



UNIL | Université de Lausanne

Unicentre

CH-1015 Lausanne

<http://serval.unil.ch>

Year : 2014

A Confluence of Risks: Control and Compliance in the World of Unstructured Data, Big Data and the Cloud

SIMMS David

SIMMS David, 2014, A Confluence of Risks: Control and Compliance in the World of Unstructured Data, Big Data and the Cloud

Originally published at : Thesis, University of Lausanne

Posted at the University of Lausanne Open Archive <http://serval.unil.ch>

Document URN : urn:nbn:ch:serval-BIB_A2BDF9375EDF5

Droits d'auteur

L'Université de Lausanne attire expressément l'attention des utilisateurs sur le fait que tous les documents publiés dans l'Archive SERVAL sont protégés par le droit d'auteur, conformément à la loi fédérale sur le droit d'auteur et les droits voisins (LDA). A ce titre, il est indispensable d'obtenir le consentement préalable de l'auteur et/ou de l'éditeur avant toute utilisation d'une oeuvre ou d'une partie d'une oeuvre ne relevant pas d'une utilisation à des fins personnelles au sens de la LDA (art. 19, al. 1 lettre a). A défaut, tout contrevenant s'expose aux sanctions prévues par cette loi. Nous déclinons toute responsabilité en la matière.

Copyright

The University of Lausanne expressly draws the attention of users to the fact that all documents published in the SERVAL Archive are protected by copyright in accordance with federal law on copyright and similar rights (LDA). Accordingly it is indispensable to obtain prior consent from the author and/or publisher before any use of a work or part of a work for purposes other than personal use within the meaning of LDA (art. 19, para. 1 letter a). Failure to do so will expose offenders to the sanctions laid down by this law. We accept no liability in this respect.



UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES
DÉPARTEMENT DES SYSTÈMES D'INFORMATION

**A Confluence of Risks:
Control and Compliance in
the World of Unstructured
Data, Big Data and the
Cloud**

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Etudes Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Docteur en Systèmes d'Information

par

David SIMMS

Directrice de thèse
Prof. Solange Ghernaoui

Jury

Prof. Alessandro Villa, Président
Prof. Valérie Chavez, experte interne
Prof. Jean Henry Morin, expert externe
Dr Igli Tashi, expert externe

LAUSANNE
2014



UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES
DÉPARTEMENT DES SYSTÈMES D'INFORMATION

**A Confluence of Risks:
Control and Compliance in
the World of Unstructured
Data, Big Data and the
Cloud**

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Etudes Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Docteur en Systèmes d'Information

par

David SIMMS

Directrice de thèse
Prof. Solange Ghernaoui

Jury

Prof. Alessandro Villa, Président
Prof. Valérie Chavez, experte interne
Prof. Jean Henry Morin, expert externe
Dr Igli Tashi, expert externe

LAUSANNE
2014

IMPRIMATUR

Sans se prononcer sur les opinions de l'auteur, la Faculté des Hautes Etudes Commerciales de l'Université de Lausanne autorise l'impression de la thèse de Monsieur David SIMMS, licencié en allemand et irlandais de l'University College of Wales (Royaume-Uni), titulaire d'un master en informatique de l'University of Kent at Canterbury (Royaume-Uni), en vue de l'obtention du grade de docteur en Systèmes d'Information.

La thèse est intitulée :

**A CONFLUENCE OF RISKS:
CONTROL AND COMPLIANCE IN THE WORLD OF UNSTRUCTURED DATA,
BIG DATA AND THE CLOUD**

Lausanne, le 28 mai 2014

p.o. 
Le doyen

Thomas von Ungern-Sternberg

Directrice de thèse

Professeure Solange Ghernaoui - Professeur ordinaire, Faculté des Hautes Etudes Commerciales, Université de Lausanne

Membres du jury

Professeur Alessandro Villa - Professeur ordinaire, Faculté des Hautes Etudes Commerciales, Université de Lausanne ; Président du Jury

Professeure Valérie Chavez - Professeure assistante en PTC, Département des opérations, Faculté des Hautes Etudes Commerciales, Université de Lausanne ;
experte interne

Professeur Jean-Henry Morin - Professeur associé en Systèmes d'Information et Services Informationnels, Faculté des Sciences Économiques et Sociales, Université de Genève ; expert externe

Docteur Igli Tashi - Corporate Security, Kudelski Security ; expert externe

Université de Lausanne
Faculté des Hautes Etudes Commerciales

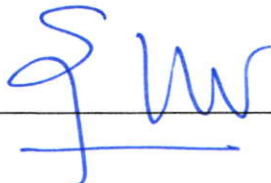
Doctorat en Systèmes d'Information

Par la présente, je certifie avoir examiné la thèse de doctorat de

David SIMMS

Sa thèse remplit les exigences liées à un travail de doctorat.

Toutes les révisions que les membres du jury et le-la soussigné-e ont
demandées durant le colloque de thèse ont été prises en considération et
reçoivent ici mon approbation.

Signature :  _____ Date : 22 mai 2014

Prof. Solange GHERNAOUTI
Directrice de thèse

Université de Lausanne
Faculté des Hautes Etudes Commerciales

Doctorat en Systèmes d'Information

Par la présente, je certifie avoir examiné la thèse de doctorat de

David SIMMS

Sa thèse remplit les exigences liées à un travail de doctorat.

Toutes les révisions que les membres du jury et le-la soussigné-e ont
demandées durant le colloque de thèse ont été prises en considération et
reçoivent ici mon approbation

Signature : _____



Date : _____

07.05.2014

Prof. Valérie CHAVEZ
Membre interne du jury

Université de Lausanne
Faculté des Hautes Etudes Commerciales


Doctorat en Systèmes d'Information

Par la présente, je certifie avoir examiné la thèse de doctorat de

David SIMMS

Sa thèse remplit les exigences liées à un travail de doctorat.
Toutes les révisions que les membres du jury et le-la soussigné-e ont
demandées durant le colloque de thèse ont été prises en considération et
reçoivent ici mon approbation

Signature : _____



Date : _____

15.05.2014

Prof. Jean Henry MORIN

Membre externe du jury

Université de Lausanne
Faculté des Hautes Etudes Commerciales

Doctorat en Systèmes d'Information

Par la présente, je certifie avoir examiné la thèse de doctorat de

David SIMMS

Sa thèse remplit les exigences liées à un travail de doctorat.

Toutes les révisions que les membres du jury et le-la soussigné-e ont
demandées durant le colloque de thèse ont été prises en considération et
'reçoivent ici mon approbation

Signature :  Date : 22.05.2014

Dr. Igli TASHI

Membre externe du jury

Acknowledgements

I would like to thank my thesis supervisor, Professor Solange Ghernaouti, for her guidance, advice and encouragement throughout the five years during which this thesis was in preparation.

I would also like to thank the colleagues and clients, throughout my career, whose suggestions, opinions and criticisms have helped shape my thoughts.

My greatest thanks, however, are due to my family, Sylvaine and Camille, for their encouragement, tolerance, patience and understanding.

Lausanne, 7 May 2014

Abstract

The emergence of powerful new technologies, the existence of large quantities of data, and increasing demands for the extraction of added value from these technologies and data have created a number of significant challenges for those charged with both corporate and information technology management. The possibilities are great, the expectations high, and the risks significant. Organisations seeking to employ cloud technologies and exploit the value of the data to which they have access, be this in the form of “Big Data” available from different external sources or data held within the organisation, in structured or unstructured formats, need to understand the risks involved in such activities. Data owners have responsibilities towards the subjects of the data and must also, frequently, demonstrate that they are in compliance with current standards, laws and regulations.

This thesis sets out to explore the nature of the technologies that organisations might utilise, identify the most pertinent constraints and risks, and propose a framework for the management of data from discovery to external hosting that will allow the most significant risks to be managed through the definition, implementation, and performance of appropriate internal control activities.

Keywords

Audit; security; compliance; data management; unstructured data; Big Data; internal controls; risk management; cloud computing; governance.

Résumé

L'émergence de nouvelles technologies puissantes, l'existence de grandes quantités de données, et le besoin de valoriser ces données et technologies, mettent les organisations face à de nombreux défis, tant pour leurs dirigeants que pour leurs responsables informatique. Le potentiel est grand, les attentes sont élevées, et les risques sont importants. Les organisations qui cherchent à utiliser les technologies du Nuage et à exploiter la valeur des données auxquelles ils ont accès, que ce soit sous la forme de « Big Data » disponibles depuis des sources externes différentes ou sous la forme de données déjà présentes au sein de l'organisation, dans des formats structurés ou non-structurés, doivent comprendre les risques associés à de telles activités. Les propriétaires de données ont des responsabilités envers les sujets des données et doivent fréquemment démontrer qu'ils sont en conformité avec des normes, lois et réglementations en vigueur.

Cette thèse a pour objectif d'explorer la nature des technologies que les organisations pourraient utiliser, d'identifier les contraintes et risques les plus pertinents par rapport à cette utilisation, et enfin de proposer un cadre pour la gestion de données, depuis leur identification jusqu'à leur exploitation dans le contexte d'hébergement externe, afin de gérer les risques les plus significatifs suivant un processus de définition, d'implémentation, et de performance des activités de contrôle interne.

Mots-clés

Audit; sécurité; conformité; gestion de données; données non-structurées; Big Data; contrôles internes; gestion de risques; informatique dans le nuage; gouvernance.

Contents

Acknowledgements	i
Abstract	ii
Keywords	ii
Résumé	iii
Mots-clés	iii
List of Figures	xi
List of Tables	12
Chapter 1 Introduction	13
1.1 Introduction	13
1.2 Description of the problem	13
1.2.1 Control and compliance	13
1.2.2 Risks and opportunities	14
1.2.3 Requirement for solutions	15
1.3 The structure of this thesis	15
1.4 Notes on vocabulary and style	16
Chapter 2 The Significance of the Problem	17
2.1 Introduction	17
2.2 A changing environment	17
2.3 Increasing quantities of data	19
2.4 Old certainties no longer apply	19
2.5 A cartography of risks	20
2.6 Specific risks related to unstructured data, big data, and the cloud	22
2.7 Consequences of the failure to address risks	23
2.8 Conclusion	25
Chapter 3 Analysis of Existing Research I: Unstructured Data	26
3.1 Introduction	26
3.2 Unstructured data: definitions and developments	27
3.3 Big data: definitions, significance, use and management	37
3.4 Data storage and processing	41

3.5	The Security of Unstructured Data.....	44
3.6	Privacy in the Digital Age.....	55
3.6.1	Definition	55
3.6.2	Responsibilities	56
3.7	Toxic data	58
3.8	Conclusion	59
Chapter 4	Analysis of Existing Research II: Cloud Computing.....	61
4.1	Introduction	61
4.2	Definition of terms	61
4.3	Risks and governance, opportunities and benefits	66
4.4	Security considerations	69
4.4.1	Visibility of third-party activities.....	70
4.4.2	Open standards and consistency	75
4.4.3	The benefits of outsourcing.....	78
4.4.4	Legal and regulatory issues.....	79
4.4.5	Service management processes	81
4.4.6	The overall benefits of cloud computing	82
4.4.7	Security aspects	83
4.4.8	Risk management	86
4.4.9	Cloud security design principles	87
4.4.10	Risk management	89
4.5	Security management processes	90
4.5.1	Security management standards	91
4.5.2	Security-focused internal controls.....	91
4.6	Information classification.....	92
4.7	User account management.....	94
4.8	Security awareness, training and education	95
4.9	Identity management.....	96
4.10	The importance of auditing and testing.....	97
4.10.1	Governance, Risk and Compliance	98
4.10.2	Penetration testing.....	99
4.11	Contingency planning.....	101
4.12	Conclusion.....	102
Chapter 5	Identification of Needs, Trends and Awareness: User and Market Surveys	104
5.1	Introduction	104
5.2	The situation in Switzerland : original research into current trends.....	105
5.2.1	Background to the survey.....	105

