. Company reputation		Cloud customer	VERY HIGH
A2. Customer trust	Includes goodwill, can be measured by complaints	Cloud customer	VERY HIGH
A3. Employee loyalty and experience		Cloud customer	HIGH
A4. Intellectual property		Cloud customer	HIGH

BENEFITS, RISKS AND RECOMMENDATIONS FOR INFORMATION SECURITY

A5. Personal sensitive data	(as defined in European Data Protection Directive)	CP / Cloud customer	VERY HIGH (as it includes data about who is using the home care system)
A6. Personal data	(as defined in European Data Protection Directive)	CP / Cloud customer	MEDIUM (operational value) / HIGH (value if lost)
A7. Personal data - critical	(all data included in the Personal Data category according to the European Data Protection Directive, and are classified or considered CRITICAL by the organization or company)	CP / Cloud customer	HIGH (operational value) / HIGH (value if lost)
A8. HR data	Data that are relevant from an operational perspective, beside the Data Protection requirements	Cloud customer	HIGH
A9. Service delivery – real time services	All those services that are time critical and that need a level of availability close to 100%	CP / Cloud customer	VERY HIGH
A10. Service delivery		CP / Cloud customer	MEDIUM
A11. Access control / authentication / authorization (root/admin v others)		CP / Cloud customer	HIGH
A12. Credentials	Of patients and of staff that access the system	Cloud customer	VERY HIGH
A13. User directory (data)	If it does not work then nobody gets in	Cloud customer	HIGH

A14. Cloud service management interface	This is the management interface (either web based or remote shell, or ...) that manages all the services provided through the cloud.	CP / Cloud customer	VERY HIGH
A15. Management interface APIs		CP / Cloud customer / EuropeanHealth	MEDIUM
A16. Network (connections, etc)	Includes the connections that are intra- and extra-cloud	CP / Cloud customer	HIGH
A17. Physical hardware		CP / Cloud customer	LOW (depends on how much you lose) / MEDIUM (could be serious if stolen and not protected)
A18. Physical buildings		CP / Cloud customer	HIGH
A19. CP Application (source code)		CP / Cloud customer	HIGH
A20. Certification	ISO, PCI DSS, etc	CP / Cloud customer	HIGH
A21. Operational logs (customer and cloud provider)	Those logs used to sustain and optimise business processes and for auditing purposes	CP / Cloud customer	MEDIUM
A22. Security logs	Useful as evidence of security breaches and forensics	CP / Cloud customer	MEDIUM
A23. Backup or archive data		CP / Cloud customer	MEDIUM

6. RECOMMENDATIONS AND KEY MESSAGES

This section includes the main set of recommendations and key messages:

- An Information Assurance Framework - a standard check-list of questions which can be used to obtain (by cloud customers) or provide (by cloud providers) assurance
- Legal recommendations
- Research recommendations.

INFORMATION ASSURANCE FRAMEWORK

INTRODUCTION

One of the most important recommendations of this report is a set of assurance criteria designed:

1. to assess the risk of adopting cloud services (comparing the risks in maintaining a 'classical' organization and architecture with the risks of migrating to a cloud computing environment).
2. to compare different cloud provider offers.
3. to obtain assurance from the selected cloud providers. The preparation of effective security questionnaires for third party service providers is a significant drain on resources for cloud customers and one which is difficult to achieve without expertise in cloud-specific architectures.
4. to reduce the assurance burden on cloud providers. A very important risk specific to cloud infrastructures is introduced by the requirement for NIS assurance. Many cloud providers find that a large number of customers request audits of their infrastructure and policies. This can create a critically high burden on security personnel and it also increases the number of people with access to the infrastructure, which significantly increases the risk of attack due to misuse of security-critical information, theft of critical or sensitive data, etc. Cloud providers will need to deal with this by establishing a clear framework for handling such requests.

This section of the recommendations provides a set of questions that an organisation can ask a cloud provider to assure themselves that they are sufficiently protecting the information entrusted to them.

These questions are intended to provide a minimum baseline. Any organisation may therefore have additional specific requirements not covered within the baseline.

Equally, this document does not provide a standard response format for the cloud provider, so responses are in a free text format. However it is intended to feed the questions into a more detailed

comprehensive framework which will be developed as a follow-up to this work; this will allow for a consistent, comparable set of responses. Such responses will provide a quantifiable metric as to the Information-assurance maturity of the provider.

It is intended for the aforementioned metric to be consistent against other providers that allow a comparison for end-user organisations.

DIVISION OF LIABILITIES

The following table shows the expected division of liabilities between customer and provider.

	Customer	Provider
Lawfulness of content	Full liability	Intermediary liability with Liability exemptions under the terms of the E-commerce Directive (1) and its interpretation. ¹
Security incidents (including data leakage, use of account to launch an attack)	Responsibility for due diligence for what is under its control according to contractual conditions	Responsibility for due diligence for what is under its control
European Data Protection Law status	Data controller	Data processor (external)

DIVISION OF RESPONSIBILITIES

With respect to security incidents, there needs to be a clear definition and understanding between the customer and the provider of security-relevant roles and responsibilities. The lines of such a division will vary greatly between SaaS offerings and IaaS offerings, with the latter delegating more responsibility to the customer. A typical and rational division of responsibility is shown in the following table. *In any case, for each type of service, the customer and provider should clearly define which of*

¹ Cf. definition of information society services as provided for in Art. 2 of Directive 98/48/EC as well as Art. 2 of Directive 2000/31/EC, in conjunction with exemptions contained in Articles 12-15 of Directive 2000/31/EC (e-Commerce Directive).

them is responsible for all the items on the list below. In the case of standard terms of service (ie, no negotiation possible), cloud customers should verify what lies within their responsibility.

SOFTWARE AS A SERVICE

Customer	Provider
<ul style="list-style-type: none"> • Compliance with data protection law in respect of customer data collected and processed • Maintenance of identity management system • Management of identity management system • Management of authentication platform (including enforcing password policy) 	<ul style="list-style-type: none"> • Physical support infrastructure (facilities, rack space, power, cooling, cabling, etc) • Physical infrastructure security and availability (servers, storage, network bandwidth, etc) • OS patch management and hardening procedures (check also any conflict between customer hardening procedure and provider security policy) • Security platform configuration (Firewall rules, IDS/IPS tuning, etc) • Systems monitoring • Security platform maintenance (Firewall, Host IDS/IPS, antivirus, packet filtering) • Log collection and security monitoring

PLATFORM AS A SERVICE

Customer	Provider
<ul style="list-style-type: none"> • Maintenance of identity management system • Management of identity management system • Management of authentication platform (including enforcing password policy) 	<ul style="list-style-type: none"> • Physical support infrastructure (facilities, rack space, power, cooling, cabling, etc) • Physical infrastructure security and availability (servers, storage, network bandwidth, etc) • OS patch management and hardening procedures (check also any conflict between customer hardening procedure

	<p>and provider security policy)</p> <ul style="list-style-type: none"> • Security platform configuration (firewall rules, IDS/IPS tuning, etc) • Systems monitoring • Security platform maintenance (firewall, Host IDS/IPS, antivirus, packet filtering) • Log collection and security monitoring
--	---

INFRASTRUCTURE AS A SERVICE

Customer	Provider
<ul style="list-style-type: none"> • Maintenance of identity management system • Management of identity management system • Management of authentication platform (including enforcing password policy) • Management of guest OS patch and hardening procedures (check also any conflict between customer hardening procedure and provider security policy) • Configuration of guest security platform (firewall rules, IDS/IPS tuning, etc) • Guest systems monitoring • Security platform maintenance (firewall, Host IDS/IPS, antivirus, packet filtering) • Log collection and security monitoring 	<ul style="list-style-type: none"> • Physical support infrastructure (facilities, rack space, power, cooling, cabling, etc) • Physical infrastructure security and availability (servers, storage, network bandwidth, etc) • Host Systems (hypervisor, virtual firewall, etc)

Where cloud customers are responsible for the security of their Infrastructures (in IaaS), they should consider the following:

APPLICATION SECURITY IN INFRASTRUCTURE AS A SERVICE

IaaS application providers treat the applications within the customer virtual instance as a 'black box' and therefore are completely agnostic concerning the operation and management of a customer's applications. The entire 'stack' – customer application, run time application platform (.Net, Java, Ruby, PHP, etc) – is run on the customers' server (on provider infrastructure) and is managed by the customers themselves. For this reason it is vitally important to note that customers must take full responsibility for securing their cloud-deployed applications. Here is a brief checklist and description relating to best practices for secure application design and management:

- Cloud deployed applications must be designed for the Internet threat model (even if they are deployed as part of a VPC - virtual private cloud).
- They must be designed or be embedded with standard security countermeasures to guard against the common web vulnerabilities (see OWASP top ten (40)).
- Customers are responsible for keeping their applications up to date – and must therefore ensure they have a patch strategy (to ensure that their applications are screened from malware and hackers scanning for vulnerabilities that allow unauthorised access to their data within the cloud to be gained).
- Customers should not be tempted to use custom implementations of authentication, authorisation and accounting (AAA) as these can become weak if not properly implemented.

In summary: enterprise distributed cloud applications must run with many controls in place to secure the host (and network – see previous section), user access, and application level controls (see OWASP (41) guides relating to secure web/online application design). Also, please note many main stream vendors such as Microsoft, Oracle, Sun, etc, publish comprehensive documentation on how to secure the configuration of their products.

METHODOLOGY

The key sections of this document are based on the broad classes of controls found in the ISO 27001/2 (42), (43) and BS25999 (44) standards. Details within these sections are derived from both the standards, as well as the requirements of industry best practices. Throughout, we have selected only those controls which are relevant to cloud providers and third party outsourcers.

The detailed framework scheduled for release in 2010 is intended to include additional standards such as NIST SP 800-53 (45).

NOTE OF CAUTION

The series of questions detailed within the following section is a selection of common controls. It is not intended to be an exhaustive list; equally, certain questions may not be applicable to particular implementations. Consequently this list should be used as a baseline of common controls, and further detail should be sought where required.

It is also worth noting that although it is possible to transfer many of the risks to an externally provisioned supplier, the true cost of transferring risk is very rarely realised. For example, a security incident that results in the unauthorised disclosure of customer data may result in financial loss to the provider, however the negative publicity and loss of consumer confidence, and potential regulatory penalties (PCI-DSS) would be felt by the end customer. Such a scenario highlights the importance of distinguishing risk from commercial risk. It is possible to transfer commercial risk, but the ultimate risk always remains with the end customer.

Any response to the results of a risk assessment - in particular the amount and type of investment in mitigation, should be decided on the basis of the risk appetite of the organisation and the opportunities and financial savings which are lost by following any particular risk mitigation strategy.

Cloud customers should also carry out their own, context-specific risk analysis. Some of available Risk Management / Risks Assessment methodologies can be found at: http://rm-inv.enisa.europa.eu/rm_ra_methods.html

As the business and regulatory environment changes and new risks arise, risk assessment should be a regular activity rather than a one off event.

NOTE TO GOVERNMENTS

The following controls are aimed primarily at SMEs assessing cloud providers. They may also be useful to governments with the following provisos. *The characteristics of the cloud used should be considered carefully in relation to any government body's information classification scheme.*

- The use of public clouds – even with favourable responses from the following questionnaire – is not recommended for anything but the lowest assurance classes of data.
- For higher assurance classes of data, the list of suggested checks in this report is valid but should be supplemented with additional checks. This report is not intended to cover such controls, but the following are examples of issues which should be covered:
 - Does the provider offer transparent information and full control over the current physical location of all data? High assurance data is often restricted by location.
 - Does the provider support the data classification scheme used?

- What guarantees does the provider offer that customer resources are fully isolated (e.g., no sharing of physical machines)?
- Assuming physical machines are not shared between customers, to what degree are storage, memory and other data traces fully erased before machines are reallocated.
- Does the provider support or even mandate physical token based 2-factor authentication for client access?
- Does the provider hold ISO 27001/2 certification? What is the scope of the certification?
- Do the products used by the provider have Common Criteria certifications? At which level? Which protection profile and security target for the product?

INFORMATION ASSURANCE REQUIREMENTS

PERSONNEL SECURITY

The majority of questions relating to personnel will be similar to those you would ask your own IT personnel or other personnel who are dealing with your IT. As with most assessments, there is a balance between the risks and the cost.