5.2.2	The Survey	107
5.2.3	Survey results	108
5.2.4	Interpretation and analysis of the results.....	109
5.2.5	Summary evaluations	111
5.2.6	Lessons learned	113
5.2.7	Further research issues suggested by the results and the limitations of the survey.....	114
5.2.8	Overall conclusion from the survey	114
5.2.9	Summary.....	114
5.3	Cloud Security Alliance/ISACA Survey 2012	115
5.3.1	Background to the survey.....	115
5.3.2	Relevance of this survey	115
5.3.3	Key findings and executive summary	115
5.3.4	Specific findings of interest.....	117
5.3.5	Overall conclusions from the survey	121
5.4	Comparison of the two surveys and their results	122
5.5	Overall summary of the surveys and their results	124
Chapter 6	Identification of Control Constraints: Management Theory and Security	125
6.1	Introduction	125
6.2	Generic risks and requirements	125
6.2.1	Security	125
6.3	Data security	131
6.3.1	Overall principles	132
6.3.2	Identity and Access Management (IAM)	133
6.3.3	Summary.....	137
6.4	Privacy Concerns	137
6.4.1	Key concerns and considerations	137
6.4.2	Key principles.....	138
6.4.3	Summary.....	140
6.5	Data lifecycle	140
6.5.1	Generation	140
6.5.2	Manner of use	140
6.5.3	Transfer.....	141
6.5.4	Transformation.....	141
6.5.5	Storage.....	141
6.5.6	Archival	141
6.5.7	Destruction	142
6.5.8	Summary.....	142

6.6	Conclusion	142
Chapter 7	An Analysis of Taxonomies and Models	143
7.1	Introduction	143
7.2	Security event taxonomy	143
7.2.1	Security Information and Event Management (SIEM)	144
7.2.2	Background to the taxonomy	145
7.2.3	Philosophy and structure	146
7.2.4	Description of the taxonomy	147
7.2.5	Complex security incidents versus basic security incidents.....	150
7.2.6	Comparison with other event classification models.....	151
7.2.7	Details of the ETSI model.....	153
7.2.8	Practical uses of the model.....	154
7.2.9	Summary.....	156
7.3	Security management and internal control standards.....	156
7.3.1	Security management standards	157
7.4	Conclusion	174
Chapter 8	The Need for Auditing and Compliance	175
8.1	Introduction	175
8.2	Auditing and compliance.....	175
8.2.1	Internal and external audit	175
8.2.2	Governance, Risk and Compliance.....	177
8.2.3	Summary.....	179
8.3	Legislative constraints.....	179
8.3.1	Sarbanes-Oxley Act	179
8.3.2	Auditing Standards	180
8.3.3	The impact of SOX on outsourcing and cloud computing.....	181
8.3.4	PCI DSS (Payment Card Industry Data Security Standard)	182
8.3.5	HIPAA (Health Insurance Portability and Accountability Act).....	182
8.3.6	Summary.....	182
8.4	Audit frameworks and approaches	182
8.4.1	Self-performed audit procedures	183
8.4.2	Service audit approach	184
8.4.3	Other procedures	186
8.5	Specifically IT based audit frameworks	187
8.5.1	SysTrust.....	187
8.5.2	WebTrust	187
8.5.3	ISO 27001.....	187

8.6	Summary	187
8.7	The Swiss context	187
8.7.1	Legal requirements	187
8.7.2	Audit Standards	192
8.7.3	Summary.....	195
8.8	Conclusion	195
Chapter 9	Identification of Key Technological Elements	196
9.1	Introduction	196
9.2	The use of cryptography	196
9.2.1	Encryption as an element of security	197
9.2.2	Basic cryptographic techniques	199
9.2.3	Symmetric and asymmetric encryption	200
9.2.4	Types of attacks on standing data, and possible counter-measures	202
9.2.5	Countermeasures	203
9.2.6	Data minimisation.....	204
9.2.7	Data stripping and separation	205
9.3	Directives.....	205
9.3.1	EU Data Protection Directive	205
9.3.2	e-Privacy Directive	206
9.3.3	Summary.....	207
9.4	Conclusion	207
Chapter 10	A high-level proposal to address the issues	209
10.1	Introduction	209
10.2	Understanding the concepts and related vocabulary	210
10.3	Scoping: types of data, uses of data, and storage options.....	210
10.3.1	Structured data.....	210
10.3.2	Unstructured data	211
10.3.3	Storage media.....	211
10.3.4	Central databases	211
10.3.5	File servers.....	211
10.3.6	Mail servers	211
10.3.7	Local hard drives.....	211
10.3.8	Backup media	212
10.3.9	Smart devices	212
10.3.10	Portable storage	212
10.3.11	Private cloud or remote storage services.....	212
10.3.12	Uses of the data	212

Contents

10.3.13	Quality of the data	212
10.3.14	Simplifying the scoping.....	213
10.4	Understanding the risks	213
10.4.1	Legal risks	213
10.4.2	Operational risks.....	213
10.4.3	Technical risks.....	213
10.5	Roles and responsibilities.....	213
10.6	Planning the transfer.....	214
10.6.1	Identifying the data to be moved	214
10.6.2	Cleaning the data and ensuring consistency	214
10.6.3	Formatting and portability	214
10.6.4	Determining what should happen to local copies of data being transferred.....	215
10.6.5	Determining the completeness, integrity and accuracy of the data transfer process	215
10.7	Handling other data appropriately.....	215
10.7.1	Reorganisation.....	215
10.7.2	Deletion	216
10.8	Addressing legal considerations.....	216
10.8.1	Confirming that no requirements or obligations are being contravened	216
10.8.2	Informing the subjects of data	217
10.8.3	Informing regulators and authorities	217
10.9	Contingency planning for problems with the transfer	217
10.9.1	Risk analysis	217
10.9.2	Business impact analysis.....	217
10.9.3	Design of monitoring activities.....	218
10.9.4	Design of contingency activities	218
10.9.5	Decision taking	218
10.10	Awareness of problems.....	218
10.10.1	Definition of incidents	219
10.10.2	Being informed of incidents or detecting them	219
10.10.3	Incidents affecting the confidentiality of data	219
10.10.4	Incidents affecting the integrity of data	219
10.10.5	Incidents affecting the accuracy of data	219
10.10.6	Respect of the reporting lines	219
10.11	Problem tracking over time.....	220
10.11.1	Recording incidents	220
10.11.2	Follow-up and closure	220
10.12	Monitoring the quality of service	220

Contents

10.12.1	Regular reporting	220
10.12.2	Local dashboard	221
10.12.3	Access to statistics.....	221
10.12.4	Management meetings	221
10.13	Monitoring of technical and compliance developments.....	221
10.13.1	Technical developments.....	221
10.13.2	Compliance developments	222
10.14	Obtaining periodic and ongoing comfort about the service provider’s security management	223
10.14.1	Comfort obtained from third parties	223
10.14.2	Comfort obtained through internal procedures	223
10.15	Framework for Management Comfort	223
10.16	The Cloud Security Alliance Matrix	229
10.17	Conclusion	230
Chapter 11	Merits and Limitations of the Proposed Solution	231
11.1	Introduction	231
11.1.1	General versus specific	231
11.1.2	Audit and internal control perspective.....	232
11.1.3	Supportive rather than prescriptive	233
11.1.4	Project management methodology.....	234
11.1.5	Evolution of assumptions	234
11.1.6	Direct linkage to security objectives.....	234
11.2	The feasibility of the model.....	235
11.2.1	Data identification and preparation	235
11.2.2	Risk evaluation and data identification	237
11.3	Future activities and developments.....	238
11.3.1	Walkthrough or implementation.....	238
11.3.2	Updates of the theoretical and structural basis	239
11.4	Overall conclusion	241
References	244
Glossary	250
List of Appendices	253

List of Figures

Figure 3.1 Three approaches to data security, from Kelley (2008)	45
Figure 3.2 A three-step approach to data-centric security, from Kelley (2008)	45
Figure 3.3 Data management approaches, cited by Ely (2010)	53
Figure 3.4 Use of formal encryption policies, cited by Ely (2010)	54
Figure 4.1 Cloud Service Models, showing levels of abstraction, from ISACA p. 14.....	70
Figure 5.1 Responses received by organisation (population of 43).....	105
Figure 5.2 Ownership of responding entities (population of 34)	106
Figure 5.3 Size of responding entities in terms of number of employees (population of 34)	106
Figure 5.4 Selected responses to Section A (population of 41).....	109
Figure 5.5 Positive responses to questions 6 to 8 (population of 41)	110
Figure 6.1 Access management sub-processes	133
Figure 7.1 The three kinds of residual risk, from ETSI	144
Figure 7.2 Definition of cyber-defence terms, from ETSI	145
Figure 7.3 The Internal Control - Integrated Framework, from COSO (1992)	169
Figure 7.4 COSO Enterprise Risk Management - Integrated Framework (2004)	172

List of Tables

Table 2.1 An example of Control Objectives, Risks and Control Activities	22
Table 3.1 Definition of key terms	42
Table 3.2 Breakdown of unstructured data by file type.....	47
Table 3.3 Leading inhibitors in respect of security.....	48
Table 3.4 Leading solutions to security issues.....	48
Table 3.5 Measures to address security concerns.....	50
Table 4.1 SWOT analysis in respect of the cloud.....	68
Table 4.2 Cloud storage risk factors	73
Table 4.3 Risk factors specific to cloud deployment models.....	74
Table 4.4 Value propositions arising from the Open Cloud Manifesto	76
Table 4.5 Key objectives in respect of information security.....	85
Table 4.6 Cloud security design principles	88
Table 4.7 Relevant NIST security principles.....	89
Table 5.1 Survey questions.....	107
Table 5.2 Summary survey results.....	108
Table 5.3 Major findings of the two surveys	123
Table 6.1 Factors inhibiting the analysis of cybercrime	131
Table 7.1 Top-level choices in respect of incident areas	149
Table 7.2 Top-level choices in respect of vulnerabilities	150
Table 7.3 The ISO 27000 family of standards	161
Table 9.1 Security requirements through the data lifecycle	199
Table 10.1 Template for management comfort	228

Chapter 1 Introduction

1.1 Introduction

圖難於其易

“Anticipate the difficult by managing the easy”¹

It is a fundamental principle of governance and management that organisations need to have effective internal controls in place to ensure the sustainability, stability, adaptability and, if necessary, the profitability of their activities. A great deal of effort has been exerted in this direction, particularly over recent years: academics and auditors have developed and published frameworks and templates in respect of internal controls; regulators and governments have insisted on the application of these structures in order to ensure, as far as can be possible, some level of consistency in performance, behaviour and reporting; and the managers and staff of countless organisations, commercial or not-for-profit, publicly or privately owned, international or very local, have attempted to comply with the requirements and the methodologies and to implement reliable and measurable internal control systems.

It should be admitted very early on in this work that attitudes towards compliance and towards the existence and operation of effective internal controls are not always as positive and committed as stakeholders might wish to believe. Not every organisation feels enthusiastic about committing the resources to ensure that good and demonstrable corporate governance is in place, while even those who are obliged under national or international or industry-based frameworks to demonstrate a certain level of compliance may only ever do the strict minimum to obtain the necessary certifications or sign-offs. In seeking to develop methods for improved quality and efficiency in governance and internal processes, however, it is necessary to assume that best practice is valued and that good faith is present on the part of corporate management. Without this, no amount of analysis and recommendation will drive improvements.

1.2 Description of the problem

The problem can be broken down into two aspects: control and compliance; and risks and opportunities.

1.2.1 Control and compliance

The control and compliance of information systems have become increasingly time-consuming and complex tasks over recent years. There are a number of reasons for this:

¹ Lao Tzu, “Tao Te Ching”, Chapter 63

- Systems have become increasingly complex and increasingly fundamental to the operations of organisations;
- Information systems have moved out of the basements of organisations and the strict and arcane control of information technology departments and become desktop tools, operated – and in theory managed – by users and user departments and used as the basis for a range of decisions and key processes;
- The interconnectedness of systems has increased in a number of fundamental ways. Different systems within organisations are linked by automated and semi-automated interfaces, while there are ever increasing links with external partners, suppliers and customers, and new technologies are being employed to facilitate such connections;
- The sharing of data and resources between organisations has become fashionable and desirable. The idea of big data and the ability to extract value from large and disparate datasets has driven developments in analytical methods;
- Technological advances have led to rapid increases in both the quantities of data that can be stored and the amount and nature of processing of these data that can be envisaged;
- Demands for flexibility in data storage aligned with the inevitable desire of financial managers to reduce costs to a minimum have led to the offer and uptake of cloud services, promising efficiency and quality of service without the management effort or capital costs involved in maintaining additional infrastructure resources in-house;
- Compliance requirements have increased greatly as a result of global trends towards enforcing effective corporate governance on organisations and insisting that this corporate governance applies to the management and operation of the information systems that underpin business and reporting activities.

As a result, of these changes and developments, driven largely by technological advances in respect of storage capacities, bandwidth, mobile communications and means of interoperability, corporate management is confronted by an operating environment in which old certainties no longer apply and in which new challenges in respect of internal control have arisen.