- What policies and procedures do you have in place when hiring your IT administrators or others with system access? These should include:
 - pre-employment checks (identity, nationality or status, employment history and references, criminal convictions, and vetting (for senior personnel in high privilege roles)).
- Are there different policies depending on where the data is stored or applications are run?
 - For example, hiring policies in one region may be different from those in another.
 - Practices need to be consistent across regions.
 - It may be that sensitive data is stored in one particular region with appropriate personnel.
- What security education program do you run for all staff?
- Is there a process of continuous evaluation?
 - How often does this occur?
 - Further interviews
 - Security access and privilege reviews
 - Policy and procedure reviews.

SUPPLY-CHAIN ASSURANCE

The following questions apply where the cloud provider subcontracts some operations that are key to the security of the operation to third parties (e.g., a SaaS provider outsourcing the underlying platform to a third party provider, a cloud provider outsourcing the security services to a managed security services provider, use of an external provider for identity management of operating systems, etc). It also includes third parties with physical or remote access to the cloud provider infrastructure. It is assumed that this entire questionnaire may be applied recursively to third (or nth) party cloud service providers.

- Define those services that are outsourced or subcontracted in your service delivery supply chain that are key to the security (including availability) of your operations.
- Detail the procedures used to assure third parties accessing your infrastructure (physical and/or logical).
 - Do you audit your outsourcers and subcontractors and how often?
- Are any SLA provisions guaranteed by outsourcers lower than the SLAs you offer to your customers? If not, do you have supplier redundancy in place?
- What measures are taken to ensure third party service levels are met and maintained?
- Can the cloud provider confirm that security policy and controls are applied (contractually) to their third party providers?

OPERATIONAL SECURITY

It is expected that any commercial agreement with external providers will include service levels for all network services. However, in addition to the defined agreements, the end customer should still ensure that the provider employs appropriate controls to mitigate unauthorised disclosure.

- Detail your change control procedure and policy. This should also include the process used to re-assess risks as a result of changes and clarify whether the outputs are available to end customers.
- Define the remote access policy.
- Does the provider maintain documented operating procedures for information systems?
- Is there a staged environment to reduce risk, e.g., development, test and operational environments, and are they separated?
- Define the host and network controls employed to protect the systems hosting the applications and information for the end customer. These should include details of certification against external standards (e.g., ISO 27001/2).
- Specify the controls used to protect against malicious code.

- Are secure configurations deployed to only allow the execution of authorised mobile code and authorised functionality (e.g., only execute specific commands)?
- Detail policies and procedures for backup. This should include procedures for the management of removable media and methods for securely destroying media no longer required. (Depending on his business requirements, the customer may wish to put in place an independent backup strategy. This is particularly relevant where time-critical access to back-up is required.)

Audit logs are used in the event of an incident requiring investigation; they can also be used for troubleshooting. For these purposes, the end customer will need assurance that such information is available:

- Can the provider detail what information is recorded within audit logs?
 - For what period is this data retained?
 - Is it possible to segment data within audit logs so they can be made available to the end customer and/or law enforcement without compromising other customers and still be admissible in court?
 - What controls are employed to protect logs from unauthorised access or tampering?
 - What method is used to check and protect the integrity of audit logs?
- How are audit logs reviewed? What recorded events result in action being taken?
- What time source is used to synchronise systems and provide accurate audit log time stamping?

SOFTWARE ASSURANCE

- Define controls used to protect the integrity of the operating system and applications software used. Include any standards that are followed, e.g., OWASP (46), SANS Checklist (47), SAFECODE (48).
- How do you validate that new releases are fit-for-purpose or do not have risks (backdoors, Trojans, etc)? Are these reviewed before use?
- What practices are followed to keep the applications safe?
- Is a software release penetration tested to ensure it does not contain vulnerabilities? If vulnerabilities are discovered, what is the process for remedying these?

PATCH MANAGEMENT

- Provide details of the patch management procedure followed.
- Can you ensure that the patch management process covers all layers of the cloud delivery technologies – ie, network (infrastructure components, routers and switches, etc), server operating systems, virtualisation software, applications and security subsystems (firewalls, antivirus gateways, intrusion detection systems, etc)?

NETWORK ARCHITECTURE CONTROLS

- Define the controls used to mitigate DDoS (distributed denial-of-service) attacks.
 - Defence in depth (deep packet analysis, traffic throttling, packet black-holing, etc)

- Do you have defences against ‘internal’ (originating from the cloud providers networks) attacks as well as external (originating from the Internet or customer networks) attacks?
- What levels of isolation are used?
 - for virtual machines, physical machines, network, storage (e.g., storage area networks), management networks and management support systems, etc.
- Does the architecture support continued operation from the cloud when the company is separated from the service provider and vice versa (e.g., is there a critical dependency on the customer LDAP system)?
- Is the virtual network infrastructure used by cloud providers (in PVLANS and VLAN tagging 802.1q (49) architecture) secured to vendor and/or best practice specific standards (e.g., are MAC spoofing, ARP poisoning attacks, etc, prevented via a specific security configuration)?

HOST ARCHITECTURE

- Does the provider ensure virtual images are hardened by default?
- Is the hardened virtual image protected from unauthorized access?
- Can the provider confirm that the virtualised image does not contain the authentication credentials?
- Is the host firewall run with only the minimum ports necessary to support the services within the virtual instance?
- Can a host-based intrusion prevention service (IPS) be run in the virtual instance?

PAAS – APPLICATION SECURITY

Generally speaking, PaaS service providers are responsible for the security of the platform software stack, and the recommendations throughout this document are a good foundation for ensuring a PaaS provider has considered security principles when designing and managing their PaaS platform. It is often difficult to obtain detailed information from PaaS providers on exactly how they secure their platforms – however the following questions, along with other sections within this document, should be of assistance in assessing their offerings.

- Request information on how multi-tenanted applications are isolated from each other – a high level description of containment and isolation measures is required.
- What assurance can the PaaS provider give that access to your data is restricted to your enterprise users and to the applications you own?
- The platform architecture should be classic ‘sandbox’ – does the provider ensure that the PaaS platform sandbox is monitored for new bugs and vulnerabilities?

- PaaS providers should be able to offer a set of security features (re-useable amongst their clients) – do these include user authentication, single sign on, authorisation (privilege management), and SSL/TLS (made available via an API)?

SAAS – APPLICATION SECURITY

The SaaS model dictates that the provider manages the entire suite of applications delivered to end-users. Therefore SaaS providers are mainly responsible for securing these applications. Customers are normally responsible for operational security processes (user and access management). However the following questions, along with other sections within this document, should assist in assessing their offerings:

- What administration controls are provided and can these be used to assign read and write privileges to other users?
- Is the SaaS access control fine grained and can it be customised to your organisations policy?

RESOURCE PROVISIONING

- In the event of resource overload (processing, memory, storage, network)?
 - What information is given about the relative priority assigned to my request in the event of a failure in provisioning?
 - Is there a lead time on service levels and changes in requirements?
- How much can you scale up? Does the provider offer guarantees on maximum available resources within a minimum period?
- How fast can you scale up? Does the provider offer guarantees on the availability of supplementary resources within a minimum period?
- What processes are in place for handling large-scale trends in resource usage (e.g., seasonal effects)?

IDENTITY AND ACCESS MANAGEMENT

The following controls apply to the cloud provider's identity and access management systems (those under their control):

AUTHORISATION

- Do any accounts have system-wide privileges for the entire cloud system and, if so, for what operations (read/write/delete)?
- How are the accounts with the highest level of privilege authenticated and managed?
- How are the most critical decisions (e.g., simultaneous de-provisioning of large resource blocks) authorised (single or dual, and by which roles within the organisation)?

- Are any high-privilege roles allocated to the same person? Does this allocation break the segregation of duties or least privilege rules?
- Do you use role-based access control (RBAC)? Is the principle of least privilege followed?
- What changes, if any, are made to administrator privileges and roles to allow for extraordinary access in the event of an emergency?
- Is there an 'administrator' role for the customer? For example, does the customer administrator have a role in adding new users (but without allowing him to change the underlying storage!)?

IDENTITY PROVISIONING

- What checks are made on the identity of user accounts at registration? Are any standards followed? For example, the e-Government Interoperability Framework?
- Are there different levels of identity checks based on the resources required?
- What processes are in place for de-provisioning credentials?
- Are credentials provisioned and de-provisioned simultaneously throughout the cloud system, or are there any risks in de-provisioning them across multiple geographically distributed locations?

MANAGEMENT OF PERSONAL DATA

- What data storage and protection controls apply to the user directory (e.g., AD, LDAP) and access to it?
- Is user directory data exportable in an interoperable format?
- Is need-to-know the basis for access to customer data within the cloud provider?

KEY MANAGEMENT

For keys under the control of the cloud provider:

- Are security controls in place for reading and writing those keys? For example, strong password policies, keys stored in a separate system, hardware security modules (HSM) for root certificate keys, smart card based authentication, direct shielded access to storage, short key lifetime, etc.
- Are security controls in place for using those keys to sign and encrypt data?
- Are procedures in place in the event of a key compromise? For example, key revocation lists.
- Is key revocation able to deal with simultaneity issues for multiple sites?
- Are customer system images protected or encrypted?

ENCRYPTION

- Encryption can be used in multiple places – where is it used?
 - data in transit

- data at rest
 - data in processor or memory?
- Usernames and passwords?
- Is there a well-defined policy for what should be encrypted and what should not be encrypted?
- Who holds the access keys?
- How are the keys protected?

AUTHENTICATION

- What forms of authentication are used for operations requiring high assurance? This may include login to management interfaces, key creation, access to multiple-user accounts, firewall configuration, remote access, etc.
- Is two-factor authentication used to manage critical components within the infrastructure, such as firewalls, etc?

CREDENTIAL COMPROMISE OR THEFT

- Do you provide anomaly detection (the ability to spot unusual and potentially malicious IP traffic and user or support team behaviour)? For example, analysis of failed and successful logins, unusual time of day, and multiple logins, etc.
- What provisions exist in the event of the theft of a customer's credentials (detection, revocation, evidence for actions)?

IDENTITY AND ACCESS MANAGEMENT SYSTEMS OFFERED TO THE CLOUD CUSTOMER

The following questions apply to the identity and access management systems which are offered by the cloud provider for use and control by the cloud customer:

IDENTITY MANAGEMENT FRAMEWORKS

- Does the system allow for a federated IDM infrastructure which is interoperable both for high assurance (OTP systems, where required) and low assurance (eg. username and password)?
- Is the cloud provider interoperable with third party identity providers?
- Is there the ability to incorporate single sign-on?

ACCESS CONTROL

- Does the client credential system allow for the separation of roles and responsibilities and for multiple domains (or a single key for multiple domains, roles and responsibilities)?
- How do you manage access to customer system images – and ensure that the authentication and cryptographic keys are not contained within in them?

AUTHENTICATION

- How does the cloud provider identify itself to the customer (ie, is there mutual authentication)?
 - when the customer sends API commands?
 - when the customer logs into the management interface?
- Do you support a federated mechanism for authentication?

ASSET MANAGEMENT

It is important to ensure the provider maintains a current list of hardware and software (applications) assets under the cloud providers control. This enables checks that all systems have appropriate controls employed, and that systems cannot be used as a backdoor into the infrastructure.

- Does the provider have an automated means to inventory all assets, which facilitates their appropriate management?
- Is there a list of assets that the customer has used over a specific period of time?

The following questions are to be used where the end customer is deploying data that would require additional protection (i.e.. deemed as sensitive).

- Are assets classified in terms of sensitivity and criticality?
 - If so, does the provider employ appropriate segregation between systems with different classifications and for a single customer who has systems with different security classifications?

DATA AND SERVICES PORTABILITY

This set of questions should be considered in order to understand the risks related to vendor lock-in.

- Are there documented procedures and APIs for exporting data from the cloud?
- Does the vendor provide interoperable export formats for all data stored within the cloud?
- In the case of SaaS, are the API interfaces used standardised?
- Are there any provisions for exporting user-created applications in a standard format?
- Are there processes for testing that data can be exported to another cloud provider – should the client wish to change provider, for example?
- Can the client perform their own data extraction to verify that the format is universal and is capable of being migrated to another cloud provider?

BUSINESS CONTINUITY MANAGEMENT

Providing continuity is important to an organisation. Although it is possible to set service level agreements detailing the minimum amount of time systems are available, there remain a number of additional considerations.

- Does the provider maintain a documented method that details the impact of a disruption?
 - What are the RPO (recovery point objective) and RTO (recovery time objective) for services? Detail according to the criticality of the service.
 - Are information security activities appropriately addressed in the restoration process?
 - What are the lines of communication to end customers in the event of a disruption?
 - Are the roles and responsibilities of teams clearly identified when dealing with a disruption?
- Has the provider categorised the priority for recovery, and what would be our relative priority (the end customer) to be restored? Note: this may be a category (HIGH/MED/LOW).
- What dependencies relevant to the restoration process exist? Include suppliers and outsource partners.
- In the event of the primary site being made unavailable, what is the minimum separation for the location of the secondary site?

INCIDENT MANAGEMENT AND RESPONSE

Incident management and response is a part of business continuity management. The goal of this process is to contain the impact of unexpected and potentially disrupting events to an acceptable level for an organization.

To evaluate the capacity of an organization to minimize the probability of occurrence or reduce the negative impact of an information security incident, the following questions should be asked to a cloud provider:

- Does the provider have a formal process in place for detecting, identifying, analyzing and responding to incidents?
- Is this process rehearsed to check that incident handling processes are effective? Does the provider also ensure, during the rehearsal, that everyone within the cloud provider's support organisation is aware of the processes and of their roles during incident handling (both during the incident and post analysis)?
- How are the detection capabilities structured?
 - How can the cloud customer report anomalies and security events to the provider?
 - What facilities does the provider allow for customer-selected third party RTSM services to intervene in their systems (where appropriate) or to co-ordinate incident response capabilities with the cloud provider?
 - Is there a real time security monitoring (RTSM) service in place? Is the service outsourced? What kind of parameters and services are monitored?
 - Do you provide (upon request) a periodical report on security incidents (e.g., according to the ITIL definition)?

- For how long are the security logs retained? Are those logs securely stored? Who has access to the logs?
- Is it possible for the customer to build a HIPS/HIDS in the virtual machine image? Is it possible to integrate the information collected by the intrusion detection and prevention systems of the customer into the RTSM service of the cloud provider or that of a third party?
- How are severity levels defined?
- How are escalation procedures defined? When (if ever) is the cloud customer involved?
- How are incidents documented and evidence collected?
- Besides authentication, accounting and audit, what other controls are in place to prevent (or minimize the impact of) malicious activities by insiders?
- Does the provider offer the customer (upon request) a forensic image of the virtual machine?
- Does the provider collect incident metrics and indicators (ie,. number of detected or reported incidents per months, number of incidents caused by the cloud provider's subcontractors and the total number of such incidents, average time to respond and to resolve, etc)?).
 - Which of these does the provider make publicly available (NB not all incident reporting data can be made public since it may compromise customer confidentiality and reveal security critical information)??)
- How often does the provider test disaster recovery and business continuity plans?
- Does the provider collect data on the levels of satisfaction with SLAs?
- Does the provider carry out help desk tests? For example:
 - Impersonation tests (is the person at the end of the phone requesting a password reset, really who they say they are?) or so called 'social engineering' attacks.
- Does the provider carry out penetration testing? How often? What are actually tested during the penetration test – for example, do they test the security isolation of each image to ensure it is not possible to 'break out' of one image into another and also gain access to the host infrastructure?. The tests should also check to see if it is possible to gain access, via the virtual image, to the cloud providers management and support systems (e.g., example the provisioning and admin access control systems).
- Does the provider carry out vulnerability testing? How often?
- What is the process for rectifying vulnerabilities (hot fixes, re-configuration, uplift to later versions of software, etc)?

PHYSICAL SECURITY

As with personnel security, many of the potential issues arise because the IT infrastructure is under the control of a third party – like traditional outsourcing, the effect of a physical security breach can have an impact on multiple customers (organizations).

BENEFITS, RISKS AND RECOMMENDATIONS FOR INFORMATION SECURITY

- What assurance can you provide to the customer regarding the physical security of the location? Please provide examples, and any standards that are adhered to, e.g., Section 9 of ISO 27001/2.
 - Who, other than authorised IT personnel, has unescorted (physical) access to IT infrastructure?
 - For example, cleaners, managers, 'physical security' staff, contractors, consultants, vendors, etc.
 - How often are access rights reviewed?
 - How quickly can access rights be revoked?
 - Do you assess security risks and evaluate perimeters on a regular basis?
 - How frequently?
 - Do you carry out regular risk assessments which include things such as neighboring buildings?
 - Do you control or monitor personnel (including third parties) who access secure areas?
 - What policies or procedures do you have for loading, unloading and installing equipment?
 - Are deliveries inspected for risks before installation?
 - Is there an up-to-date physical inventory of items in the data centre?
 - Do network cables run through public access areas?
 - Do you use armoured cabling or conduits?
 - Do you regularly survey premises to look for unauthorized equipment?
 - Is there any off-site equipment?
 - How is this protected?
 - Do your personnel use portable equipment (e.g., laptops, smart phones) which can give access to the data centre?
 - How are these protected?
 - What measures are in place to control access cards?
 - What processes or procedures are in place to destroy old media or systems when required to do so?
 - data overwritten?
 - physical destruction?
 - What authorization processes are in place for the movement of equipment from one site to another?
 - How do you identify staff (or contractors) who are authorized to do this?
 - How often are equipment audits carried out to monitor for unauthorised equipment removal?
 - How often are checks made to ensure that the environment complies with the appropriate legal and regulatory requirements?

ENVIRONMENTAL CONTROLS

- What procedures or policies are in place to ensure that environmental issues do not cause an interruption to service?

- What methods do you use to prevent damage from a fire, flood, earthquake, etc?
 - In the event of a disaster, what additional security measures are put in place to protect physical access?
 - Both at the primary as well as at the secondary sites?
- Do you monitor the temperature and humidity in the data centre?
 - Air-conditioning considerations or monitoring?
- Do you protect your buildings from lightning strikes?
 - Including electrical and communication lines?
- Do you have stand-alone generators in the event of a power failure?
 - For how long can they run?
 - Are there adequate fuel supplies?
 - Are there failover generators?
 - How often do you check UPS equipment?
 - How often do you check your generators?
 - Do you have multiple power suppliers?
- Are all utilities (electricity, water, etc) capable of supporting your environment?

How often is this re-evaluated and tested?

- Is your air-conditioning capable of supporting your environment?
 - How often is it tested?
- Do you follow manufacturers recommended maintenance schedules?
- Do you only allow authorised maintenance or repair staff onto the site?
 - How do you check their identity?
- When equipment is sent away for repair, is the data cleaned from it first?
 - How is this done?

LEGAL REQUIREMENTS

Customers and potential customers of cloud provider services should have regard to their respective national and supra-national obligations for compliance with regulatory frameworks and ensure that any such obligations are appropriately complied with.

The key legal questions the customer should ask the cloud provider are:

- In what country is the cloud provider located?
- Is the cloud provider's infrastructure located in the same country or in different countries?
- Will the cloud provider use other companies whose infrastructure is located outside that of the cloud provider?
- Where will the data be physically located?
- Will jurisdiction over the contract terms and over the data be divided?
- Will any of the cloud provider's services be subcontracted out?

- Will any of the cloud provider's services be outsourced?
- How will the data provided by the customer and the customer's customers, be collected, processed and transferred?
- What happens to the data sent to the cloud provider upon termination of the contract?

LEGAL RECOMMENDATIONS

At the present time, most of the legal issues involved in cloud computing will be resolved during the evaluation of contracts, ToUs, User Licensing Agreements (ULAs) and SLAs by the customer. It is important to differentiate between the case of a small to medium sized organisation which would make a choice between different contracts offered on the market, and a larger organisation, which would be in a position to negotiate clauses. In the legal analysis of this paper, we take the perspective of the small-to-medium sized organisation which is assessing different contracts, SLAs, etc, offered on the market, since this is the more common case. This is because the business model of cloud computing differs from that of outsourcing: in order to deliver some of the benefits to its customers, cloud computing relies on the economies of scale from providing a low cost, commodity service, as opposed to a service specifically tailored to a customer's needs. Larger organisations may however use the same considerations when negotiating contracts. While past experiences with similar Internet technologies provide some guidance to allow customers and cloud providers to assess the security risks involved in cloud computing, it is necessary for both to consider the unique nature of cloud computing when evaluating these risks.

Although there is much common ground, certain standard contract clauses may deserve additional review because of the nature of cloud computing. Particular attention should be paid to rights and obligations relating to notifications of breaches in security, data transfer, creation of derivative works, change of control, and access to data by law enforcement entities. Because the cloud can be used to outsource critical internal infrastructure, and the interruption of that infrastructure may have wide ranging effects, attention should be paid to whether the standard limitations of liability adequately represent allocations of liability, given the parties' usage of the cloud, or the allocation of responsibilities for infrastructure [see **Division of responsibilities**].

Until legal precedent clarifies concerns in relation to data security that are specific to cloud computing, customers and cloud providers alike should look to the terms of their contract to effectively address risks.

The following is a list of areas the customer should pay particular attention to when assessing SLAs, ToUs, ULAs and other agreements for cloud services:

1. **Data Protection:** attention should be paid to choosing a processor that provides sufficient

technical security measures and organisational measures governing the processing to be carried out, and ensuring compliance with those measures

2. **Data Security:** attention should be paid to mandatory data security measures that potentially cause either the cloud provider or the customer to be subject to regulatory and judicial measures if the contract does not address these obligations.
3. **Data Transfer:** attention should be paid to what information is provided to the customer regarding how data is transferred within the cloud provider's proprietary cloud, outside that cloud, and within and outside the European Economic Area.
4. **Law Enforcement Access:** each country has unique restrictions on, and requirements providing for, law enforcement access to data. The customer should pay attention to information available from the provider about the jurisdictions in which data may be stored and processed and evaluate any risks resulting from the jurisdictions which may apply.
5. **Confidentiality and Non-disclosure:** the duties and obligations related to this issue should be reviewed.
6. **Intellectual property:** in the case of IaaS and PaaS, intellectual property, including original works created using the cloud infrastructure, may be stored. The cloud customer should ensure that the contract respects their rights to any intellectual property or original works as far as possible without compromising the quality of service offered (e.g. backups may be a necessary part of offering a good service level).
7. **Risk Allocation and limitation of liability:** when reviewing their respective contract obligations, the parties should underscore those obligations that present significant risk to them by including monetary remediation clauses, or obligations to indemnify, for the other party's breach of that contract obligation. Furthermore, any standard clauses covering limitations of liability should be evaluated carefully.
8. **Change of Control:** transparency concerning the cloud provider's continuing ability to honour their contract obligations in the case of a change of control, as well as any possibility to rescind the contract.

The legal recommendations expressed are generally from the cloud customer perspective.

LEGAL RECOMMENDATIONS TO THE EUROPEAN COMMISSION

We recommend that the European Commission study or clarify the following:

1. Certain issues related to the Data Protection Directive and Article 29 Data Protection Working Party recommendations warrant clarification. In particular:

2. Under which circumstances the Cloud Provider may be classified as a Joint Controller;
3. The application of Section 25(2) of the Data Protection Directive as applied to the processing of data in countries outside the European Economic Area during the data's transfer from one cloud provider to another, or within the company's cloud. ²
4. The impact of data transfers to, and from, countries outside the European Economic Areas, if those countries do not ensure an adequate level of protection for the data.
5. Whether the concept of "transferring data" should be re-examined in the light of technological advancements since the Directive was originally drafted, particularly in light of an accountability-based legal approach (e.g. as proposed by the Galway project (51)).
6. Whether cloud providers should have an obligation to notify their customers of data security breaches, and what information those customers should be required to pass on to end customers. This could also be accomplished through contractual clauses so it should be investigated which means would be more effective. For example, legislation on breach reporting could be difficult to enforce and could even act as a disincentive to transparency.
7. Whether it is necessary for Member States to clarify how the intermediary liability exemptions of the eCommerce Directive (articles 12-15) apply to Cloud providers.
8. The differences in Member States regarding laws governing enforcement requests by various public authorities for data stored in the cloud, in particular with a view to evaluate the differences of the level of protection vis-à-vis government requests of personal data stored on premise (home or businesses), and personal data stored in the cloud.

How best to support a minimum data protection standards and privacy certification schemes, based on accountability concepts which is common across the globe or at least all the EU Member States.

More details on the five legal issues can be found in [ANNEX I](#).

RESEARCH RECOMMENDATIONS

We recommend the following as priority areas for research in order to improve the security of cloud computing technologies:

² The ePrivacy directive, as revised in 2009 (<http://register.consilium.europa.eu/pdf/en/09/st03/st03674.en09.pdf>), requires Member States to introduce a security breach notification scheme. Note that this scheme will be applicable to electronic communication networks and electronic communication services, not to information society services such as cloud computing services

BUILDING TRUST IN THE CLOUD

- Certification processes and standards for clouds: more generally, cloud computing security lifecycle standards that can be certified against cloud specific provisions for governance standards – COBIT (52), ITIL (53), etc;
- Metrics for security in cloud computing;
- Return on security investments (ROSI): the measures cloud computing can enable to improve the accuracy of ROI for security;
- Effects of different forms of reporting breaches on security;
- Techniques for increasing transparency while maintaining appropriate levels of security:
 - Tagging, e.g., location tagging, data type tagging, policy tagging
 - Privacy preserving data provenance systems, e.g., tracing data end-to-end through systems;
- End-to-end data confidentiality in the cloud and beyond:
 - Encrypted search (long term)
 - Encrypted processing schemes (long term)
 - Encryption and confidentiality tools for social applications in the cloud
 - Trusted computing in clouds, e.g., trusted boot sequences for virtual machine stacks;
- Higher assurance clouds, virtual private clouds, etc;
- Extending cloud-based trust to client-based data and applications.