1.2.2 Risks and opportunities

Corporate governance is a question of the management and balance of risks and opportunities. Recent developments and trends have highlighted three areas of opportunity, among others, that have the potential to bring about fundamental changes to the way organisations operate, both internally and in relation with others. These opportunities, discussed in detail in Chapters 2 to 4 of this thesis, are the use of unstructured data, big data, and the cloud.

From an audit and control perspective, these opportunities are automatically accompanied by risks. Each opportunity carries with it its own specific risks, risks of exposure, loss, and lack of compliance. There is also the combined risk – the confluence of risks from which this thesis takes its title – that arises when two or more of these opportunities are seized simultaneously. Risk management is difficult at a corporate level even when one-dimensional and the risks related to a specific activity are recognised, quantified and

prioritised. Dealing with multi-dimensional risk, especially when the context is new to an organisation, makes the process of risk management even more difficult.

1.2.3 Requirement for solutions

There is therefore a need for the consolidation of best practices and technical understanding into robust and applicable frameworks that can be used as the cornerstone of projects related to the introduction and implementation of new technologies and operating models.

The model proposed in this thesis is very much presented from an internal control and external audit perspective, reflecting my professional interests and experience. It is also intended to be platform, industry and sector independent.

1.3 The structure of this thesis

This thesis can be broken down into a number of sections.

In Chapter 2 I discuss the significance of the problem, considering the earliest days of information management and discussing the basic principles of governance and control.

In Chapters 3 and 4 I review the technical and professional literature related to the three opportunities mentioned above and summarise their importance. This forms an essential basis for the further work presented in this thesis: considered separately, a great deal of analysis and discussion has been documented in respect of each of subject and some consensus has emerged over best practices, genuine risks, and management objectives. No work aiming to consolidate existing work and combine current thinking in respect of risks, challenges and responses can be presented without a clear analysis of the relevant literature, the conclusions from which provide a framework for the rest of this thesis.

In Chapter 5 I present the results of two surveys performed in respect of the use of unstructured data and the cloud. The first survey is original research, seeking to determine the state of the market and the attitudes of senior and technical management to the questions raised by the potential use of these technologies. The second survey, performed in the United States by two professional organisations, provides additional information and perspectives on a similar range of questions. The analysis of the results of these surveys is a key step firstly in determining whether problems exist for organisations, and secondly in designing proposed solutions for these problems.

In Chapters 6 to 9 I discuss the constraints in place over organisations when seeking responses to the risks presented by the proposed adoption of one or more of the opportunities in question.

In Chapter 10 I present an analysis of the requirements for a global, high-level solution to the need for a control and compliance framework and propose such a framework for project management and oversight. The framework uses the vocabulary and principles of the auditor but is designed to be used by management in order to identify and address their own risks.

In Chapter 11 I discuss the limitations and weaknesses of the proposed framework, present the results of its partial implementation at two organisations in French-speaking Switzerland, and suggest directions for further development and improvement.

There then follow four appendices, containing figures that illustrate statements or references made in the body of the text, the detailed results of my own survey of organisations, my conference papers from 2009 to 2013 that are related to the questions of control, compliance and data management, and the text of a book chapter that draws heavily on the analytical work presented in Chapter 5 and the theoretical work presented in Chapters 2 and 8.

1.4 Notes on vocabulary and style

It has been necessary to take decisions in respect of four aspects of vocabulary and presentation when writing this thesis.

Firstly, I have chosen to employ standard British English spelling, as this is the form of the language with which I grew up. I have retained the original spelling in direct quotations from sources, however, and because of the preponderance of American English spelling in the published literature the reader should not be surprised to see alternate spellings of certain words depending on the context.

Secondly, I have chosen to refer to “organisations” throughout. In many cases the context will indicate that the discussion refers primarily to commercial businesses, but the underlying principles of the analysis apply, in general, to all kinds of organisations, be they businesses, not-for-profit organisations, administrations or associations.

Thirdly, in a perhaps misguided display of pedantry and old-fashioned thinking, I have consistently treated “data” as a mass noun and thus, in the context of storage and management, to be plural. This runs counter to the predominant usage in the technical literature, where “data” is most commonly treated as a singular (count) noun. I believe, however, that there is value in retaining the difference between the singular “datum” and plural “data”, and note with approval that this distinction is still made in French, the working language of my department. As noted above I have retained the exact text of quoted sources, and so “data” will appear as singular when included in direct quotations.

Fourthly, and perhaps inconsistently, I prefer to consider nouns such as “management” to be implying more than one individual and so have chosen to conjugate verbs in the plural when referring to management activities.

Chapter 2 The Significance of the Problem

2.1 Introduction

“Errant consilia nostra, quia non habent quo derigantur; ignoranti quem portum petat nullus suus ventus est”

“Our plans miscarry because they have no aim. When a man does not know what harbour he is making for, no wind is the right wind”²

Organisations today are being confronted by pressing and complicated questions related to their ever-increasing generation of, and reliance on, large quantities of data. As different pressures combine to oblige organisations to address their management and use of data, there is a need for understanding, structure and clarity in adapting management activities and internal control procedures to correspond to current risks.

2.2 A changing environment

The accumulation of information by organisations is nothing new. Throughout recorded history, administrations, businesses and organisations have discovered that creating, recording, retaining and reviewing information is a natural consequence of any kind of activity, and the challenge of managing this information has been a consistent driver for procedural and technological developments.

Archaeology and assiduous historical and linguistic research have provided numerous examples of how cultures developed techniques for recording information.

The first attempts to make a record of historical events in China of which historians are aware date from the Shang dynasty of the second millennium BCE (there is some uncertainty over precise chronology but there is general agreement that the dynasty was in power from c. 1600 BCE to c. 1046 BCE). The earliest surviving records were inscribed on bones or tortoise shells, which permitted their survival, while in later centuries, chroniclers left detailed accounts on paper or silk. A particularly important source of physical material for historians has been the Ruins of Yin, close to modern day Anyang, which has been identified as the last Shang capital. There tens of thousands of artefacts have been discovered providing insight into the history of the period: valuable information on the politics, economics and religious practices of the time, as well as insights into art and medicine.

The Royal Library of Ashurbanipal, named after the last great king of the Neo-Assyrian Empire, was established in Nineveh some time during the seventh century BCE. Its rediscovery and excavation during the middle decades of the nineteenth century uncovered thousands of clay tablets and fragments, the

² Seneca the Younger, “Epistulae Morales ad Lucilium (Moral Letters to Lucilius)”, Letter LXXI: On the supreme good, line 3

majority of which are now in the British Museum in London. The collection database counts 30,943 items, which given the number of fragments probably relate to around 10,000 separate texts [1]. The tablets cover a wide range of subject matters, ranging on the administrative side from legislation, foreign correspondence and engagements to aristocratic declarations and financial matters. Religious texts contain divinations, omens, incantations and hymns to various gods, while yet others are concerned with medicine, astronomy, and literature, this last category including such famous texts as the Epic of Gilgamesh. Ashurbanipal was renowned as both a ferocious and successful war leader and an educated and intellectually curious man, and he accumulated texts from throughout his empire and as booty and tribute from neighbouring states. His thirst for learning and knowledge was so wide-ranging that the information had to be preserved in written form to be able to be preserved and assimilated.

Tradition has it that Alexander the Great was so inspired by the scale and reputation of this library that he determined to build his own; given the three centuries that passed between the destruction of Nineveh and Alexander's reign, however, such old Persian and Armenian traditions are at best doubtful. What is certain is that a great library was constructed either in the reign of Alexander's friend and successor Ptolemy I Soter or his son Ptolemy II, and was fittingly located in the city of Alexandria in Egypt. Part of a larger research institution called the Musaeum of Alexandria, where many of the most famous thinkers of the ancient world studied, the library was conceived as a symbol of the wealth and power of Egypt. It employed many scribes to borrow books from around the known world, copy them, and return them – or, in the case of books found on board ships visiting the port of Alexandria, copy them and give the copy back to the ship, keeping the original.

Most of the books were kept as papyrus scrolls. It is unknown how many scrolls were housed at any given time, although estimates of hundreds of thousands of scrolls are commonly accepted. It is important to note that as a single work could span several scrolls, and that the Library sought to obtain multiple copies of major works for the purposes of comparison and textual criticism, the number of individual works held in the collection was likely to have been in the tens of thousands.

An important development in information management came with the compilation by Callimachus (310/305–240 BC) of his *Pinakes* ("tables"), a critical inventory of literary works widely considered to be the first library catalogue (although possibly inspired by the methods and practices of earlier Mesopotamian libraries) and based on the holdings of the Library of Alexandria during his time working there. The papyrus rolls in the Library were grouped together by subject matter and stored in bins, each of which carried a label with painted tablets hung above the stored papyri. The *Pinakes*, named after these tablets, are a set of books or rolls of index lists. The tablets on the bins gave bibliographical information for every roll, typically starting with a title, the author's name, birthplace, his father's name, any teachers he trained under, and his educational background, a brief biography of the author and a list of the author's publications. The entry also included the first line of the work, a summary of its contents, and information about the origin of the roll.

Callimachus' system divided works into six genres and five sections of prose. These were rhetoric, law, epic, tragedy, comedy, lyric poetry, history, medicine, mathematics, natural science and miscellanies. Each category was alphabetised by author.

This is important because it shows how early in the intellectual tradition that links us to the Assyrians and the Greeks the necessity of cataloguing and indexing became apparent, and how at the same time the

utility of metadata – data about data – became clear. Even in the earliest days of document management, when documents took the form of clay tablets or papyrus scrolls, those responsible for maintaining libraries and records started to need to consider how to manage the information they had, deciding how to store it physically, how to be able to retrieve it, and how to record what is stored in each location.

Even as the technologies used for information storage have developed and evolved, driven most notably by the development of movable type printing and the electronic revolution, the principles and challenges have remained the same.

2.3 Increasing quantities of data

As the quantities of data and information have continued to grow, so have the needs for the efficient and practical means of managing these data. The judicious use of metadata remains a key element of such efforts, reducing the quantity of data that needs to be recorded in cataloguing systems to uniquely identify both sources and items of data, and seeking to impose some kind of structure on potentially wildly diverse data elements.

The digital revolution of recent decades has accelerated the rate of growth of data. Hilbert and Lopez [2] report that the global capacity to store data, which is driven by the need for storage, has roughly doubled every month since the 1980s, with general purpose computing capacity increasing by an annual average of 58%. They give a graphic illustration of the rapid increase in the quantities of data in circulation: “The total amount of information grew from 2.6 optimally compressed exabytes in 1986 to 15.8 in 1993, over 54.5 in 2000, and to 295 optimally compressed exabytes in 2007. This is equivalent to less than one 730-MB CD-ROM per person in 1986 (539 MB per person), roughly 4 CD-ROM per person of 1993, 12 CD-ROM per person in the year 2000, and almost 61 CD-ROM per person in 2007. Piling up the imagined 404 billion CD-ROM from 2007 would create a stack from the earth to the moon and a quarter of this distance beyond (with 1.2 mm thickness per CD).”

Individual organisations are confronted with corresponding increases in the sizes of their data stores. At the beginning of the 1990s, a server with a disk capacity of 300 megabytes was not uncommon to support the central systems and file server needs of a medium-sized company, and quarter-inch analogue tapes with a capacity of 250 megabytes were sufficient for the purposes of a full system backup. Similar businesses in 2014 are using several terabytes of storage for their daily operations.

Such increases have already inspired many organisations to consider, in a more or less structured way, their strategy in respect of overall data management. In many circumstances it is impractical to retain all data in live production systems, both because of the disk space being filled and because of the impacts on processing time. An appropriate strategy therefore needs to be defined, perhaps adopting a policy of regular archiving of data, or even

2.4 Old certainties no longer apply

There is an immediate and obvious impact of these changes. For centuries, the growth in the quantities of information being created or received by organisations was predictable and manageable, and techniques for managing, accessing and archiving the information could be developed and allowed to evolve gradually while continuing to address operational risks. Management could justifiably be confident in their continued ability to manage their information-related risks. The explosion in the quantities of data

created, starting in the 1980s, accompanied by the revolutions in digital technologies for storing and processing these data, has led to drastic changes in both the risk profiles for organisations and the technical and organisational tools available to address these risks.

The increased quantities of data can be divided into three categories.

The first, ordinary structured data generated during standard business activities as a result of daily operations, tends to grow at a predictable and manageable rate and does not present any novel or disproportionate operational risks.

The second, unstructured data, arises from the extraction of data from central systems and the operation of non-centralised applications and systems. Unstructured data are defined fully and discussed in detail in Chapter 3.

The third can be labelled with the umbrella term Big Data. This term, also discussed in detail in Chapter 3, covers the accumulation and processing of large datasets, often from disparate sources, in order to generate useful information for competitive advantage or more efficient operations.