DATA PROTECTION IN LARGE-SCALE CROSS-ORGANIZATIONAL SYSTEMS

The following areas require further research with respect to cloud computing:

- Data destruction and lifecycle management
- Integrity verification - of backups and archives in the cloud and their version management
- Forensics and evidence gathering mechanisms
- Incident handling - monitoring and traceability
- Dispute resolution and rules of evidence
- International differences in relevant regulations, including data protection and privacy
 - Legal means to facilitate the smooth functioning of multi-national cloud infrastructures
 - Automated means to mitigate problems with different jurisdictions.

LARGE-SCALE COMPUTER SYSTEMS ENGINEERING

- Security in depth within large-scale distributed computer systems;
- Security services in the cloud – de-perimeterisation of security technologies and the adaptation of traditional security perimeter control technologies to the cloud, e.g., HSMs, web filters, firewalls, IDSs, etc;
- Resource isolation mechanisms - data, processing, memory, logs, etc;

- Interoperability between cloud providers;
- Portability of VM, data and VM security settings from one cloud provider to another (to avoid vendor lock-in), and maintaining state and session in VM backups and the long distance live migration of virtual machines;
- Standardization of interfaces to feed data, applications and whole systems to the cloud – so that every OS can develop the corresponding client interface;
- Resource (bandwidth and CPU, etc) provisioning and allocation at scale (elasticity);
- Scalable security management (policy and operating procedures) within cloud platforms:
 - automatic enforcement of security and data protection policies
 - secure operating processes of providers - the implementation of governance processes;
- Resilience of cloud computing - how to improve the resilience of a cloud:
 - use of cloud architectures at the client side (edge networks, p2p, etc)
 - aggregating multiple client networks
 - client-based redundancy and backup;
 - cloud bursting and global scale resilience in clouds.

Another useful source of information for research recommendation will be the PROCENT (Priorities of Research on Current & Emerging Network Technologies) report, due to be published in December 2009. Please consult the following: <http://www.enisa.europa.eu/act/res/technologies/procent>.

GLOSSARY AND ABBREVIATIONS

AAA	Authentication, authorization and accounting
AD	Active directory
API	Application programming interface - specification of interface published by software supplier
ARP	Address resolution protocol (2)
Asset	The target of protection in a security analysis
Availability	The proportion of time for which a system can perform its function
BS	British Standard
CA	Certification authority
CC	Common Criteria
Confidentiality	Ensuring that information is accessible only to those authorized to have access (ISO 17799)
Co-residence	Sharing of hardware or software resources by cloud customers
CP	Cloud provider
CRL	Certificate revocation list
CRM	Customer relationship management
Data controller	The 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; where the purposes and means of processing are determined by national or Community laws or regulations, the controller or the specific criteria

	for his nomination may be designated by national or Community law.
Data processor	A natural or legal person, public authority, agency or any other body which processes personal data on behalf of the controller.
Data subject	Identified or identifiable natural person (see EU Directive 95/46/EC) from whom data is collected and/or about whom that data is processed
DDoS	Distributed denial of service
De-provision	The process of enforcing the removal of a resource from use, or disallowing its use by a set of users
Edge network	In this context, a network of computers which is able to process and store data for delivery close to the final destination
EDoS	Economic denial of service
Escrow	The storage of a resource by a third party which has access to that resource when certain well-defined conditions are satisfied
FIM	Federated identity management
Guest OS	An OS under the control of the cloud customer, running in a virtualised environment
Host OS	The operating system of the cloud provider which runs multiple guest OSs
HSM	Hardware security module
Https	Http connection using TLS or SSL
Hypervisor	Computer software or hardware platform virtualization software that allows multiple

	operating systems to run on a host computer concurrently
IDS	Intrusion detection system
Integrity	The property that data has not been maliciously or accidentally altered during storage or transmission
IP	Internet protocol
IPS	Intrusion protection system
ISO	International Organization for Standardization
LDAP	Lightweight Directory Access Protocol
MAC	Media access control (address of a network node in IP protocol)
MITM	Man in the middle (a form of attack)
MSS	Managed security services
NIS	Network and information security
NIST	National Institute of Standards and Technology (US)
Non-repudiation	The property whereby a party in a dispute cannot repudiate or refute the validity of a statement or contract
OCSP	Online Certificate Status Protocol
OS	Operating system
OTP	One-time password (type of authentication token)
OVF	Open virtualisation format
Perimeterisation	The control of access to an asset or group of

BENEFITS, RISKS AND RECOMMENDATIONS FOR INFORMATION SECURITY

	assets
Port scan	Probing a network host to determine which ports are open and what services they offer
Protection profile	A document specifying security evaluation criteria to substantiate vendors' claims of a given family of information system products (a term used in Common Criteria)
Provision	The issuing of a resource
PV LAN	Private VLAN
QoS	Quality of service
RBAC	Role-based access control
Resilience	The ability of a system to provide and maintain an acceptable level of service in the face of faults (unintentional, intentional, or naturally caused)
ROI	Return on investment
ROSI	Return on security investment
RPO	Recovery point objective
RTO	Recovery time objective
RTSM	Real-time security monitoring
Security target	A document specifying security evaluation criteria to substantiate the vendor's claims for the product's security properties (a term used in Common Criteria)
Service engine	The system responsible for delivering cloud services
Side channel attack	Any attack based on information gained from the physical implementation of a system; e.g., timing

	information, power consumption, electromagnetic leaks or even sound can provide an extra source of information which can be exploited to break the system.
SLA	Service level agreement
SSL	Secure Sockets Layer (used for encrypting traffic between web servers and browsers)
Subpoena	In this context, a legal authority to confiscate evidence
TLS	Transport Layer Security (used for encrypting traffic between web servers and browsers)
ToU	Terms of use
UPS	Uninterruptable power supply
VLAN	Virtual local area network
VM	Virtual machine
VPC	Virtual private cloud
VPN	Virtual private network
Vulnerability	Any circumstance or event with the potential to adversely impact an asset through unauthorized access, destruction, disclosure, modification of data, and/or denial of service
XML	Extensible Mark-up Language

BIBLIOGRAPHY

1. **IDC Cloud Computing 2010 - An IDC Update**, Frank Gens, Robert P Mahowald, Richard L Villars, Sep 2009 - Doc # TB20090929, 2009
2. — *Western European Software-as-a-Service Forecast, 2009–2013*, David Bradshaw, Apr 2009 - Doc # LT02R9, 2009
3. **General Services Administration US - GSA** [Online]
http://www.gsa.gov/Portal/gsa/ep/contentView.do?pageTypeId=8199&channelId=-24825&P=&contentId=28477&contentType=GSA_BASIC
4. **PCI Security Standards Council** [Online]
https://www.pcisecuritystandards.org/security_standards/pci_dss.shtml
5. **NIST** [Online] <http://csrc.nist.gov/groups/SNS/cloud-computing/index.html>
6. **Wikipedia** [Online] http://en.wikipedia.org/wiki/Cloud_computing
7. **Craig Balding** *cloudsecurity.org*. [Online] <http://cloudsecurity.org/2008/07/21/assessing-the-security-benefits-of-cloud-computing/>
8. **SUN - Project Kenai** [Online]
http://kenai.com/projects/suncloudapis/pages>HelloCloud#Examining_the_Virtual_Data_Center
9. **EC - European Commission** [Online] http://ec.europa.eu/enterprise/policies/sme/small-business-act/index_en.htm
10. **ISO/IEC. ISO/IEC 27001:2008** *Information technology - Security Techniques - Information security risk management; Annex E: Information security risks assessment approaches*, 2008
11. **Wikipedia** [Online] http://en.wikipedia.org/wiki/Open_Virtualization_Format
12. **MITRE** [Online] <http://cwe.mitre.org/data/definitions/400.html>
13. **BBC** [Online] http://news.bbc.co.uk/2/hi/uk_news/scotland/glasgow_and_west/6089736.stm
14. **www.retailresearch.org** [Online] <http://www.retailresearch.org/reports/fightinternalfraud.php>
15. **NY Daily News** [Online] http://www.nydailynews.com/gossip/2009/08/23/2009-08-23_outted_blogger_rosemary_port_blames_model_liskula_cohen_for_skank_stink.html

16. **Enterprise Storage Forum** [Online]
<http://www.enterprisestorageforum.com/continuity/news/article.php/3800226>
17. **Electronic Discovery Navigator** [Online] <http://www.ediscoverynavigator.com/statutesrules/>
18. **Find Law** <http://technology.findlaw.com> [Online]
<http://technology.findlaw.com/articles/01059/011253.html>
19. **CBS 11 TV** [Online] <http://cbs11tv.com/local/Core.IP.Networks.2.974706.html>
20. **WIRED** www.wired.com/ [Online] <http://www.wired.com/threatlevel/2009/04/company-caught/>
21. **Samuel T King, Peter M Chen, Yi-Min Wang, Chad Verbowski, Helen J Wang, Jacob R Lorch**
SubVirt: Implementing malware with virtual machines. 2006
22. **Secunia** [Online] <http://secunia.com/advisories/37081/>
23. — [Online] <http://secunia.com/advisories/36389/>
24. **Kortchinsky, Kostya** <http://www.immunityinc.com> [Online]
<http://www.immunityinc.com/documentation/cloudburst-vista.html>.
25. **Ormandy, Tavis** [Online] <http://taviso.decsystem.org/virtsec.pdf>
26. **Thomas Ristenpart, Eran Tromer, Hovav Shacham, Stefan Savage** [Online]
<http://people.csail.mit.edu/tromer/papers/cloudsec.pdf>
27. **Gentry, Craig** [Online] <http://delivery.acm.org/10.1145/1540000/1536440/p169-gentry.pdf?key1=1536440&key2=6166986521&coll=GUIDE&dl=GUIDE&CFID=60359435&CFTOKEN=10086693>
28. **Schneier, Bruce** [Online]
http://www.schneier.com/blog/archives/2009/07/homomorphic_enc.html
29. **www.spywarewarrior.com** [Online] <http://www.spywarewarrior.com/uiuc/ss/revoke/pgp-revoke.htm>
30. **RSA Laboratories, PKCS#11** [Online] <http://www.rsa.com/rsalabs/node.asp?id=2133>
31. **Jun Zhou, Mingxing He** [Online] http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=4716141
32. **Clulow, Tyler Moore and Jolyon** [Online] <http://people.seas.harvard.edu/~tmoore/ifipsec-pres.pdf>

33. **Andrew Bechere, Alex Stamos, Nathan Wilcox** [Online]
<http://www.slideshare.net/astamos/cloud-computing-security>
34. **Wikipedia** [Online] http://en.wikipedia.org/wiki/Token_bucket
35. — [Online] http://en.wikipedia.org/wiki/Fair_queueing
36. — [Online] http://en.wikipedia.org/wiki/Class-based_queueing
37. **Devera, Martin** [Online] <http://luxik.cdi.cz/~devik/qos/htb/old/htbtheory.htm>
38. **Open Source Xen Community** <http://xen.org/> [Online]
39. **Common Criteria Recognition Agreement (CCRA)** <http://www.commoncriteriaportal.org/> [Online]
40. **OWASP** [Online] http://www.owasp.org/index.php/OWASP_Top_Ten_Project
41. — [Online] http://www.owasp.org/index.php/Category:OWASP_Guide_Project
42. **27001:2005, ISO/IEC Information technology -- Security techniques -- Information security management systems -- Requirements**
43. **27002:2005, ISO/IEC Information technology -- Security techniques -- Code of practice for information security management**
44. **Group, BSI BS 25999 Business Continuity**
45. **NIST Special Publication 800-53, Revision 2 Recommended Security Controls for Federal Information Systems**
46. **OWASP** [Online] http://www.owasp.org/index.php/Main_Page
47. **SANS Institute** [Online]
http://www.sans.org/reading_room/whitepapers/securecode/a_security_checklist_for_web_application_design_1389?show=1389.php&cat=securecode
48. **Software Assurance Forum for Excellence in Code (SAFECode)** [Online] <http://www.safecode.org>
49. **IEEE Standards Association** [Online] <http://standards.ieee.org/getieee802/download/802.1Q-2005.pdf>
50. **The European Privacy Seal** [Online] <https://www.european-privacy-seal.eu/>

51. **EDRI - European Digital Rights** [Online] <http://www.edri.org/edri-gram/number7.2/international-standards-data-protection>
52. **ISACA** [Online]
http://www.isaca.org/Content/NavigationMenu/Members_and_Leaders1/COBIT6/COBIT_Publications/COBIT_Products.htm
53. **Office of Government Commerce (OGC)** [Online] <http://www.iti-officialsite.com/home/home.asp>
54. **Luis M. Vaquero, Luis Roderó-Merino, Juan Cáceres, Maik Lindner** *A Break in the Clouds: Towards a Cloud Definition*
55. **Cloud Security Alliance**, Security Guidance for Critical Areas of Focus in Cloud Computing, April 2009, <http://www.cloudsecurityalliance.org/guidance/csaguide.pdf>
56. **Jericho Forum**, *Cloud Cube Model: Selecting Cloud Formations for Secure Collaboration, April 2009*, http://www.opengroup.org/jericho/cloud_cube_model_v1.0.pdf
57. **Gartner**, *Assessing the Security Risks of Cloud Computing*, June 2008, <http://www.gartner.com/DisplayDocument?id=685308>
58. **Data Liberation Front**, Google, <http://www.dataliberation.org/>

ANNEX I – CLOUD COMPUTING – KEY LEGAL ISSUES

- I. Five key legal issues have been identified which are common across all the scenarios:
 1. data protection
 - a. availability and integrity
 - b. minimum standard or guarantee
 2. confidentiality
 3. intellectual property
 4. professional negligence
 5. outsourcing services and changes in control.
- II. Most of the issues identified in this discussion are not unique to cloud computing. Indeed, customers of cloud computing services may find it helpful to use the legal analysis applied to other Internet services as a foundation upon which to base their legal analysis of the security risks posed by cloud computing. To avoid repeating prior analysis, we have focused on those aspects of cloud computing security that we believe present new legal challenges or material changes from the analysis applied to prior Internet technologies.
- III. We believe that potential customers of cloud services will be quite concerned about issues related to data protection. Accordingly, in this legal analysis we have focused on these issues in more detail than on others.
- IV. While this document sets out five key legal issues, a theme that is consistent across all scenarios and in all of the discussions about cloud computing is the need for cloud computing providers to have highly detailed and product-specific contracts and other agreements and disclosures, and for customers to carefully review these contracts. or related documentation Both parties should also pay close attention to service level agreements (SLAs), since consideration of many legal issues associated with cloud computing are resolved in, or at least mitigated by, SLAs.
- V. Before getting into legal details it is worth noting that the customers of cloud providers may vary in type (from private to public entities) and size (from SMEs to large companies) and, thus, in the extent to which they are in a position to negotiate. This is very relevant from the legal point of view, because the relationship between the cloud providers and their customers will be mostly regulated by means of contracts. Because of the lack of specific regulations, reciprocal duties and obligations will be set forth in either standard general terms and

conditions, unilaterally drafted by the cloud provider, and either (more commonly) simply accepted by the customers without modification or negotiated in specific agreements.

- VI. The following table summarises the three possibilities in terms of negotiating contracts and agreements between the customer and the cloud provider.

CLOUD PROVIDER	CUSTOMER
A) Large company – strong ability to negotiate contract clauses	SME – Weak or lacking ability to negotiate contract clauses
B) Both the customer and the provider have the ability to negotiate contract clauses	
C) SME – Weak ability to negotiate contract clauses	Large company or public administration - may negotiate contract clauses

Depending on the particular case (whether it is A, B or C), the way to tackle the issues identified in *subsection I* may differ significantly.

- VII. It is important to differentiate between the case of a small to medium sized organisation, which would make a choice between different contracts offered on the market, and a larger organisation, which would be in a position to negotiate clauses. It is foreseeable that the main commercial benefit of cloud computing will come from the fact that cloud computing will likely be a bulk or commodity service that can be bought at short notice or on a pay-per use basis (e.g., case A: large cloud provider - SME customer). This assumes standardisation of services and thus of legal conditions. Therefore, in the legal analysis of this paper, we describe the issues primarily from the perspective of the small-to-medium organisation which is assessing different contracts, SLAs, etc, offered on the market.

Nevertheless, there may be situations in which cloud computing services will be tailored to large customers, ie, large companies and public administrations (e.g., case B). This assumes specific, tailored contracts. Case C is likely to be less common. In that case, there will be space for negotiation as in case B. Larger organisations may however use the same considerations when negotiating contracts. For this reason we have included a discussion of recommendations for negotiation, where this is possible.

It is also worth noting that even where a customer cannot negotiate terms of a contract with a specific provider, the customer is still *free to select between alternative offerings on the market*. In the case of an SME, therefore, recommendations for specific contractual clauses should be understood in terms of preferences between offerings in the market.

VIII. The analysis that follows captures and highlights how these five selected key legal issues may be addressed across the three different bargaining scenarios set out in paragraph VI.

1. Data protection

This section deals with legal issues of data protection that will often arise with the use of a cloud computing service and aims at giving guidance based on the wording of the Directive 95/46/EC of the European parliament and of the Council of 24 October 1995 on the Protection of Individuals with regard to the Processing of Personal Data and on the Free Movement of such Data³ (hereinafter: the “Data Protection Directive”). However, as these issues will be directly governed by national laws implementing the Data Protection Directive, customers of cloud computing services are advised to re-examine these issues based on the applicable national law

Glossary

The following definitions are set out in the Directive 95/46/EC of the European parliament and of the Council of 24 October 1995 on the *Protection of Individuals with regard to the Processing of Personal Data and on the Free Movement of such Data* (hereinafter: the ‘Data Protection Directive’).

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

Sensitive data means personal data allowing the disclosure of racial or ethnic origin, religious, philosophical or other beliefs, political opinions, membership of parties, trade unions, associations or organizations of a religious, philosophical, political or trade-unionist character, as well as personal data disclosing health and sex life.

Processing of personal data (Processing) means any operation or set of operations which is performed upon personal data, whether or not by automatic means, such as collection, recording, organization, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, blocking, erasure or destruction.

³ The official text of the Directive 95/46/EC and the status of implementation is available at http://ec.europa.eu/justice_home/fsj/privacy/law/index_en.htm

Controller means the 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; where the purposes and means of processing are determined by national or Community laws or regulations, the controller or the specific criteria for his nomination may be designated by national or Community law.

Processor means a natural or legal person, public authority, agency or any other body which processes personal data on behalf of the controller.

Defining the issues

- 1.1. Considering that the services provided by cloud providers generally consist of email, messaging, desktops, projects management, payroll, accounts and finance, CRM, sales management, custom application development, custom applications, telemedicine, and customers' billing, personal data (including sensitive data) will be processed. This data may belong to a number of persons (data subjects), e.g., employees, clients, suppliers, patients and, more generally, business partners.
- 1.2 Given the fact that personal data is undoubtedly processed, it is relevant to understand exactly when the Data Protection Directive applies. Section 4 states: '1. Each Member State shall apply the national provisions it adopts pursuant to this Directive to the processing of personal data where: (a) the processing is carried out in the context of the activities of an establishment of the controller on the territory of the Member State; when the same controller is established on the territory of several member States, he must take the necessary measures to ensure that each of these establishments complies with the obligations laid down by the national law applicable; (b) the controller is not established on the member State's territory, but in a place where its national law applies by virtue of international public law; (c) the controller is not established on Community territory and, for the purposes of processing personal data, makes use of equipment, automated or otherwise, situated on the territory of the said Member State, unless such equipment is used only for the purposes of transit through the territory of the Community.'
- 1.3 From an analysis of Section 4 of the Data Protection Directive, it follows that:
 - a) the place where the controller is established is relevant to the application of the Data Protection Directive;
 - b) What is not relevant for the application of the Data Protection Directive is the place of processing of the personal data or the residence of the data subject.

- 1.4 The Data Protection Directive will then apply if the Controller is established in the EU and if the Controller is not established in the EU but uses equipment located in the EU for processing of personal data (e.g., data centres for storage and remote processing of personal data situated on the territory of a Member State, computers, terminals, servers), unless such equipment is used only for purposes of transit through the territory of the Community.⁴
- 1.5 Once it is determined that the Data Protection Directive applies, the next question is: Who is the Controller and who is the Processor? If the customer of the cloud provider determines the purposes and means of the processing of personal data he is the Controller and if the cloud provider processes personal data on behalf of his customer he is an External Processor.⁵ In fact, the classification as a Processor or Controller is very different regarding the compliance duties and obligations and related liabilities. In our analysis we assume that the customer of the cloud provider is the Controller and the cloud provider an External Processor.
- 1.6 The main duties and obligations for the Controller set forth in the Data Protection Directive are:
- a) processing the personal data according to the principles of Fairness, Lawfulness, Finality, Adequacy, Proportionality, Necessity and Data Minimisation (Section 6 of the Data Protection Directive);
 - b) Obtaining unambiguous consent from the data subject where point a. of article 7 of the 95/46 directive applies⁶.

⁴ For further guidance on the issue of establishment and use of equipment and determinants for the applicability of the Data Protection Directive see Article 29 Data Protection Working Party's opinions on online social networking and search engines – respectively Opinion 5/2009 on online social networking; Opinion 1/2008 on data protection issues related to search engines; available at: http://ec.europa.eu/justice_home/fsj/privacy/workinggroup/wpdocs/2009_en.htm.

⁵ External' because in the specific case the Processor is a subject outside the Controller's company/organization.

⁶Article 7 of the 95/46 directive states:
Member States shall provide that personal data may be processed only if:
(a) the data subject has unambiguously given his consent; or

- c) processing the personal data after having provided the data subject with the necessary information (Section 10 of the Data Protection Directive);
- d) guaranteeing the data subject the rights laid down in Section 12 of the Data Protection Directive - e.g., to obtain confirmation as to whether or not data relating to the data subject is being processed, to obtain information on the purposes of the processing, the categories of data concerned, the recipient or categories of the recipients to whom the data are disclosed; to rectify, erase or block the data processed in a way which is not compliant with the provision of the Directive; etc. – (Section 12 of the Data protection Directive);
- e) implementing appropriate technical and organizational security measures to protect personal data against accidental loss, alteration, unauthorised disclosure or access and against all other unlawful forms of processing (Section 17 of the Data Protection Directive);
- f) choosing a Processor that provides sufficient guarantees with respect to the technical security measures and organisational measures governing the processing to be carried out, and ensuring compliance with those measures;
- g) transferring of personal data to ‘third countries which do not ensure an adequate level of protection within the meaning of Section 25 (2) of the Data Protection Directive only in case the data subject has given the previous consent unambiguously to the proposed transfer or under the condition that other procedures are in place as per Section 26 (e.g., ‘Standard Contractual Clauses’ or – if the data are transferred to the United States – ‘Safe Harbor Principles’).

(b) processing is necessary for the performance of a contract to which the data subject is party or in order to take steps at the request of the data subject prior to entering into a contract; or

(c) processing is necessary for compliance with a legal obligation to which the controller is subject; or

(d) processing is necessary in order to protect the vital interests of the data subject; or

(e) processing is necessary for the performance of a task carried out in the public interest or in the exercise of official authority vested in the controller or in a third party to whom the data are disclosed; or

(f) processing is necessary for the purposes of the legitimate interests pursued by the controller or by the third party or parties to whom the data are disclosed, except where such interests are overridden by the interests for fundamental rights and freedoms of the data subject which require protection under Article 1 (1).

- 1.7 The data controller (in this analysis, the cloud customer) should provide the data subjects (end-users of the cloud customer) with all the mandatory information related to the data processing. The cloud customer will be required under the Data Protection Directive to inform their customers about the circumstances of the transfer to the cloud provider, the quality of the cloud provider (ie, external processor), and the purposes of the transfer. In fact, externalising the services mentioned above in 1.1 necessarily implies the communication and transfer of such data to third parties⁷, (ie, the cloud providers),⁸ who may be located in Europe but also in countries outside the European Economic Area (third-countries). These countries may not offer an adequate level of protection of personal data within the meaning of Section 25 (2) of the Data Protection Directive. It is crucial that those who collect data subject to the Data Protection Directive ensure that they understand the application of the Directive to the use and transfer of that data. In this respect, controllers not currently engaging in cloud computing are advised to seek informed consent from the data subjects to data processing and transfer outside the European Economic Area. Those currently engaged in cloud computing are advised to ensure that this consent has been procured and that it adequately describes the nature and extent of processing and transfer. The alternative would be to have in place one of the procedures set forth in Section 26 (e.g., ‘Standard Contractual Clauses’ or ‘Safe Harbor Principles’ – if the data is transferred to the United States and the cloud provider participates in such a programme). Actually, this second way may present some advantages on the transfer of data based on the consent of the data subject because such consent may be withdrawn at any time by the data subject.
- 1.8 It is recommended that the Commission clarify the application of Section 25 (2) of the Directive as it applies to the possible processing of data in countries outside the European Economic Area during its transfer from one cloud computing provider to

⁷ Note that in some national laws (e.g. the German Data Protection Act) the terms “transfer” and “third party” are defined legal terms which carry specific legal implications. The use of these terms herein is not intended to carry such implications.