The rapid growth in the quantities of data being generated and processed can in many cases oblige organisations to reconsider their strategies in respect of data storage. Many will increase their efforts to control data by improving or reinforcing data management processes, including backups and archiving. Some will undertake a process of identifying data that can be deleted. Such decisions will of course need to take account of regulatory and compliance requirements, such as the need to retain certain accounting information for ten years in various jurisdictions. Still others will consider the likely quantities of data that they will be handling, review their capacity to manage in effectively, and consider the use of an outsourcing service provider, often in such circumstances a cloud service provider, as a possible solution. The impacts of such decisions are also described and discussed in Chapter 4.

2.5 A cartography of risks

In addressing the new and modified risks presented by the twin challenges of unstructured data and big data, the management of organisations need to adopt a systematic and well-structured approach.

A model that is widely used in industry for ensuring that risks are properly and completely identified is a simple three-column one. This model can apply across an entire organisation or to individual business units or business cycles; frequently in practice they exist as a set of nested documents to which senior management can refer as a whole in order to ensure completeness while departmental management will be responsible for the documents relating to the cycles or processes with which they are occupied. In respect of IT systems, for example, the highest-level objectives in respect of systems and data may well be included in high-level matrices for the whole organisation related to key business operations and continuity, and then there may be a subset of documents relating to the key domains of IT management such as Changes to Applications, Security, and Operations. These domains may also require, for the sake of clarity, separate documents relating to subdomains such as user security, data security, physical security, and others.

The first column details the Control Objectives for the organisation. These set out the fixed objectives of management in respect of the unit or cycle in question. In the context of data management, the objectives

are likely to address the issues of security, availability and consistency of data, to name three topics as examples.

The second column details the Risks, the potential obstacles to achieving the Control Objectives. Crucially, every risk must relate to a specific control objective, while each control objective can have more than one related risk. This is essential both in ensuring that the risks identified are relevant and relate to genuine objectives, and that all the potential obstacles to achieving each objective have been identified.

The third column details the detailed Control Activities to be undertaken to address the risks that have been identified. Clearly such activities can only be defined after a proper evaluation of the risks, because not every risk will necessarily be best addressed by some kind of mitigating activity.

Conventional risk management theory in the business process context and as applied in such frameworks as COBIT [3] allows for four different options when confronted by a risk. These options are as follows:

- i. Mitigate: design and perform control activities that allow the risk to be brought down to an acceptable level for the organisation. This is the approach that gives rise to traditional internal control activities. An example would be addressing the risk of inappropriate and unauthorised access to systems and data by implementing unique user IDs with a supporting password regime.
- ii. Transfer: in some circumstances risk can be transferred onto a third party, usually in return for some consideration. A typical example of this would be insurance, where in return for premiums the risk of significant financial loss, for example, can be transferred to an insurer.
- iii. Avoid: in some circumstances the evaluation of a risk leads to the conclusion that the most appropriate response is to avoid activities that lead to exposure to that risk. An example of this would be a company declining to do business in certain markets or countries because of the risk of terrorism, kidnapping or civil war.
- iv. Accept: in other circumstances analysis of the risk may lead to the conclusion that although the risk is valid, the potential consequences of that risk being realised (likelihood multiplied by impact) are sufficiently low that the risk can simply be accepted without any action being taken. Examples of this are traditionally difficult to find because by their nature acceptable risks are often unlikely and viewed as improbable; one example might be the decision to locate company premises close to an airport with the potential risk that went with it of being affected by a serious plane crash on take-off or landing.

As with the relationship between Control Objectives and Risks, the relationship between Risks and Control Activities is one to many. In practice it is rare that one single control activity can address all aspects of a risk, especially when considering the design of Control Activities from the perspective of business process control objectives. Under the model commonly employed throughout audit methodologies, it is important to ensure, throughout the processing of business transaction, the completeness, accuracy and validity of all transactions, as well as restricted access to systems and data. It is uncommon to find any internal control that can cover all of these objectives; even when the first three are covered, it is usually by means of a resource-intensive activity such as a one-to-one verification of a transaction before posting or approval, which may not be efficient in a context of high volumes of transactions.

We thus have a model in which each Control Objective will be linked to one or more Risks, and each Risk is linked to one or more Control Objectives. There is therefore a direct link in both directions between Control Objectives and Control Activities. This is an essential element in ensuring the quality of the risk management and control environments. When such linkage is clearly documented and reviewable, the reader can be certain that every Control Objective has at least one relevant Control Activity relating to it. The reader can also see why each Control Activity exists, because it will be directly linked to a Risk and thus to a Control Objective. Hence redundant or unnecessary internal controls can be identified and at a later stage decommissioned or removed.

A good and appropriate example of this methodology can be given in respect of data security. Clearly there are many aspects to data security and these will be discussed in detail in Chapter 6, but a subset of the possible Control Objectives can serve as an illustrative – and far from exhaustive - example of the theory described above.

Control Objective	Risk	Control Activities
Data Security		
Only authorised users should have access to systems and data	Unauthorised users may copy, modify or delete sensitive data	All users have unique user IDs
		Passwords of minimum length and complexity apply to every user ID
		All requests for user access follow a formal and documented approval process
		User access rights are regularly reviewed for appropriateness
	The activities of unauthorised users may be difficult to trace and attribute	All system activities are logged in audit files
		System log files are reviewed on a regular basis to detect anomalous activities

Table 2.1 An example of Control Objectives, Risks and Control Activities

It can be seen that this one example Control Objective ultimately has six Control Activities linked to it, and that the reason for the existence of each can be clearly related to a defined and genuine Risk.

2.6 Specific risks related to unstructured data, big data, and the cloud

A number of risks specific to each change in the operating environment can be identified. These will be discussed in detail in Chapters 3 to 10, but an indication of their scale and scope can usefully presented here.

In general risks can be classified in terms of how they prevent an organisation from achieving its objectives. Some can be considered problems of action, internal or external, accidental or deliberate, that cause processes to fail or targets not to be met. Others can be considered problems of inaction or ignorance, where opportunities are not identified and actions are not undertaken that would help in meeting objectives.

In respect of unstructured data, the most significant risks based on my experience and current thinking in the audit profession include the following:

- Taking decisions based on incomplete, inconsistent or irrelevant data
- Losing control of data because the data do not reside on a platform or within an application to which corporate security standards apply
- Failing to leverage information and knowledge that reside within an organisation because of a lack of knowledge of what is there and a lack of ability to process and use this information.

In respect of big data, the most significant risks based on my experience and current thinking in the audit profession include the following:

- Taking decisions based on inconsistent or inappropriate data
- Not appreciating the requirements for managing the security of the data accumulated and used

In respect of cloud services, the most significant risks based on my experience and current thinking in the audit profession include the following:

- Being unable to ensure the confidentiality, integrity and availability of data not held within the traditional control perimeter
- Being unaware of issues in processing
- Being locked into the use of one service provider

The above points constitute only a small part of the risk universe in which organisations operate. Clearly every organisation will have its own unique risk profile, and I will return to this principle throughout this thesis when insisting on the necessity of each management team performing its own detailed risk assessment rather than trying to rely on a pre-populated list of risks obtained from a methodology or professional adviser. The key points identified above are sufficiently fundamental to apply to the majority of organisations, however.

2.7 Consequences of the failure to address risks.

As will be discussed in detail in Chapters 6 to 10, the failure to address the risks related to the management and use of data can have a number of consequences of varying severity. These include:

- Problems with regulators resulting from non-compliance with standards and legal requirements
- Problems with clients and partners resulting from security breaches
- Problems arising from the use of incomplete, inconsistent or incorrect data for business operations or decision making
- Problems in ensuring the integrity of data

- Problems in maintaining access to necessary systems and data

It is often a useful motivating factor in driving change to encourage corporate management to consider and attempt to quantify the impacts on their activities and operations of such problems. Clearly every organisation will have its own priorities and different requirements: certain industries such as banking and insurance operate within very specific regulatory environments and the effects of non-compliance or security breaches will be, *prima facie*, more significant than a problem of data leakage for a small manufacturing company.

The more fundamental and all-encompassing issues can have similar effects on all organisations, however. I can cite, from my own professional experience, two cases of weak controls over data quality leading to significant financial losses and cash-flow difficulties for the organisations in question.

The first case involved a privately-owned property company with a medium and long term business model that was essentially based on deciding whether to build property or to buy and renovate existing property. The key tool used for the build or buy decision was a Microsoft Excel spreadsheet, containing multiple worksheets and drawing in data from files created by extracting data from the central systems, as well relying upon manual input to update assumptions and statistics. When this model was professionally audited, at the request of a board member who had doubts about the accuracy of the results that the model was producing, it was found that one row of figures in one worksheet was consistently inaccurate to a factor of one hundred as a result of an error in the data extraction process. Reperforming the analysis for the five previous years using corrected data showed such significant changes in the overall results that the decision taken in each of those years would have been reversed.

The second case involved another privately-owned company that included a specialist treasury unit that was dedicated to placing the company's cash reserves on the markets, reviewing their positions and making changes every day. Similarly to the previous case, one of their key tools was a Microsoft Excel spreadsheet of great complexity that was updated through a combination of data feeds and manual input. Once again, an external audit commissioned because of a sense of unease with some of the results revealed errors in some of the source data being used, errors that led to final results that were skewed by amounts rising into the millions of pounds. Fortunately for the company in question, any losses incurred because of these systematic errors were potential, in the sense of missed opportunities to invest more effectively rather than actual losses caused by inappropriate investments.

Both organisations subsequently understood the potential impact of data quality and accuracy problems in their business models and designed policies and procedures to minimise the risk of such problems recurring.

The discussion of these cases should not pass without an educated comment on another common factor, which was the use in both cases as a key management tool of a Microsoft Excel spreadsheet. Such models, frequently unvalidated, unaudited, undocumented and without clear ownership, present opportunities for error whenever used.

The other major class of problems mentioned above that can easily have serious consequences for an organisation is the lack of availability of systems and data. This is not uniquely a problem related to unstructured data, big data or the cloud, but it is of particular relevance to any organisation that relies upon its data for ongoing operations; logically organisations that are invested in the processing of large

quantities of data and/or are using cloud services for storing it fall into this category. Should data be unavailable in a reliable form for any length of time, the very future of the organisation could be endangered. This could be the result, for example, of a loss of clientele resulting from an inability to respond to or process orders, a breakdown of relationships resulting from an inability to pay amounts due, or legal consequences arising from an inability to fulfil regulatory requirements.

Management should thus be aware of the range and importance of the potential consequences of failure to address risk issues effectively, and thus develop and implement appropriate risk management procedures.

2.8 Conclusion

The adoption of new technologies is a balancing act between opportunity and risk. In order to take the most appropriate strategic and tactical decisions, managers and experts need to understand the nature of the technologies and initiatives that are being presented to them, to appreciate the strengths, weaknesses, opportunities and threats that they bring, to understand how to evaluate the risks and the benefits arising from the opportunities, and to design, implement and operate effective and targeted frameworks of internal controls capable of providing the necessary assurance that risks are being adequately addressed while activities continue.

Chapter 3 Analysis of Existing Research I: Unstructured Data

“Data have become a torrent flowing into every area of the global economy”³

3.1 Introduction

The topics of unstructured data and Big Data have been the subject of numerous technical and popular articles in recent years, with the two subjects perhaps understandably being treated together as the full extent of the potential for amalgamating disparate datasets and extracting information and value from the processing of such datasets has started to become apparent.

In this chapter relevant literature relating to both unstructured and big data is discussed and the common themes and conclusions identified.

The choice of literature to consult was driven by a philosophy of obtaining as wide a range of perspectives as possible. My approach was therefore to consider not only traditional academic publications, although these were widely consulted, with particular care taken to consult older publications in order to obtain and reflect upon the way that the concepts of unstructured data and big data have evolved over time. In addition I have widely consulted both articles in the more popular technical press and white papers published by putative service and product providers. Clearly such publications need to be read with a sceptical eye, because the motivation behind their creation is unmistakably commercial. On the other hand, though, they provide useful sources of information in respect of what is actually happening for businesses and organisations; where their concerns might lay, which directions they might be considering, which risks are taking on significance for them, and which kinds of solutions might be attractive and appropriate.

This chapter is divided into several sub-sections. Initially I discuss the definitions of both unstructured data and big data and consider how their meanings and significance have developed over time. Then I consider the essential questions that the use of such data raises for organisations: what are the privacy implications of data storage and manipulation; how can data be managed effectively outside a traditional structured database schema; how can big data be managed and controlled; and what are the consequences of the existence of so-called “toxic data.”

The objectives of this analysis are to demonstrate the scale of the impact of the existence and use of such data for modern organisations. Effective risk management, a theme that will recur in this thesis, is necessarily grounded on a wide-ranging and realistic understanding of the circumstances in which an organisation is operating. In order to identify any risks related to the storage and use of data, management need to know what data they are holding and what their impact can be, both on internal operations and on the subjects of those data.

³ The Economist, “A special report on managing information: Data, data everywhere”, 25 February 2010

3.2 Unstructured data: definitions and developments

Buneman et al [4] discuss the requirements for and advantages of holding data in an unstructured format, citing the examples of biological research and recording data relating to films and television series. “There are two good reasons to query and manipulate data whose structure is not constrained by a schema. First, some systems and proposals have recently emerged in which the schema is absent or, when it exists, only loosely constrains the data; second, for the purposes of browsing it may be convenient to forget the schema, even though one exists.”

A real-life and financially significant example of how the storage of data in unstructured and often incompatible formats can cause problems in one of the contexts Buneman mentions, film and television series data, can be drawn from professional audit work undertaken by the author of this thesis at AGICOA⁴, the Geneva-based not-for-profit organisation that was “established thirty years ago to track and distribute royalties on retransmission of the products of independent producers” [5]. The principles of what this organisation tries to do are simple, if daunting in scale. The reality is of complexity and difficulty: “Our main competitor is frustration. The all-too frequent reaction to administering secondary property rights across kaleidoscopic variations of language registrations, rules, formats, ownership and time periods is to give up. Revenues are uncollected and the returns on creative investment are much less than they should be” [6]. At the core of the operational difficulties faced by AGICOA are differences in the type, nature, format and completeness of the data provided by rights holders, making matching, tracing and reporting an involved process that is difficult to automate with any degree of robustness.

Seiner [7] describes his definition of unstructured data, which he refers to as “artifacts.” “Artifacts includes data/documents/content recorded in electronic format that can be managed and leveraged for the benefit of your company, your customers, your suppliers, etc. Artifacts include word processing files, html files (web pages), project plans, presentation files, spreadsheets, graphics, audio files, video files, emails... any data that is not in tabular or delimited format.”

Suciu [8] discusses how the nature of data has evolved to include much data stored outside database management systems and discusses the potential of XML to handle machine-readable dynamic data sources. He makes the distinction between data formats such as HTML which are designed to be human readable and those such as XML which are primarily about data classification and argues for the provision of the right tools for analysis and for methods to optimise the storage of data.

Jhingran et al [9] describe data integration as the driving force of IT spending during the last decade, especially in the context of distributed data, an effort designed to overcome the silo-based nature of data storage. They discuss the major trends that make such integration difficult – the heterogeneity of data, the “federation” and “distribution” of data, and the increasing use of data for competitive advantage, and go on to point out the two different perspectives: “Thus it is clear that there are two slightly different perspectives - well-formed structured schema and the relatively poorly structured world of documents. Bringing these two worldviews together is the “holy grail” of information integration” (p.557).

Chu et al [10] comment on the “growing consensus that it is desirable to query over the structure implicit in unstructured documents.” Currently “users have to rely on a combination of keyword search, browsing,

⁴ www.agicoa.org

and possibly predefined search options. Although these mechanisms are easy to use and often lead to what users are looking for eventually, they cannot leverage the potentially rich set of structures embedded in text.” They identify interesting outstanding problems such as: “how to handle updates to the unstructured data; how to record the evolution of data...; how to help users write queries that exploit the structure discovered...; how to optimise queries....”

Considering the conceptual and practical differences between unstructured and semi-structured data, Buneman [11] gives as a characteristic of semi-structured data that “the information that is normally associated with a schema is contained within the data, which is sometimes called “self-describing”.” He points out that “even when dealing with structured data, it may be helpful to view it as semi-structured for the purposes of browsing” (p.117).

Franklin et al [12] discuss the challenges of data management: “providing search and query capability; enforcing rules, integrity constraints, naming conventions, etc; tracking lineage; providing availability, recovery and access control; and managing evolution of data and metadata” (p 27). They emphasise the importance of modelling dataspace as a set of participants (the data sources) and relationships, especially when the data do not share common formats and characteristics.

Raghuveer et al [13] state that as “supporting exhaustive search is not among the primary design criteria of today’s file systems and database systems... the problem of where and how to store unstructured data in order to search it efficiently needs a fresh re-assessment” (p. 951). Their search-oriented analysis discusses the need to integrate structured and unstructured data, and how this challenge is made more difficult by their different storage access techniques: structured data are best held within the page structures of databases, while file systems are better for unstructured data. Bringing disparate datasets together for useful comparison and extraction of knowledge is therefore complicated on technical and conceptual grounds.

Doan et al [14] report on a lengthy project to manage unstructured data using extraction, integration and user interaction. In particular they discuss building unstructured data management systems that “extract structures (e.g. person names, locations) from the raw text data, integrate the structures... to build a structured database, and then leverage the database to provide a host of user services” (p.1).

They report that building such a system did raise many challenges in the areas of information extraction, information integration, and user interaction, many of which had not been apparent at the planning and design stages of the project. In their approach it was essential to distinguish and manage separately a number of key elements of the unstructured data management system, such as the physical layer, the data storage layer, the processing layer, and the user layer. This illustrates the complexity in such projects: whereas technologies and techniques for managing structured data have evolved to the extent that many issues relating to these areas simply do not arise, or at least are dealt with in a straightforward and almost inadvertent way during the planning and design phases of data management projects, in the field of unstructured data a great deal of structured and careful planning does need to be performed.

Taking a wider view, Maluf and Tran [15] discuss the key assumptions in approaching the management and integration of enterprise data. They state a number of key assumptions relevant when building a single system: “1 Data must always be stored and managed in DBMS systems... 2 The database must always provide for and manage the structure and semantics of the data through formal schemas... 3 Managing

multiple schemas from several independent data sources and their interrelationships between them is inevitable and unavoidable: thus produces “schema chaos” (p. 2).

The application of these fundamentals to the world of unstructured data provokes questions of realism and validity: is the idea of schema chaos necessarily inevitable, or can it be avoided or at least partially side-stepped by the application of metadata management techniques or other approaches to data perception and handling?

In a substantial review of information extraction research, Sarawagi [16] talks of the “interest in converting our personal desktops to structured databases” (p. 261), an ambition clearly shared by many technologists and professionals who find themselves submerged with data and with few convenient tools for managing the information that they have generated and acquired. Two key considerations are identified: Customer Care and Data Cleaning; thereby focusing attention on the objectives for the successful end use of the information being extracted. The article then classifies the types of sources of unstructured data along two axes: the basic unit of granularity (record or sentence, or paragraph and document); and the heterogeneity of the sources (machine generated ages, partially structured domain specific sources, open ended sources), in order to give a sense of meta-structure to the possibilities for obtaining and processing unstructured data.

In an article arguing for a structured approach to managing unstructured data, Doan et al [17] give as a basis for their work the prominence of “unstructured data, which we take to include text documents, Web pages, emails and so forth” (p. 1). They note that “dealing with this kind of unstructured data may require fundamental changes to the entire end-to-end systems we use to manage the data” (p. 1). Discussing the current and potential future trends in the field, in which developments are likely to be driven by commercial imperatives and the search for profits, they comment that “web companies large and small have found a new business model: they develop such applications (and often also the hosting platforms), then invite developers to use them to build compelling Web services” (p. 6). This leads them to the tentative conclusion that “the above scenario offers an interesting vision for the evolution of the Web: the Web will become increasingly structured, but in a bottom-up fashion” (p. 6).

This is an interesting perspective on the wider question of unstructured data because rather than concentrate narrowly on the detailed managerial and technical questions of how to handle modern unstructured data sets within organisations, it takes a step back and considers how the whole data and information ecosystem might evolve, with structure and order emerging from the chaos.

In an article on the business uses of unstructured text, Kuechler [18] quotes “A widely touted IT factoid states that 80% of the information produced by and contained in most organisations is stored in the form of unstructured data” (p. 86) thereby positioning the subject firmly in the context of modern corporate activities. He then remarks that “Prior to the legal necessity of monitoring regulatory compliance, the data was largely ignored by both corporate IS departments and corporate management for several reasons” (p. 87): these reasons included their dubious – or at least unrecognized - value, the fact that their expertise needed to tap their value was rarely to hand, and the way that the storage and processing requirements were prohibitively expensive.

This analysis does rather gloss over a fundamental weakness in corporation management practices over a couple of decades, starting when desktop applications became feasible and widespread and continuing, at least in contexts subject to strict regulation, until compliance and conformity and risk management became

the order of the day; this weakness related to the lack of appreciation of how much sensitive data was being stored within, and how many critical decisions were being taken based on, desktop applications that were not necessarily subject to development standards, security measures, or review and approval of the confidentiality, availability and integrity of data. The reason for this was, in this author's experience, a combination of a lack of awareness of what was going on in user departments and of a willingness to accept answers that had not cost too much to be generated, without asking too many questions about the underlying processes and assumptions.

In an article featured on popular technology websites, Jowitt [19] reports on the conflicting statistics in circulation in respect of the quantities of unstructured data present in companies across Europe. Research published by Hewlett Packard showed that on average European companies believed that 25% of their data was unstructured, with a specifically United Kingdom average of 29.4%. HP contrasted these numbers with research from industry analysts suggesting that more than 70% of information is in the form of unstructured data. This demonstrates a significant issue in developing solutions to the question of unstructured data: knowing what the dimensions and significance of the phenomenon are. The article then considers other factors that can act as brakes upon effective information management, citing research that the central bottleneck is information classification. This is particularly difficult to apply after the event, once data have already been created and stored, but methods for efficiently and accurately classifying data on the fly are also unreliable, creating potential issues for organisations that need to comply with legislation such as the Freedom of Information Act in the UK. Without data stored in a structured and classified format, demonstrable compliance would be impossible.

Taking a wider perspective, Abadi [20] reports that "Data management applications are potential candidates for deployment in the cloud. This is because an on-premises enterprise database system typically comes with a large, sometimes prohibitive up-front cost, both in hardware and software" (p. 3). He then goes on to cite the characteristics of the cloud that would lend themselves particularly to large-scale data management tasks; "Compute power is elastic, but only if workload is parallelizable" (p. 4); "Data is stored at an untrusted host" (p. 5); and "Data is replicated, often across large geographic distances" (p. 5). He then outlines the scope and nature of practical data management applications in the cloud, supporting transactional data management and analytical data management, such as querying a data store for planning, problem solving and decision support.

Beyer et al [21] report that "As the volume of semi-structured content within enterprises continues to grow, there is increasing interest and commercial value in harnessing this content to power the next generation of search and business intelligence applications" (p. 28), identifying the main driver behind applied research in the field. They give numerous examples of enterprise repositories of unstructured and semi-structured content, including email archives, call center transcripts, customer feedback databases, enterprise intranets and collaboration and document-management systems; identifying platforms, systems and applications in which potentially valuable information has a tendency to gather without being catalogued or classified and without its utility being recognised. They go on to stress the importance of analytics: "Transcripts of customer calls can yield valuable business insights in areas such as product perception, customer sentiment and customer support effectiveness. However, analytics is essential to extract the appropriate information from the raw text of the transcripts and transform this information into a form that can be consumed by BI analysis and reporting applications" (p. 28-29); this is the key to the whole topic of identifying and exploiting sources of data. They also correctly note that "increased

legislation around corporate data governance is requiring enterprises to invest in applications for regulatory compliance and legal discovery” (p. 29).

The author of this thesis is aware of two instances among corporations with which he had professional dealings of the legal discovery process uncovering documents and information of which management and in-house counsel were unaware and which proved to be of great significance in the legal proceedings that were underway. One case involved drafts of a document written in Microsoft Word, stored in backups of a departmental fileserver, that contained review comments and traces of modifications invisible in the final PDF file shared with third parties, while the other involved copies of email exchanges that likewise had been preserved on backups and that demonstrated that individuals within the corporation had in fact been aware of a potential exposure issue and had decided not to act. Both of these situations resulted in embarrassment and financial loss for the corporations in question.