⁸ Unfortunately, there does not seem to be an official definition of transfer of data. However, from Section 4 of the Directive 95/46/EC it may be possible to gather that transit of data through the territories is not relevant from the legal point of view. For example, if data are transferred from the UK to the US, whether the data flows through network links that run via Iceland, Greenland and Canada seems to be irrelevant from the legal point of view.

another, or within the one company's cloud, should that cloud be located in multiple jurisdictions with one jurisdiction outside the European Economic Area.

- 1.9 All the parties involved in the data processing (data subjects, controllers and processors) should understand their respective rights and obligations relating to the processing of data as defined in the Data Protection Directive and the relevant statutory instruments by which the Directive has been implemented in the various EU Member States.⁹ Furthermore, these parties should also understand the right for respect for private life as set out in Article 8 of the European Convention on Human Rights and Fundamental Freedoms where the countries involved are signatories to the Convention or have implemented enabling domestic legislation.
- 1.10 To apply the Data Protection Directive adequately, the availability and integrity of data are key, leading the discussion to data security measures. There are unavoidable trade-offs here. More data security is likely to lead to reduced availability. The customer of the cloud provider may thus want to take a close look at the security measures the cloud provider has in place and the data availability guaranteed. It has to be born in mind that in most European countries there are mandatory data security requirements. The customer of the cloud provider will have to make sure those measures are complied with. In some cases (eHealth and, possibly, Resilience scenarios, when sensitive data and financial data are processed) the customer may even want to ensure even stricter data security measures as to the storage of data, communication or transfer of data, data disaster recovery and onward transfer.
- 1.11 It has to be clear at this point that the customer – when classified as sole data Controller - will be the entity responsible for the processing of personal data in relation to the data subjects. The customer will also be responsible for this data when such processing is carried out by the Cloud Provider in the role of external Processor. Failure to comply with the Data Protection Directive may lead to administrative, civil and also criminal sanctions, which vary from country to country, for the data controller. Such sanctions are mainly detailed in the relevant statutory instruments by which Directive 95/46/EC has been implemented in the various EU member States.

Dealing with the issues

⁹ See Status of implementation of Directive 95/46 on the Protection of Individuals with regard to the Processing of Personal Data, available. Available at:
<http://ec.europa.eu/justice_home/fsj/privacy/law/implementation_en.htm#italy>.

- 1.12 The issues discussed above may all be dealt with contractually. Apart from themselves ensuring to collect any personal data in compliance with Sections 7 and 10 of the Data Protection Directive, i.e. having previously duly informed the data subjects and obtained their consent (if required by Section 7), cloud customers should look for the presence of a Data Protection clause in the contract between the customer and the cloud provider. This clause should set forth the relevant parties' duties and obligations. The cloud customer should consider the following in evaluating such clauses:
- a. Bearing in mind that the cloud customer is classified as a data controller under the provisions of EU data protection law the customer is legally responsible for fairness, lawfulness, finality, etc..., clauses should be sought which support the customer's compliance with the principles of the Data Protection Directive.
 - b. The cloud provider should cooperate with the controller in order to assure that the latter can effectively guarantee the data subject's rights in accordance with Section 12 of the Data Protection Directive.
 - c. The cloud provider should have in place adequate security measures pursuant to Section 17 and the cloud provider should promptly notify the controller of any breach of data security and cooperate swiftly to solve the problem.
 - d. Possible transfers of personal data to third countries which do not ensure an adequate level of protection within the meaning of Section 25 (2) of the Data Protection Directive should be done on the basis of either prior unambiguous consent from the data subject to the proposed transfer or other procedures in accordance with Section 26 (e.g., 'Standard Contractual Clauses' or 'Safe Harbor Principles' – if the data are transferred to the United States and the cloud provider participates in such a programme). It should be borne in mind that cloud computing may consist of data transfers. It may be difficult to address this issue contractually. We recommend that the issue be addressed by the European Commission.
- 1.13 **Note that in case A** (see *Introduction*, paragraph VI), the contract, including the data protection clause, will be drafted by the cloud provider because of the impossibility of negotiating contractual clauses between a large provider and numerous small customers. So the potential customer should carefully analyse the provision to determine whether the clause gives the customer sufficient guarantees of lawful data processing by the cloud provider. and adequate remedies for contractual damages

- 1.14 **In cases B and C** (see *Introduction*, paragraph VI), the data protection clause will be subject to negotiation. In addition, security measures may be addressed in annexes and SLAs. In addressing security issues, the parties should keep in mind that they may not be able to detail all security measures to be addressed. Because IT security is an ongoing race to deal with new issues, contract terms need to be free to develop accordingly.
- 1.15 **In cases B and C (high value contracts with possibility to negotiate)**, it may also be advisable for the customer to negotiate adequate remedies for contractual damages should the Data Protection clause be breached. Last but not least, if the cloud provider's breach is substantial it may be included in the list of instances which lead to unilateral termination of the agreement
- 1.16 If the cloud provider is in a country outside the European Economic Area and that country does not offer an adequate level of data protection, it is advisable to have in place procedures in accordance with Section 26 (e.g., 'Standard Contractual Clauses' or 'Safe Harbor Principles' – if the data are transferred to the United States and the cloud provider participates in such a programme), rather than basing the transfer on the consent of the data subject (for the reason pointed out in *subsection 1.7*). However, it has to be stressed that the transfer of data within the territory of Member States is not without problems. Indeed, despite the fact that personal data can freely circulate within Member States, the laws are not consistent across countries. This inconsistency may create obvious difficulties in compliance and thus liability issues. We recommend that the Commission take steps towards the standardization of minimum data protection requirements in Europe. This is particularly important in the light of the fact that the Data Protection Directive is currently under revision. Moreover, a data protection certification scheme based on minimum data protection standards, which are common across the Member States, may be extremely useful.

2. Confidentiality

Defining the issues

- 2.1 Confidentiality concerns are also raised by the scenarios considered by this paper. In fact, secret information and 'know-how' may be processed in clouds. Any leakage of information caused by voluntary communication by the Cloud Provider or clouds' security breach may jeopardise customer business/services. NB in this context, It is crucial to distinguish between processing of data as in computational operations over that data, and the storage or transmission of data without altering it, since processing in this sense usually requires the data to be in unencrypted form.
- 2.2 Having a closer look at the concept of know-how and the possible ways to protect it seems worthwhile.

Know-how is defined as a body of information that is secret, substantial and identified in any appropriate form.¹⁰ The term 'secret' means that the know-how package as a body, or in the precise configuration and assembly of its components, is not generally known or easily accessible. The term 'identified' means that the know-how is described or recorded in such a manner as to make it possible to verify that it fulfils the criteria of secrecy and substantiality. To this purpose 'substantial' means that the know-how includes information which is of importance for the whole or a significant part of:

- i a manufacturing process, or
 - ii a product or service, or
 - iii for the development thereof and excludes information which is trivial.
- 2.3 There do not seem to be any European regulations applicable to such scenarios. European regulations regarding know-how, set out in the definition above, apply principally to licensing and activities involving the transfer and exploitation of information.

Dealing with the issues

- 2.4 Keeping regulations in mind, and in order to preserve the economic value of know-how and secret information in general, including research results, customer and

¹⁰ See Commission Regulation (EC) No 772/2004 of 27 April 2004 on the application of Article 81(3) of the Treaty to categories of technology transfer agreements.

project-related information, we recommend that customers seek contractual terms covering this issue. In fact, parties' duties and obligations to preserve such value should be specifically addressed in a 'confidentiality/non-disclosure clause'. Particular attention should be given to the boundaries of the responsibilities of parties and related liabilities. Technical annexes may be particularly effective places to address this issue.

- 2.5 **In case A**, the potential customer of the cloud provider should carefully analyse the confidentiality/non-disclosure clause to determine whether the cloud provider offers sufficient guarantees to protect the customer's secret information and know-how that will circulate in the cloud.
- 2.6 **In cases B and C**, we recommend that the parties negotiate a provision that reflects the damage a party may sustain should confidential or secret information be disclosed. If the disclosure is substantial, this breach may be included in the list of instances which allow the company to unilaterally terminate the agreement.

3. Intellectual property

Defining the issues

- 3.1 Intellectual property may also be at risk in the cloud computing scenarios.
- 3.2 Although an entity outsourcing services to the Cloud Provider may protect and enforce its intellectual property rights by means of the relevant legislation, which is similar in all the European Member States, a breach of Intellectual Property rights may cause immediate damage which will never be fully restored in a legal proceeding.
- 3.3 Moreover, in the unlikely case that the interactions between the customer and the cloud provider – for example, in the negotiation phase which will be possible in case B or C - may give rise to joint results which can be object of intellectual property rights (for example, techniques to better handle data). Therefore, it is wise to determine who will own these rights prior to engaging in cloud computing activities, and further determine the use that the parties can make of the objects of such rights.

Dealing with the issues

- 3.4 Intellectual Property rights should be regulated through dedicated contractual clauses: “Intellectual Property Clause” and “Confidentiality/Non Disclosure Clause”¹¹.
- 3.5 **In case A**, the potential customer of the cloud provider should carefully assess the value of its intellectual property and the risks related to cloud computing services. Having done so, the customer should carefully review any clauses governing intellectual property to determine whether the cloud provider offers sufficient guarantees and allows the customer appropriate tools to protect its information (e.g. through encryption of data), to protect the customer’s assets. The cloud customer should ensure that the contract respects their rights to any intellectual property as far as possible without compromising the quality of service offered (e.g. the creation of backup copies may be a necessary part of offering a good service level).
- 3.6 **In cases B and C**, the ‘intellectual property clause’ should be detailed enough to provide clear rules to address the issues described in paragraph 3.3 above. Moreover, it is advisable that the customer negotiate a clause in which the cloud provider is penalized should the provisions governing intellectual property be violated. Substantial breaches by the cloud provider may be included in the list of instances allowing the company to unilaterally terminate the agreement.

4. Professional negligence

Defining the issues

- 4.1 Failures in the services outsourced to the cloud provider may have a significant impact on the customer’s ability to meet its duties and obligations to its own customers. The customer may thus be exposed to contractual and tortious liability to its customers based on negligence.
- 4.2 Failures by the cloud provider may also result in liability by the customer to its employees. Because the customer is outsourcing technology that provides, or supports, critical internal functions such as email, messaging, desktops, project management, and payroll services, failure of the cloud provider, and the resulting inability of customer employees to access these functions or the data processed by them, may lead to customer liability to its employees.

¹¹ As to the “Confidentiality/Non Disclosure Clause,” paragraph 2.4 above applies.

- 4.4 A related issue is whether the terms of the contract attribute responsibility to the customer for any illegal acts carried out using the account which are authenticated by the customer's credentials but not actually carried out by the customer.

Dealing with the issues

- 4.5 **In case A**, the customer should carefully review the (standard) limitation/exclusion of liability clause in favour of the cloud provider to check whether it is sustainable.
- 4.6 **In cases B and C (i.e. in the rarer case where high value contracts are negotiated)** we recommend that the customer shift its liability for the issues mentioned above, as far as possible, towards the cloud provider if this is possible without incurring higher costs due to that liability shift. This may be accomplished by means of "Limitation of Liability" and "Indemnity" clauses. Substantial breaches by the cloud provider may be included in the list of instances allowing the customer to unilaterally terminate the agreement. It should be noted, however that the data controller always remains legally liable according to the provisions of the Data Protection Directives (1) (1), in respect of damages to data subjects, independently of any contractual clauses.
- 4.7 We recommend that legal clarification be provided to the European Community how the intermediary liability exemptions of the eCommerce directive apply to cloud providers

5. Outsourcing services and changes in control

Defining the issues

- 5.1 The agreement between the company and the cloud provider is likely to be defined as a contract *intuitu personae*. An *intuitu personae contract* is one in which a party chooses to contract with a company based on qualities that are unique to the company. For example, a customer may choose a particular cloud provider because of the conditions it offers, its reputation or professionalism, or its technical skills. As a result, the customer may be reluctant to see the cloud provider outsource all or part of the services to be provided to the customer.
- 5.2 The control of the cloud provider may also change and, as a result, the terms and conditions of the services provided by the cloud provider may change, too.