In the context of incorporating datasets into larger collections of data and of performing semantic analysis of the contents in order to identify and classify them, Ceglowski et al [22] discuss the problems encountered in machine processing and the potential requirement for human input in cleaning and processing data and in interpreting and assigning significance to results. “The... approach has two weaknesses: first, it does not account for homonyms and polysems, especially words that exist both as verbs and as nouns, a very common feature of English... Second, it fails to identify semantically meaningful multi-word constructs, such as proper names and noun phrases” (p. 1). The principle of “Garbage In Garbage Out” has been familiar in the domain of information systems since the days of the computing pioneer Charles Babbage (1791-1871) and “the term is attributed to George Fuechsel, a programming instructor who used it as a teaching device in the late 1950s (p.112) [23]; this problem is of particular significance when considering the use of data of unknown sources and inconsistency structure and formatting.

Blumberg and Atre [24] discussed the basic problems inherent in the use of unstructured data a decade ago. They wrote “The term unstructured data can mean different things in different contexts” and that “a more accurate term for many of these data types might be semi-structured data because, with the exception of text documents, the formats of these documents generally conform to a standard that offers the option of meta data” (p. 42). Already there we can see the requirement for a careful and specific definition of terms. For them, a key first step in moving towards some kind of standardization will involve raising the awareness of the users and suppliers of technologies: “Once awareness of the issue is raised, the next step is to identify the unstructured data in the organisation” (p. 43). Subsequently they identify additional key issues, such as adding context to search, classification, and taxonomy: “One stumbling block has been the difficulty of creating and maintaining taxonomies. Essentially this task requires individuals who both understand the organisations and have a degree in library science” (p. 44). The management of unstructured and semi-structured data is thus immediately identified as a challenge that is not purely technical or to be left to technologists to solve. They comment that “additionally, a well-populated taxonomy may be eight to ten levels deep and contain hundreds, even thousands, of categories... updating and maintaining such a taxonomy is both time-consuming and expensive” (p. 44), thereby identifying the resource implications of a well-structured approach to data management, and then point out the need to address content intelligence, using discovery systems and platforms for content applications in order to achieve data integration and integration with enterprise applications.

Tsai et al [25] discuss the issue of data provenance. They point out that “It is critical for a Service-Oriented Architecture (SOA) system to consider its data provenance, which concerns security, reliability and integrity of data as they are being routed in the system” (p. 223), focusing on a fundamental principle of systems and data management. Unstructured data remain data that need to be subject to appropriate controls and validation, regardless of their sources and the potential for exploitation that might lie within. They remark that “Data provenance is not a new problem as it has been studied before in other fields, such as business and medicine” (p. 224). From a theoretical perspective, they echo Groth et al [26] in citing seven requirements for a provenance system: Verifiability, Accountability, Reproducibility, Preservation, Scalability, Generality, Customizability (p. 224-225).

Mukherjee and Mao [27] discuss the technologies useful for efficient searching, very much a basic requirement of the use of large data stores. “Data must be accumulated (spidered) and indexed before it can be searched. This requires knowledge of where the critical information exists and access to these repositories, which can be secure” (p. 39), they write, emphasizing that there will always be phases of preparatory work to be performed before data can be used in any systematically valuable way. The key ideas on which they focus are data filtering and search relevance, and they emphasise the informal and unstructured nature of many of the queries that get launched in the ordinary course of business: “Many intranet queries (60 to 80 per cent) are targeted to retrieve “stuff I’ve seen”” (p. 40). They recommend the systematic and reliable use of metadata: “Metadata in semi-structured documents brings tremendous value to content search and organisation” (p. 43); but do identify key problems with how this state of affairs can be brought about: “Subject matter experts can be hired to tag or annotate documents manually. Manual tagging, however, does not scale to large volumes of information; therefore, automation is mandatory” (p. 43).

This is a key limitation in many processes for tagging data and automating as much of the process as possible: very often the only people who know what the data refer to and in which context they have been acquired and previously manipulated are the users who are too busy to perform huge amounts of tagging and classifying.

An experience that is not brought out in this paper concerns the frequent – perhaps overwhelming – use of filenames that are vague and non-specific and that are very rarely supported by realistic and practical metadata. It is possible to envisage an automated process that can function quite efficiently crawling through a file server and classifying files with names such as “Mainco Trial Balance Q3 2012”; an identical process scanning directories or attachments named “Document” would require a large amount of manual intervention. This problem becomes even more acute when considering entirely unstructured data such as the text included in emails, where the contents could be valuable but much meta-information is likely to be contextual rather than explicit.

Duffy [28] discusses the elements of knowledge management: KM software, infrastructure tools and technologies, data warehousing infrastructure, content management software, businessware management systems, collaborative applications, knowledge management access software. The significance of this is that any attempt to harness and utilize unstructured data cannot be a single project performed in a vacuum and without appropriate resources. Those commissioning such projects within organisations need to place them within a framework of effective knowledge management, with clear objectives and the tools and resources that such ambitions require.

In 2000 Coy [29] indulged in a little prediction and remarked that “The Darwinian struggle of daily business will be won by the people – and the organisations – that adapt most successfully to the new world that is unfolding” and that “Attributes that made them (corporations) ideal for the 20th century could cripple them in the 21st century.” This was a reflection on the ever-changing business environment and the requirement to evolve and adapt, to seek greater value from the resources already committed and to extract greater utility from the information to hand.

In 1996, before the subject of unstructured data began receiving detailed attention in the technical press, Guynes and Vanecek [30] discussed the nature of critical success factors in overall data management, emphasising the “movement of the data administration function from its primarily technical database orientation to a more managerial one.” Their conclusions, applicable also to the context of trying to manage and impose some kind of structure upon unstructured data, include the importance of integrating new structures into this existing data and database landscape, educating users and analysts in the importance, contents and potential uses of the new data structures, and the necessity of developing, implementing and conforming to an overall corporate data policy and data standards. Any investment in using data sources needs to ensure that the overall landscape will be simplified and improved, not made more complicated.

Rao [31] discusses how unstructured data can yield useful and valuable knowledge and comments that “Internal documents might constitute companies’ most ineffectively utilized asset today” (p. 29). “Statistics indicate that workers waste many hours searching for, sorting and assessing information, incurring a significant organisational productivity cost” (p. 29). Definitions: “Neither content nor knowledge work is truly unstructured. Content, despite often being called “unstructured data”, is shaped – first, by intrinsic aspects of representation and expression and second, by the social context in which it is produced and consumed”, “The term unstructured data refers to the difficulty of applying IT to routing and accessing content – slicing and dicing and manipulating it in all the ways typical of data stored in rows and columns in relational databases” (p. 30). Specific applications for KM: publishing, intelligence and law enforcement, pharmaceutical research and development.

Baars and Kemper [32] discuss the role, merits and importance of metadata, pointing out that “After its extraction, the metadata-based content descriptions can be handled just like any other structured data source and be stored in an integrated data repository alongside other relevant data” (p. 134). They also discuss how these metadata can be created, an issue long problematic for organisations: “Regarding their origin, relevant metadata can either be entered manually in the source system by end users, for example, with fields to identify customer satisfaction levels. Interesting metadata can also result from access and usage logs or from search queries” (p. 134).

Metadata are viewed as a fundamental tool in data management by a number of authors. Seiner [7] proposes a personal conceptual unstructured data meta-model, suggesting the range of types of metadata that could or should be recorded about unstructured data and artifacts (his term). These types of metadata include: business function; subject area, purpose; steward; location; community and audience; security; data related; time or time/date; media type; package; status and version; project and process; and event. Clearly not every type of metadata will be applicable or available in respect of every item of unstructured data, but the scope of classification provides some level of assurance that sufficient useful distinctions can be made to create a valuable and practical store of metadata.

An interview with a senior manager at data governance company Varonis [33] reveals where specialist companies view the key security risks in respect of data. “Unstructured data has no commonly used access log, and permissions are managed manually... Unstructured data... is potentially the source of a data leak in many businesses for the simple reason that it tends to be stored but not fully audited. And if you can’t audit the data... you cannot determine what data is sensitive, who’s accessing this data and what they’re doing with the data.” The interview highlights the company’s estimation that “data volumes are growing at the rate of 650% every five years with 80% of these data volumes being unstructured” (2010 estimation). This reveals the rapidly increasing significance of the data management problem. The Varonis approach to solving these problems, as outlined in the interview, is to implement a classification framework with a metadata layer that allows management to determine which data have changed or are most frequently accessed.

In their own white paper [34] Varonis present their analysis of the problems related to the management of unstructured data and define ten requirements for “any effort, system or technology whose purpose is to protect business data.” These requirements are: visibility; control; auditing; security; performance; scale; ease of installation; ease of use; ease of integration; and low total cost of ownership. This list of requirements is interesting because it does not place security above all other criteria. It recognises the need for an over-riding governance concept and also recognises the importance of consistency and usability. Any system, application or methodology that is not entirely automated and invisible to users requires buy-in, understanding and acceptance on the part of the user community in order to ensure that it will be adopted, observed and respected.

In a high-level overview of information management, Detlor [35] asks some fundamental questions about how information management applies and relates to modern requirements for managing electronic data. He explains that for some authorities “information management draws upon ideas from both librarianship and information science.” He himself defines it as “the management of the processes and systems that create, acquire, organise, store, distribute, and use information. The goal of information management is to help people and organisations access, process and use information efficiently and effectively. Doing so helps organisations operate more competitively and strategically, and helps people better accomplish their tasks and become better informed.”

Detlor provides a definition of unstructured information as “the type of information that can be found in reports, documents, email messages, and PowerPoint presentations, among others. This includes reports and documentation generated internally with a company and outside the enterprise as well.” He then links this to the idea of document management systems, also known as content management systems, which “help manage unstructured information that is created, acquired, organised, stored, distributed and used within an organisation. These systems support the electronic capture, storage, distribution, archiving and accessing of documents.”

The key to his analysis lies in his definition of the requirements for effective management information processes. For him, “A good information management program in an organisation will manage the full lifecycle of information ranging from creation to use. For example:

- When generating transactional data, steps will be taken to ensure that the data will be stored following database “normalization rules” to promote data integrity, the single sourcing of data, the reduction of wasted database space, and fast transaction processing.

- When acquiring information, such as the purchase of market research data or competitor intelligence information, steps will be taken to reduce duplicate purchases and to increase the accessibility of any purchased data and information across the enterprise.
- Any data or information that is stored will be adequately protected against unauthorized access, as security, privacy and copyright concerns exist.
- Data and information stewardship programs will be set up to identify those organisational workers or units who are responsible for the quality and management of certain data and information items.
- Data and information will be regularly backed up for recovery purposes.
- Duplicate or mirror copies of data and information items will be created to facilitate access, and reduce network congestion and/or an overload of requests on the servers on which the data and information reside.
- Old data and out-dated information will be archived and/or deleted.”

This definition is complete and addresses the essential issues of data duplication and data lifecycle management, including deletion. This latter point is a particular issue when trying to classify, clean and prepare datasets that have accumulated over time.

Detlor makes one other comment that is of interest. He writes that “nevertheless, much confusion exists over the role IT plays in the management of information in organisations, with some equating information management primarily to the management of information technology itself.” It is important to make and maintain the distinction within organisations of information as an entity and an asset and information technology as a tool and an enabler. Merely having access to modern and efficient information systems does not mean that information itself will be appropriately managed.

Cheung et al [36] discuss the importance of taxonomies within the field of knowledge management. They begin by defining knowledge management as encompassing “the identification and mapping of intellectual assets within an organisation” and then point out that “the knowledge of an organisation does not reside in structured databases only. Most of the information or knowledge in an organisation is unstructured or semi-structured such as email communications, which contain much human knowledge and details of customer relationships related to daily operations. Many companies realize the value of the knowledge inherent in unstructured information which makes up to 80–98% of all the data, information and knowledge in an organisation.” They then report that “only a limited number of companies can provide an effective way for the acquisition, production and sharing of knowledge extracted from this unstructured information. This is due to the fact that knowledge in unstructured information is inconvenient and difficult to capture and share”, for a number of reasons including that “some sort of textual analysis is often required in order to organise, classify, and manage the unstructured information and place it within a conceptual framework for acquisition and sharing. Moreover, the unstructured content usually appears in a number of different formats and this makes it difficult to access. The lack of structure makes it difficult to organise its content.”

This then is the core of the unstructured data problem. Many organisations are aware of its existence and of its potential value, but actually managing it in any practical way is a significant challenge. A core

component of the solution proposed here is in the application of an effective taxonomy “which can systematically organise and classify the unstructured knowledge from the mass of textual documents and other items. A good taxonomy is also crucial for the effective management of unstructured information. A taxonomy can be used for creating metadata, i.e. special words to describe an object for information retrieval, special terms to describe categories supporting browsing navigation, or schemas governing Web page layout and structure, and data control lists used in support of data mining.” The argument is compelling: information cannot be structured and sorted without some kind of framework, and given the variety of formats of the source files, consistent and enforced metadata is a pragmatic and structured approach to being able to record the nature and contents of data elements.

Soibelman et al [37] discuss the management of unstructured data in the specific context of the construction industry. The importance of the problem is described: “the large majority of construction documents are available as complex unstructured data sources, such as text documents, digital images, web pages, and project schedules.” Issues specific to that industry are identified, but the same principles apply to all uses of unstructured data, especially when trying to create or apply metadata as a means of managing and referring to data. “Traditionally, research in construction data analysis has been focused on transactional databases, in which each object/instance is represented by a list of values for a given set of features in a data table or a spreadsheet. Comprehensive Knowledge Discovery in Databases (KDD) processes, as well as many statistical and machine learning algorithms, have been introduced and applied in previous research.” The use of such databases is necessarily restricted, however. “Currently a large percentage of project information is still kept in these sources for two major reasons. First, the large numbers of features, or high dimensionalities of such information make it inefficient to store and manage related data, such as text files with hundreds of words or images with thousand of pixels, in transactional databases. Second, information contained in the original data formats, such as graphical architectures of schedules and texture information of images could be lost if they are transferred into a single data table or spreadsheet.” This is a fundamental challenge in the handling and management of all unstructured data, whether it is possible for any taxonomy and system of managing metadata function efficiently, being simple enough to allow efficient classification and search without losing too much detail and becoming too abstracted from the underlying data.

Summary

Unstructured data have been the subject of a great deal of discussion and analysis over recent years. There is still a tendency to define such data in term of what they are not, as results of both the wide range of files and datasets that could be considered to consist of unstructured data and of the far greater familiarity that computer scientists and technologists have with structured data, standard formats and database schemas.

The management of unstructured data is clearly a non-trivial problem for organisations. As a starting point, identifying the data can frequently prove to be a time-consuming and complicated task requiring consistency and resources. Classifying the data is the next challenge, and although different approaches exist, such as the creation and use of metadata, this is not automatically a straightforward and efficient process.

3.3 Big data: definitions, significance, use and management

Steven Baker, in his book “The Numerati” [38] cites some interesting numbers to start to give an idea of the scope and scale of the phenomenon of big data: “in a single month, Yahoo alone gathers 110 billion pieces of data about its customers” (p5). He then summarises the purpose behind the accumulation and processing of big data: “The key to this process is to find similarities and patterns” (p6). Examples are given of the power of data accumulation and mining, showing how modern techniques and technologies could have a significant impact on privacy: “A Carnegie Mellon University study recently showed that simply by disclosing gender, birthdate, and postal zip code, 87 per cent of people in the United States could be pinpointed by name” (p13). He delves into the history of the subject, beginning with its first theoretical and mathematical underpinnings at a time before digital computers: “For decades, IBM researchers have been transforming bigger and bigger pieces of the company’s business into math. The science they use, known as operations research, was born during World War II” (p30), going on to make reference to the simplex algorithm and optimization (p31-32). He also underlines the importance of metadata about communications, knowing who is communicating with whom, whether in the context of mail they are using “to” or “cc” to address messages, and discusses where this is leading (p35). He then introduces the next layer of information, the aspect that makes the handling of big data so interesting and potentially so valuable: “But if each one of the items you bought carried a bit more contextual information, what computer scientists call a layer of “semantic” detail, much more of you would pop into focus” (p44).

Baker describes the need for updated and powerful technologies to handle the quantities of data that are being generated, although the essential nature of human input in designing systems and coaching behaviour cannot be ignored: “Faced with such complexity and contradictions, machines need smart and patient teachers to guide them in making sense of us” (p45). He then discusses microtargeting and then considers the fundamental role of context and meta-information in communications: “We humans – each of us carrying a prodigious brain wired for communication – misread each other’s words and gestures daily...If you listen, we’re constantly amending what we say. Getting our meaning clear in someone else’s head and understanding what they’re trying to say is a struggle to which we devote an enormous amount of our intelligence” (p112). This needs to be replicated or replaced in some way when handling large quantities of data, for otherwise there is always the risk of accumulating apparent but pointless or invalid connections. He goes on to discuss the frequent problem in many contexts, and especially in the field on data analysis, of mistaking correlation for causation (p116).

A key question he asks is whether there are times when you can have too much data (p128). With the enthusiasm to accumulate and merge ever increasing numbers of datasets, and with the increasing availability of datasets from a wide variety of sources, there is likely to be a tendency to throw so much data together that spurious connections overwhelm the valid findings and processing time expands to become a limiting factor.

Baker also refers to the phenomenon of terrorist and criminal planning circulating on the net: “This information didn’t travel highly encrypted on secure networks. Most of it was mingling freely with the rest of the globe’s chatter” (p130). This suggests that use can be made of communications data for intelligence and security purposes, although the significant challenge for the relevant authorities will be to identify what is important for them, a needle in a haystack scenario. He widens the discussion to consider the importance of this for ordinary, civilian, every-day users of the Internet and its services, as if there will be data gathering performed by various agencies, and the PRISM revelations of 2013 demonstrate that this is

the case, “We’re going to have to re-evaluate our ideas about privacy and secrets. We all have different kinds of secrets. Some things we tell no one. Others we share with family and a friend or two. Many are secrets in name only; we blab about them all the time. But until recently, our secrets were scattered... Unless a detective was on the case, the bits of information didn’t find each other. Now they can and will” (p204).

Data merging and mining mean that every user has the responsibility to ensure the level of privacy that they feel they require, assuming that they have the necessary awareness, skills and time. This obviously is a major practical concern, for three major reasons. Firstly, individuals are very rarely aware, with good reason, of the kind of data and metadata that they publish or leave behind them on the Internet; deliberate and knowing postings or transfers of data may be recognised as such and remembered, but much interaction online is not recognised as data transfer by the average user. Secondly, users have no visibility over the kinds of metadata and connection data that is generated every time they log onto a network or connect to a site. Again, the average user is unlikely to know what form these data could take. Thirdly, users have no idea how, where, or by whom all the data relating to them are stored. Taken together, these reasons mean that in practice the idea of users actually exercising the responsibility they have over the data that concern them, in order to ensure an appropriate level of privacy, is prohibitively problematic.

In a white paper on Big Data published in 2013 [39], ISACA set out the background to the Big Data trend. They describe it as “a trend in technology that is leading the way to a new approach in understanding the world and making business decisions. These decisions are made based on very large amounts of structured, unstructured and complex data (e.g., tweets, videos, commercial transactions) which have become difficult to process using basic database and warehouse management tools.” In referring to the size and nature of the datasets in question, they conclude that “big data refers to data sets that are too large or too fast-changing to be analysed using traditional relational or multidimensional database techniques or commonly used software tools to capture, manage and process the data at a reasonable elapsed time.”

The nature of personal data is also defined, as this has significant impacts on the way datasets are used and how they should be managed. According to this definition, personal data can be categorized as:

- Volunteered data – data created and deliberately shared by individuals (for example, social network or forum profiles)
- Observed data – data that are captured when recording the actions of individuals (for example, location data when using cell phones or when taking digital photographs)
- Inferred data - data about individuals derived from the analysis of volunteered or observed information (for example, credit scores or purchasing patterns).

Quite apart from the immediate business impact and effects on process models, the use of big data can affect the following activities:

- Governance – determining what data should be included and how data governance should be defined and delivered
- Planning – anticipating the effects on processes, resource management and infrastructure strategies and management

- Utilisation – determining what use will be made of the data and how the necessary infrastructure should be provided and managed, whether this is in-house or outsourced
- Assurance – knowing how to obtain comfort over the quality and consistency of data
- Privacy – determining how data privacy will be ensured, protecting data from unauthorised users and over-controlling governments in every country where the data will be stored, processed or transferred.

The report identifies three key questions that management should be asking in respect of big data. These are:

- Where should we store the data?
- How are we going to protect the data?
- How are we going to utilize the data safely and lawfully?

The concept of big data risk management is still in its infancy for many organisations and that security policies and procedures are still developing in response to the new challenges. The importance of effective corporate data lifecycle processes is made clear, as “inaccurate, incomplete or fraudulently manipulated data pose an increasing risk as enterprises become more dependent on the data to drive decision making and assess results.”

Data vulnerabilities is a key issue to be addressed, especially in respect of the increased exposure to such vulnerabilities of organisations that rely upon the processing of personal data for their value. A number of uncertainties need to be addressed in order to lessen the impact of these vulnerabilities:

- Privacy - individual needs for privacy vary, as do approaches to adopting and enforcing legislation and regulations
- Global governance – the issue is complicated by the absence of global legal interoperability, with each state developing its own legal and regulatory frameworks, each with different scopes and ranges of applicability
- Personal data ownership – historical concepts of property rights are not easily extended to data, thereby creating challenges in establishing rights over ownership and usage
- Transparency – there is a fine balance to be struck between too much and too little transparency
- Value distribution - before value can be shared more fairly, there is a need for clarity on what constitutes value for each stakeholder.

In order to reduce the potential for damages caused by the use of and reliance upon inaccurate, incomplete, outdated or fraudulent data, organisations “should take inventory of all the data sources they are pulling into their analyses and assess each source for vulnerabilities.” Analysis of sources, user accesses, and motivations for manipulation should be performed to identify issues before data are used or further transmitted.

A number of strategies for addressing big data risk can be defined. These include:

- Aligning the technology solution to business needs, ensuring that stakeholder drivers are aligned with stakeholder needs and allowing existing governance structures to be adjusted to the particular security and assurance requirements of the new technology
- Ensuring the quality of the data, considering the three axes of intrinsic quality (accuracy, objectivity, believability, reputation), contextual and representational quality (Relevancy, completeness, currency, appropriate amount of information, concise representation, consistent representation, interpretability, understandability, ease of manipulation), and security and accessibility quality (availability/timeliness, restricted access)
- Choosing the right partner for outsourcing and/or cloud services

Appropriate data governance is an essential part of the overall corporate governance process. It can be argued that without a proper data governance process, big data projects can present organisations with a number of problems, including misleading data and unexpected costs. The advantages of effective data governance are obvious, however, as data governance programs can provide a framework for setting data-usage policies and for implementing controls that are designed to ensure that information remains accurate, consistent and accessible, thereby addressing in part the main objectives of information security.

Assurance considerations for big data can be grouped into four categories.

- Approach and understanding incorporates elements of the tone at the top principle seen in the COSO framework and requires the creation and implementation of a data policy that defines the data in scope, establishes a system of governance and assurance over data quality, and identifies both qualitative and quantitative criteria for evaluating the accuracy, reliability completeness and timeliness of data.
- Data quality should be ensured by internal controls across the data flow to evaluate data against the criteria set out above.
- Data confidentiality and privacy should be enforced by controls applied to data identified as sensitive during the data risk management process. Reference should be made to relevant laws and regulations such as the UK Data Protection Act and the US PCI-DSS, while the controls should include the usual ITGC elements such as passwords settings, data masking, physical security and encryption.
- Data availability should be ensured by the implementation of tested disaster recovery arrangements that correspond to the organisation's data recovery point objective (RPO) and recovery time objective (RTO) criteria as defined in a business impact analysis.

Summary

A major part of the challenge resulting from the desire to use big data will reside in ensuring that big data risk and concerns are not viewed as an exclusively information technology exercise. It is also important to perform periodic reviews of big data strategies and security policies and procedures, even if these appear to represent unnecessary costs. Essentially there needs to be high-level buy-in into the whole concept by

senior organisational management, and they will need to ensure that the right resources are made available, and the right authorities and responsibilities attributed and exercised, across the organisation.

3.4 Data storage and processing

Manyika et al [40] talk about developments in technology and processing. “The ability to store, aggregate, and combine data and then use the results to perform deep analyses has become ever more accessible as trends such as Moore’s Law in computing, its equivalent in digital storage, and cloud computing continue to lower costs and other technology barriers” (p2), clearly identifying some of the key technical drivers of the rapid recent development of potential uses of data. They report that: “MGI estimates that enterprises globally stored more than 7 exabytes of new data on disk drives in 2010, while consumers stored more than 6 exabytes of new data on devices such as PCs and notebooks. One exabyte is the equivalent of more than 4,000 times the information stored in the US Library of Congress” (p3). This gives some idea of the scale of the phenomenon; other statistics were previously presented in Chapter 2.

Possible techniques for analysing big data are introduced, discussing how these techniques have been developed and continue to evolve as requirements change and technologies offer greater speed and greater opportunities. This analysis is interesting as it shows how techniques from various fields (such as statistics and computer science) can be combined with data obtained from various sources (formal databases, crowdsourcing, usage data) to provide a rich range of possible analyses and means of visualizing and interpreting the results.