Dealing with the issues

- 5.3 **In case A**, we recommend that the customer determine whether services will be outsourced by the cloud providers and whether the cloud provider issues some guarantees or warranties relating to the performance of the services outsourced. However, we do not recommend that the customer look to be able to restrict the outsourcing of services by the cloud provider. We also recommend that the contract be reviewed to determine how the cloud provider will communicate changes in control to the customer. The customer may also want to consider whether the contract includes the right to terminate the contract if a change in control occurs.
- 5.4 **In cases B and C**, the customer *may* choose to require that the outsourcing of services by the cloud provider be subject to the customer's prior authorisation. To make this decision, the customer will need to be informed about the type of services that the cloud provider intends to outsource and the identity of the company to whom these will be outsourced. Even if the customer agrees to the outsourcing, it may want the cloud provider to issue some guarantees or warranties relating to the performance of the services outsourced. By the same line of reasoning, the customer may also want to have the chance to approve a change of control, or to terminate or renegotiate the contract in case of a change in the control of the cloud provider. Such options may be carefully specified in the contract between the company and the cloud provider by means of a 'third-party outsourcing' clause, a 'warranties and indemnification' clause, a 'change in control' clause, or a 'termination of agreement' clause – again depending on the bargaining power of the parties.

Conclusions

All the contractual clauses from section 1 to section 3 may be suitable for standardisation, except the relevant penalties, which depends on the parties' bargaining power. Whereas the inner content of the contractual clauses in sections 4 and 5 depends itself on the bargaining power of the parties, they are less suitable for standardisation.

ANNEX II – SME USE-CASE SCENARIO

AN SME PERSPECTIVE ON CLOUD COMPUTING

ENISA cloud computing security risk assessment

This scenario was used as a basis for the risk analysis published in the report.

Limitations and assumptions

This scenario is partially based on the results of the survey: An SME perspective on Cloud Computing [REF]. The scenario is NOT meant to be a road map for a company considering, planning or running cloud computing projects and investments.

The selection of a medium-sized company as a use-case was made to guarantee to the assessment a high enough level of IT, legal and business complexity. The aim was to expose all possible information security risks. Some of those risks are specific to medium-sized business, others are general risks that every micro, small and medium enterprise will likely face when migrating to a cloud computing environment.

The scenario is NOT intended to be completely realistic for a single organisation but all elements of the scenario are likely to occur frequently in many organisations; currently there is no single provider in the market that can cover the breadth of services described in the scenario, but all the services will be covered by several providers.

The assignment of applications to each tier (IaaS, PaaS, SaaS) is arbitrary and is done for illustrative purposes only and is NOT a recommendation.

Scenario

The company CleanFuture works in the photovoltaic business. The company produces and supplies complete solar and photovoltaic systems and key components for solar systems and heating. The company was founded in 1999 in Germany, where the main production site is located. Since then CleanFuture has been a fast growing company and turnover has been increasing at an average of 20% per year.

In 2003, a branch office was opened in Spain and, in 2004, new offices were opened in Italy. During 2005, a decision to relocate the business line producing anti-reflective solar glass to Poland was taken and by June 2006 that factory was already producing the first products. The company is also planning to explore the USA market.

Clean Future employs 93 people:

- 50 in Germany (2 different sites: headquarters (including production site, laboratory and 1 branch office)
- 34 in Poland
- 5 in Spain
- 4 in Italy.

The company also has a variable number of contractors (from 10 to 30 interim agents, sales representatives, consultants, trainees, etc).

Due to competitive pressure and the economic and financial crisis of 2008-2009, CleanFuture began an internal discussion on a near-term strategy to reduce costs and increase productivity. IT services were identified as a crucial area with a large margin for improvement.

An internal analysis was performed on IT and security requirements and the following conclusions were drawn:

1. More flexibility and scalability are needed to respond to variable demands for IT services (a variable number of employees during the year, a variable number of partners and suppliers to be dealt with, sudden changes in the market landscape, possible cooperation with a research centre and universities, possible opening of branch offices and enlargement of the sales force, etc).
2. High quality IT services (in terms of effectiveness and performance) and a high level of information security (in terms of availability, integrity and confidentiality) are required by the company. However, in order to provide internal resources (IT Dept) with such high levels of service, specific expertise is needed along with capital investment in hardware, software, IT support and information security.
3. Business continuity and disaster recovery capabilities need to be improved.
4. A test-bed for assessing new applications to support the business, as well as a cooperation environment where developers can work together with partners towards new solutions and projects, would be extremely important from the perspective of business efficiency and the capacity to innovate.
5. A physical-to-virtual (P2V) migration project would give important feedback in terms of the reliability and efficiency of the final set-up.

The services and applications identified as the ones to be affected by the new IT approach were:

- email and messaging
- desktop (office applications)
- project management
- payroll

- CRM and sales management
- accounting and finance
- running or hosting custom applications, and the development of custom applications
- identity management.

The internal working group, supported by an external consultant, proposed cloud computing technologies as a possible solution for CleanFuture's needs.

As a next step, a feasibility study on cloud computing was carried out. A report was delivered to the management board of the company: 'CleanFuture – A feasibility study on Cloud Computing: a possible implementation strategy and related business, legal and security concerns'.

Based on the analysis of the ad hoc working group, the report proposes that the identified IT services and applications be outsourced to at least three cloud providers. In the long run the three providers could constitute a so-called 'federation of clouds', but for the time being it is advisable for the sake of simplicity to use three independent providers linked throughout a 'federated identity management' solution.

1. Cloud Provider #1: will offer a cloud-based hosting service for email, messaging, desktop environments, project management and payroll (ie, a 'software as a service' (SaaS) model of cloud computing). Contractually, the data may be located and processed in different locations globally, including Asia, Europe and the USA.
2. Cloud Provider #2: will offer a cloud-based platform for custom applications hosting, often termed as the 'platform as a service' (PaaS) model of cloud computing. This custom application consists of a 'simulator' that helps customers to self-configure a photovoltaic installation, calculate the energy production (according to their geographical location) and the ROI (according to the incentive of the country where the installation will be done).
3. Cloud Provider #3: will offer a cloud-based infrastructure for HR, accounting and finance, CRM and sales management and custom application development (ie, the 'infrastructure as a service' (IaaS) model of cloud computing).

In the short-term (two years) CleanFuture will take care of business continuity and disaster recovery of data and services outsourced to PaaS and IaaS providers. That will be done by using the existing infrastructure. The SaaS provider assumes responsibility for backup and business continuity requirements for the services they provide. In both cases, the backup service is provided by both the provider and CleanFuture for a period of two years.

The strategic medium-term plan for disaster recovery is still to be defined. The following two options are to be compared:

- I. to identify a business partner with whom to create a small private cloud and share the capabilities and cost of such an infrastructure;

II. to buy business continuity and disaster recovery services from each cloud provider. The decision will be taken within two years when the current IT infrastructure will be obsolete. Until then, CleanFuture will use their in-house technology and on-site hosting for their continuity and recovery needs.

Identity management

The report recognises identity management as a component that affects all aspects of the migration. For reasons of reliability and scalability, CleanFuture should NOT rely in the long-term on an internal company directory for user authentication and account management. A scalable, resilient and future-proof solution must provide:

- a. single sign-on
- b. single sign-off
- c. a single identity directory for all services
- d. a single application for identity provisioning and de-provisioning
- e. secure management of any cryptographic keys used for authentication and signature
- f. access control policy enforcement (e.g., using XACML). A solution to guarantee that ALL the users (staff members, partners, contractors) comply with the company's security baseline requirements. These requirements are to be established according to the characteristics of the user profile and permission. The minimum requirements settled should be: updated antivirus and updated OS.

The report recommends moving to a federated identity management solution which will decouple the various accounts needed on different solution providers from identity provisioning and management services. A brief survey shows that few existing cloud solutions provide the interfaces required for a complete FIM solution. This leads to a set of important requirements on the migration:

1. Services selected should support authentication via a selected FIM framework (implementation using Liberty/Cardspace + SAML 2.0).
2. Before migrating any services and applications to the cloud, CleanFuture should implement a single-sign-on solution for all their applications, including authentication of external partners.
3. Trust properties of any key management infrastructure should be strongly verified.
4. A security client health baseline should be defined for all clients accessing all services.

Project	Phase 1 – 2008	Phase 2 – 2009	Phase 3 – 2010	Phase 4 – 2011	Phase 5 – 2012
Physical to virtual (P2V) migration	Adopt a virtualization platform inside the company, perform a physical to virtual (P2V) migration of the following applications: CRM and Sales Management, Custom Application, HR	Verify the reliability and performance of the solution in phase 1. P2V migration of the following application: Finance and Accounting Selection of FIM and key management solutions.	Verify the reliability and performance of the solution in phase 2. Migration to FIM SSO solution and key management solution.		
Migration to cloud computing: PROVIDER #1 - SaaS			Selection of the cloud provider (SaaS) and migration of the following application: Project Management*	Migration of the following applications and services: Email*, Messaging*, Desktops*, Payroll*	
Migration to cloud	----		Selection of the cloud	Verify the reliability and	Verify the reliability and

BENEFITS, RISKS AND RECOMMENDATIONS FOR INFORMATION SECURITY

<p>computing: PROVIDER #2 - PaaS</p>			<p>provider (PaaS) and migration of the following applications: CRM and Sales Management</p>	<p>performance of the PaaS Provider. Migration of the following applications: Custom Applications, Account and Finance, and HR</p>	<p>performance of the PaaS Provider. Migration of the following application: Custom Application Development</p>
<p>Development of private cloud partner for DR and BC</p>			<p>Partner identification _Project requirements definition...</p>	<p>Executive plan definition...</p>	<p>Private cloud goes live.</p>

*Please note that those applications or services were outsourced to a cloud computing provider without an internal physical to virtual migration being performed.

Existing security controls

Provider #1 (SaaS) and Provider #2 (PaaS) claim to implement a set of standard security controls which include:

- firewall
- IDS/IPS (Network and Host based)
- system hardening and in-house penetration testing
- ITIL compliant incident and patch management.

No further details are given. The selection of the providers was done by Clean Future on the basis of the good reputations of Provider #1 and Provider #2.

Provider #3 (IaaS) offers pre-configured VM instances in various standard configurations. They do not however offer pre-hardened instances by default, ie, the customer is entirely responsible for all security measures on VM instances including review of all default settings.

Provider #3 specifies background checks on all employees (with some limitations according to local laws), physical access control based on biometric smart-cards and need-to-know based data access control policies.

All connections (for IaaS, PaaS, SaaS, IDM, etc), EXCEPT those with customers (e.g., using the configuration application), are encrypted (either via VPN or SSH).

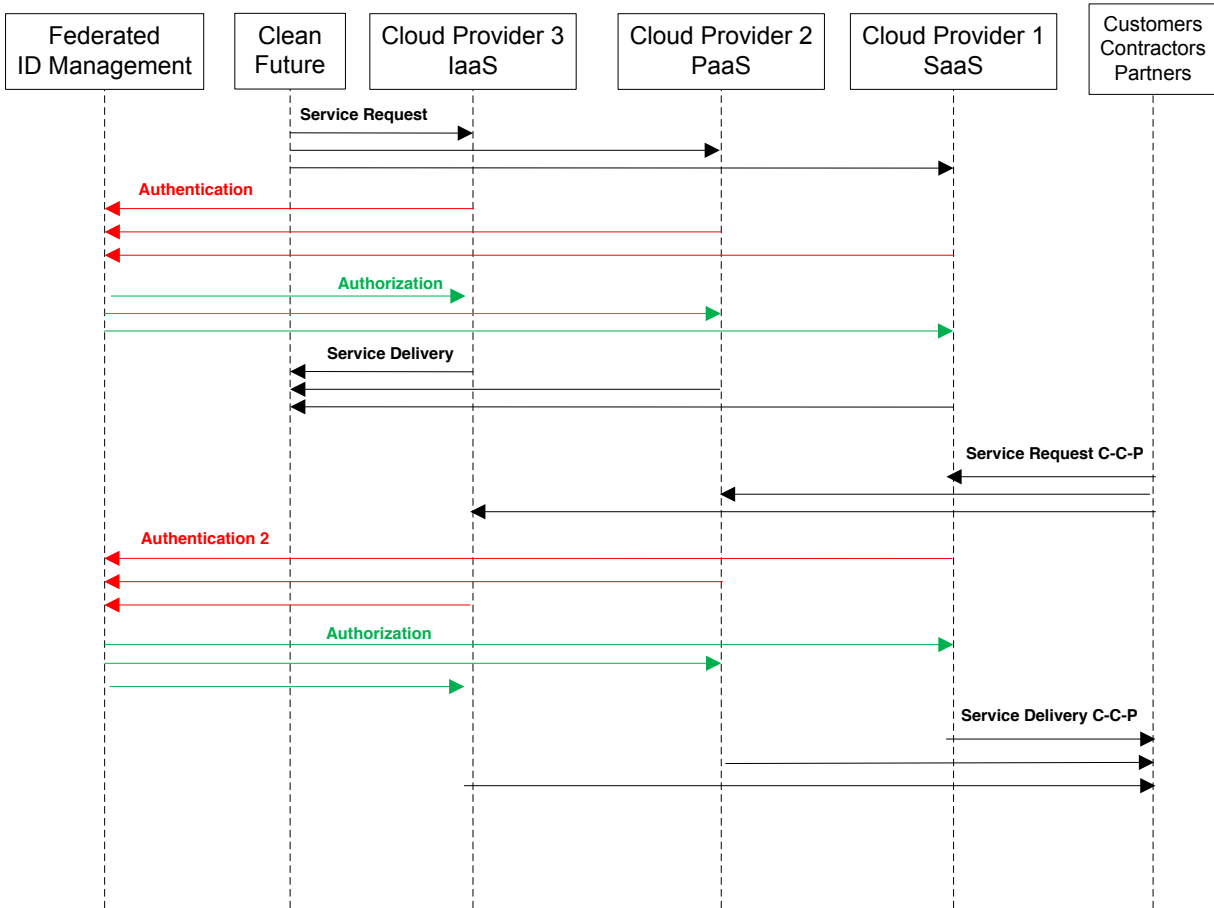
All providers are ISO 27001 compliant but none of them declare the exact scope of the certification.

The SLA with each provider includes a breach notification clause. All the providers offer premium (paid) security reporting features. Such paid reports may include: failed breaches (of the customer's assets), attacks against specific targets (per company user, per specific application, per specific physical machine, ratio of internal attacks compared to external attacks, etc), trends and statistics.

The reporting threshold for failed attempts and the scale for incidents severity are customized according to the customer's specific needs.

BENEFITS, RISKS AND RECOMMENDATIONS FOR INFORMATION SECURITY

Data flow



ANNEX III – OTHER USE-CASE SCENARIOS

Here you can find a brief summary of the Resilience and eHealth scenarios we used in our risk analysis.

RESILIENCE SCENARIO

This scenario explores how the use of cloud computing affects the resilience of services in the face of:

- sudden increases in customer demand (e.g., in periods of financial crisis)
- denial of service attacks
- localised natural disasters
- misuse of the infrastructure as an attack platform
- data leaks (malicious or careless insider or broken process).

The year is 2012. XK-Ord provides real-time e-commerce through a web-service interface, along with content delivery solutions in the form of widgets, which can be embedded into purchasing portals.

Typical use-cases are:

- real-time price data and charts for goods in purchasing portals
- historical data for use in price-prediction and analysis
- order histories and stock control reports for companies
- real-time currency conversion and FX histories
- up to date SOX & EU anti-monopoly trading reports
- financial data for more complex applications.

In addition, XK-Ord provides a platform for managing the various services and combining them into custom applications. Given their service offering, XK-Ord requires high resilience of:

- latency – delays in data provision can lead to high-value deals being lost;
- request fulfilment – e.g., highly reliable:
 - database queries and results presentation
 - web server fulfilment of http requests
 - TCP/IP infrastructure;
- data integrity – errors in data can lead to financial loss;
- confidentiality and reporting – data has financial value – therefore, if disclosed to non-paying customers, it represents a financial loss to XK-Ord;
- application integrity and vulnerabilities.

Infrastructure

In 2011, XK-Ord moved to a cloud infrastructure for reasons of cost, flexibility and reliability. XK-Ord uses CumuloNimbus Systems, a cloud provider offering IaaS for content delivery.

- Data is stored using a DaaS (database as a service) model.

- CRM and XK-Ord's customer account management, including billing, are managed by a second cloud provider, Stratocumulus. Credentials are issued and verified by XK-Ord using this service, while control access to the content is provided by CumuloNimbus resources, ie, Stratocumulus is acting as a federated identity provider providing single sign-on.
- XK-Ord's HR, payroll, office desktop applications and R&D systems are managed directly by XK-Ord and hosted on site by XK-Ord.

Network

Compared to using a data centre, the cloud providers' infrastructure offers considerable improvements in total bandwidth, processing, memory and storage capabilities, as well as the ability to scale limits quickly. For example, routers located near content delivery locations use scalable virtualised memory, logging and packet filtering resources. IPSec is implemented in parts of the network. These features considerably improve resilience against DDoS attacks.

Resource management

- Content delivery is charged to XK-Ord on a per-HTTP-request basis. Costs are capped according to a choice of policies offered by CumuloNimbus. XK-Ord customers are charged according to the number of HTTP requests to each service according to a different scheme.
- The provider operates resources on a shared-tenancy basis with other (not necessarily similar) clients over their entire infrastructure. This means that isolation between resources used by different customers must be strong. XK-Ord has the option to pay a small fee to reserve resources in advance, which increases overall reliability for the service provider and for XK-Ord.
- Short-term growth within available resources is faster than typical non-cloud infrastructures. Adding more expansion resources (ie, more hardware for the cloud provider is slow). DDoS defences in particular must scale quickly and implications for cost and resource-use should be well-defined.
- Standard SLAs and standard APIs are used to ease movement between clouds.

Security services

XK uses a security service provider, BorealisSec, for real-time security monitoring (RTM), vulnerability assessment and device management.

- BorealisSec staff manages XK systems hosted on Cumulonimbus using a VPN connection.

- Logs are collected on Cumulonimbus and sent automatically to the BorealisSec SIEM (security information and event management) platform via VPN, for analysis.

Should an incident be triggered, either:

- A BorealisSec administrator will deal with the incident directly (automatically or manually), or
- they will open a ticket with CumuloNimbus to solve the problem.

In any case, response to incidents will be contractually agreed with XK, according to severity.

The following other points are worth noting:

- Vulnerability assessment can only be performed on a test installation since the Cumulonimbus ToU prohibits proactive security testing.
- BorealisSec provides compliance and audit reports as far as possible within the CumuloNimbus ToU.
- BorealisSec is responsible for maintaining patches to software outside the remit of the cloud provider.

SLA: XK-Fin -> customers

XK-Ord offers a service level agreement (SLA) to its customers in order to compete with other financial data companies offering SLAs. It is worth noting that XK's SLA may offer higher levels of reliability than Cumulonimbus, despite the dependency between the two. This may be because XK is willing to accept a higher level of risk.

Goal	KPI	Value	Penalties
Service availability	% Uptime per month	99.99	20% reduction in bill for every factor of 10
Latency (NB this is the time from when the stock market publishes data)	Average response time over 100 requests over 1 day	1 sec	5% reduction in bill for every violation
Administration	Time to respond to request in minutes	60 min	5% reduction in bill for every violation
Alerting	Minutes to alert customer of a violation of service (not including this one...)	5 min	5% reduction in bill for every violation
Time to recover from fault	Hours	2 hours	5% reduction in bill for every violation

EHEALTH SCENARIO

This scenario explores the use of cloud computing by large government bodies which have to satisfy strict regulatory requirements and are very sensitive to negative public perception. A key consideration – when using cloud services – will be a public perception that there has potentially been a lack of consideration of security or privacy issues. This would be especially true should ‘public’ cloud services be used.

EuropeanHealth represents a large government health service in Europe *but does not describe any specific national health service*. EuropeanHealth is composed of public organisations and private suppliers providing eHealth services. It is a very large organisation spread across several sites and it caters to 60 million citizens. Prior to using any kind of cloud infrastructure, it has over 20 IT service providers and more than 50 data centres.

Specific scenario

The specific scenario involves an eHealth platform that provides care and monitoring of patients with chronic illnesses in their homes. This general process is described in more detail as follows:

1. A monitoring centre uses an independent Internet-based platform deploying in-home sensors to monitor and interact with elderly patients at home.
2. The monitored variables are analysed for anomalies based on a profile. A monitoring centre decides when more specialized services are needed (doctors, nurses, etc).
3. Patients may also choose to make information available to external eHealth service providers. Such private information is provided via a centralized database.
4. Services are provided to elderly patients at home using a multimodal interface that adapts to the abilities of the elderly. Avatars and speech synthesis can be used.

The monitored data is available to doctors and hospitals through the unique patient medical record service. Patient information can be accessed through the unique patient identifier. This service provides documentation of a patient's medical history and care.

Gov-cloud

To deliver these services using a cloud infrastructure, EuropeanHealth uses **Gov-Cloud**, a cloud service provided by national governments for government services as a whole. This is a hybrid private-partner cloud since it is used by trusted partners only and only government organisations have administrative access (e.g., public administration, healthcare). It uses a dedicated network infrastructure, which is physically independent of the public Internet. The Gov-cloud is hosted in multiple geographical locations but virtual machines can be migrated from one site to another.

All services in our specific scenario run on the Gov-Cloud with the security properties described below. For example:

- some of the services running at home are running on the cloud using IaaS;

- the services running at the monitoring centre are running on the cloud using IaaS;
- the monitored data is also stored in the cloud using DaaS (database as a service).

Gov-Cloud also provides a means of transferring patient data securely (previously it was quite difficult) using a customised email service for doctors and nurses. This is provided by a third party but designed by EuropeanHealth.

Data protection

All data collected by EuropeanHealth must satisfy the following requirements:

- Data (including sensitive personal information) must be encrypted in transit and at rest where potentially at risk (e.g., on mobile devices).
- Data processing must satisfy European data protection law (e.g., definition of 'data processor' for all operations).
- National law applies certain restrictions on the processing of the data (e.g., data should not leave the original country of collection at any time).
- Clinical safety has to be paramount with certain applications; this means that integrity and availability have to be 'guaranteed' in some instances.
- Sensitive data should be destroyed at a specified time in its lifecycle (e.g., by destruction of hard disks at the 'end of life' of equipment).
- Physical security controls in data centres where data is stored must be adequately assured (some of this is covered currently via ISO27001 submissions from suppliers).
- Senior staff is given special responsibility for the confidentiality of 'patient and service-user information'.

Compliance with laws, regulations and best practices

- All suppliers must demonstrate compliance with ISO27001. They are NOT required to be accredited but compliance is verified through yearly submission of their information security management system and associated policy documents.
- Additional certifications and accreditations assist EuropeanHealth organisations in choosing appropriate providers, e.g., ISO20000 (Service Management), ISO9001 (Quality), etc, but these are not required.
- In terms of audit and compliance to regulations or the nominated standards of the service provider, cloud computing service providers must ensure that they are able and willing to allow the right to audit their policies, processes, systems and services.

Governance

A basic set of security controls is provided by Gov-Cloud and additional controls are optionally provided by management services or in-house management for each user of Gov-Cloud (such as EuropeanHealth). Governance standards, such as ITIL, are used.

EuropeanHealth cannot mandate internal departments to adopt specific technologies, but only recommend technologies to be adopted. EuropeanHealth departments remain free to implement the technology that best meets their needs.

EuropeanHealth can request participating organisations to provide documentation showing that their recommendations have been followed, e.g., proof that all data on laptops is encrypted. For external suppliers, there are specific requirements for organisations to connect to the EuropeanHealth network and remain connected.

Access control and audit trails

EuropeanHealth provides Single Sign On (SSO) for their applications and services using smart cards as authentication tokens. EuropeanHealth organisations may use many other forms of authentication or multiple forms for different purposes (ie, single factor, two factors, biometric and so forth). Gov-Cloud third party suppliers interface with EuropeanHealth PKI using smart cards.

There are absolute requirements in terms of audit to ensure that it is clear who has accessed what personal data or sensitive personal data and for what purposes.

Service Level Agreements

SLAs would need to be contractual and embedded within any cloud service offering to EuropeanHealth organisations. The key will likely be 24/7 availability (but dependent on the type of service, application or data being hosted).

- A concern for EuropeanHealth organisations will be the potential loss of control that they will feel (e.g., of infrastructure, the services, the data and provisioning, etc). The ability of cloud service providers to prove that there is 'no' loss of control will be a key consideration for take up.