It is always essential to have a clear definition of terms when discussing technologies and applications. The key items in this analysis can be summarised and presented as follows:

Term	Definition
Business intelligence (BI)	A type of application software designed to report, analyse, and present data. BI tools are often used to read data that have been previously stored in a data warehouse or data mart. BI tools can also be used to create standard reports that are generated on a periodic basis, or to display information on real-time management dashboards, i.e., integrated displays of metrics that measure the performance of a system – this is an example of existing processes and business requirements that are being revolutionized by the adaption of newer technologies.
Cloud computing	A computing paradigm in which highly scalable computing resources, often configured as a distributed system, are provided as a service through a network.
Extract, transform, and load (ETL)	Software tools used to extract data from outside sources, transform them to fit operational needs, and load them into a database or data warehouse
Metadata	Data that describes the content and context of data files, e.g., means of creation, purpose, time and date of creation, and author
Semi-structured data.	Data that do not conform to fixed fields but contain tags and other markers to separate data elements. Examples of semi-structured data include XML or HTML-tagged text.
Structured data.	Data that reside in fixed fields. Examples of structured data include relational databases or data in spreadsheets.

Term	Definition
Unstructured data.	Data that do not reside in fixed fields. Examples include free-form text (e.g., books, articles, body of e-mail messages), untagged audio, image and video data.

Table 3.1 Definition of key terms

Such definitions are particularly important when discussing and trying to describe the scope and nature of the different data structures and contents that are being accumulated and processed in different environments.

When using big data, any project needs to follow traditional project management processes. “Organisations should begin a process of identifying and prioritizing these opportunities; business line and function leaders should kick off processes in their respective areas of responsibility. It is worthwhile noting that identifying opportunities and potential sources of valuable data (see previous section), especially external sources of data, will often be an iterative, rather than sequential, process” (p112). It is emphasised the fact that the use of big data for business intelligence is never going to be a one-stop, short-term project, as requirements will evolve and additional or replacement data sources will always become available.

The fundamental concepts of privacy and security, which dominate the controlled and managed use of real-life data, need to be systematically addressed. “Addressing privacy and security issues will become paramount as more data increasingly travel across boundaries for various purposes. Privacy, in particular, not only requires attention to compliance with laws and regulations, but also is fundamental to an organisation’s trust relationships with its customers, business partners, employees, and other stakeholders. Certainly, organisations will need policies that comply with privacy laws and any government privacy regulations.” This is a useful summary of what could be considered a top-down approach to security and privacy, considering above all other factors the legal and regulatory aspects of data management.

In practice, however, there are other requirements. “But in developing a privacy policy, organisations will need to thoughtfully consider what kind of legal agreements, and, more importantly, trust expectations, it wants to establish with its stakeholders. And it will need to communicate its policies clearly to its stakeholders, especially customers, as they become increasingly savvy and concerned about what is known about them and how that information can potentially be used” (p116). This recognises that the push for clarity and demonstrable conformity will come from a combination of regulators, data providers and the subjects of the data being used, particularly in a climate of increased sensitivity about what data are out there and what they are being used for. Given that users and the subjects of data will increasingly be encouraged to identify their own security and privacy requirements and work towards ensuring that these are met, a philosophy of cooperation and understanding would appear to be a good starting point for addressing these issues in a mutually satisfactory way.

Bollier [41] notes that “The growth of cluster computing systems and cloud computing facilities are also providing a hospitable context for the growth of inferential data techniques” (p3), referring to [42]. He talks about the motivations for researching and investing in improved data handling and analysis, citing real cases from industry: “Baker noted that companies are often built “on revenue streams that come from imprecise data methods that are often wrong.” The company may or may not need to decide whether to “move from what works to truth.” It may not be worth trying to do so. This leads Baker to wonder if “truth

could be just something that we deal with in our spare time because it's not really part of the business model.'" (p12).

He then discusses specific cases of the problems associated with large-scale and distributed data handling: "Jeff Jonas of IBM Software Group noted, that "There is no outbound record-level accountability—organisations transfer data out and they don't know where they sent it. That means when they go and correct something, they actually don't know who to tell to correct it downstream. That's how things work in most every information-sharing system. Recipients know where they get their records, for the most part, but the issuer doesn't. So they can't keep the ecosystem of data current." One partial solution to the problems information sharing and data protection, suggested Jonas, is a process that he calls "analytics in the anonymized data space." By this, he means that data holders who intend to share data with another party must deliberately "anonymize" the data first, before sharing it. Then analytics are performed on the data while it remains in the anonymized form. Such analytics are emerging and in some settings they are producing materially similar results as analytics performed on clear text" (p31-32).

This is an interesting approach because it requires a great deal of proactive and carefully managed data cleaning before data can be used, an effort that would require structured, monitored and reviewed processes to identify and classify the types of data held, the full extent of their contents, the nature of sensitive or confidential fields, and analysis of what could be removed, masked or hidden to ensure that the dataset retained its value. Such processes would surely be expensive and time-consuming to design and implement and require no small amount of manual intervention, at least at the design and monitoring stages.

Bollier then discusses conference input from specialists in data usage. "Greg Skibiski of Sense Networks believes we need a "New Deal on Data", referring to [43]. "By this, he means that the end users should be able to own their data and dictate how it is used. This should apply to "any data that we collect about you and metadata that we make out of it," he said. He also urged that data should have a "lifespan," so that it is routinely purged after a given period of time. Otherwise, data that is saved is more likely to be abused." This is a critical issue in data management, ensuring consistency and coherence. "But Kim Taipale has doubts that such an "ownership" standard would be practical given the scale and range of Big Data uses today. The point should be to empower users to have greater control over the access and use of their data. One of the best ways to prevent abuses, said Stefaan Verhulst, is for companies to conduct "information audits" so that they only collect the data that they need in the first place. By asking better questions upfront, companies will decrease their legal vulnerability and the likelihood of privacy violations" (p35-36). There are two interesting ideas emerging from this: the notion of user ownership of data, which is surely desirable but not necessarily easily achievable; and then the importance of auditing to ensure that proper controls are in place.

Summary

The storage and processing of big data is fraught with risks and complications. It is clear that structure and order need to be designed and implemented before data are acquired and processed, because ad hoc attempts to impose structure on existing collections of data are likely to prove prohibitively complex. Organisations need to be able to identify their internal requirements and external obligations before accumulating and processing data in order to ensure that objectives and requirements are being met in a systematic and verifiable way.

3.5 The Security of Unstructured Data

In a magazine article Talmor [44] discusses the applicability and limitations of data-centric security. He describes the reinforced perimeter approach adopted in traditional security methods: “The continuously escalating and mutating threat environment has led many firms to layer security countermeasures one upon another; starting with firewalls, companies have added intrusion detection systems, malware filters, client-side firewalls, and encrypted network tunnels. Networked business can create a virtual fortress around its infrastructure but still must share information with mobile employees, external business partners, and remote customers.” This summarises very briefly the current situation and its most immediate significant limitation, which is that security measures are systematically applied at the perimeter of organisations, although the flow of information and data in and out of organisations is a core element of daily activities.

He also refers to the scale of the unstructured data question in modern businesses: “In most organisations, 70-90% of business data is in an unstructured or semi-structured state and recent research indicates that only 23% of organisations feel this data is properly protected. Unstructured data includes Word and Excel documents, images, BLOBs (Binary Large Objects), not to mention the billions of emails and instant messages generated every day. Much of this is sensitive data, such as personally identifiable information (PII) and intellectual property (IP) that must be protected with appropriate measures.” Although the sources of the numbers mentioned are not given, these numbers are broadly consistent with other sources quoting the findings of research. It is also interesting that instant messaging transcripts are included in the definition of unstructured data, as this is a type of data not frequently included in other definitions and examples. Given the popularity of the technologies, however, and its free text nature, it should clearly be included.

The change in philosophy in respect of the security of information is then discussed. Given that data will be entering and leaving the organisation on a regular basis and thus rendering perimeter security inappropriate, Talmor writes that “Many businesses now realise that rather than continuing to add layers of infrastructure security, it’s more effective to protect critical data throughout its lifecycle, regardless of where it resides or moves. This concept of protecting data rather than devices is known as data-centric security.” The logical consequence of this is that data-centric security must provide protection for data at rest (storage) and in transit. In order to do so, unstructured data that require protection are encrypted before they are transferred or stored.

Emphasising that in order for security to be maintained, data that have been decrypted for processing need to be systematically destroyed after use, Talmor insists upon two additional requirements for data-centric technologies. Firstly, they must include data rights management, real-time strong authentication and encryption, in order to be complete. Secondly, they need to be easy to use systematically. “Most users concentrate on getting their work done, not on the underlying technology powering that work. And when security solutions are deemed too difficult to use, many users will circumvent the solution as well as the security.” This latter point of course applies to all security and internal control measures, but is likely to be more exaggerated in respect of technologies such as encryption where the complexity is likely to be visible and daunting for ordinary users.

In a 2008 white paper Kelley [45] also considers the challenge of protecting unstructured data, similarly viewing things from a data-centric perspective. She discusses the phenomenon of de-perimeterisation as

information systems and their uses evolve, and contrasts the three approaches to data security that have been most favoured over the last twenty years:



Figure 3.1 Three approaches to data security, from Kelley (2008)

(Spelling mistake in the source material)

Perimeter access controls fell out of favour when networks and infrastructures started becoming widely interconnected, while application specific controls were found to be complex and expensive without necessarily providing the full range of protection that was required.

Kelley presents a simple three-step approach to data-centric security.



Figure 3.2 A three-step approach to data-centric security, from Kelley (2008)

This model is very straightforward and high-level and is based on first principles. It does address the key issues, though, insisting that the planning and analysis of data types are performed before data are located and classified, and only then can thought be given to how best to protect the data.

There are a number of business and technical requirements related to data-centric protection approaches.

Among business requirements can be included:

- Compliance – ensuring that the solution chosen is consistent with relevant requirements such as PCI DSS
- Time to Deploy – ensuring that the lead time for implementation is appropriate and practical
- Business Process Alignment – ensuring that the solution corresponds to and does not disrupt existing business processes
- User Experience – ensuring that the solution is invisible to and empowers users
- Cost – ensuring that costs, including hidden costs such as additional strains on the helpdesk, are manageable and appropriate.

Among technical requirements can be included:

- Availability – the solution chosen must work reliably across all platforms and infrastructures
- Architecture – the solution must work within the existing architecture and be sufficiently robust to adapt and grow alongside the core architecture
- Extensibility/Interoperability – the solution needs to demonstrate flexibility and extensibility going forwards and permit transparent integration with other core systems
- Certifications – if the operating environment requires certification of the quality and reliability of the encryption methods in place, the solution should be certifiable
- Scalability – the solution should be scalable as user populations and infrastructures grow.

In a 2010 white paper on emerging trends in information security [46], Verizon highlights the importance of unstructured data for modern businesses. They note that “Organisations are beginning to see the power in extracting intelligence from unstructured data. This is now being done by combining expressions of both traditional, structured data, such as a customer list, with unstructured textual data, such as call center notes. New semantic tools enable textual analysis to uncover trends in conversations. The exploding numbers of social networking and Web 2.0 tools, including blogs, RSS feeds, Twitter, and Facebook, represent increased information opportunities for marketing, product, and customer service managers. This contextual intelligence gained from the combination of structured and unstructured data yields richer insights that help with analyzing the correlation between customer sentiment and buying behavior.” This is a positive analysis of the potential benefits to be gained from mining unstructured data.

The report discusses the risks associated with the retention and handling of unstructured data. “Unstructured data and increased data volumes present major risks. The first, and probably most likely, is the unintended leakage of secretive information and intellectual property. This information, often referred to as “secrets,” has intrinsic value for the company. It comes in the form of new product details, source code, construction plans, factory layouts, sales and revenue forecasts, and other types of competitive

intelligence.” This is key and a fundamental reason why unstructured data need to be protected, even if their precise nature and contents are not immediately obvious to management and users.

The report presents a three-stage process for addressing the issue. It recommends starting with a data governance program, performing a data discovery exercise, and then establishing policies for data use and protection. It is interesting that there is no mention of carrying out a structured and tailored risk assessment before establishing the policies, as good practice would suggest, but this is surely a fundamental part of the model.

An online article for Channel Insider [47] refers to a white paper produced by the Aberdeen Group in which eight steps towards securing unstructured data are proposed. These are: identify and classify data; prioritise security control objectives; establish consistent policies; select and deploy data protection solutions; invest in documentation, awareness and training; assign clear ownership and accountability; automate enforcement; and measure and monitor. These steps are consistent with long-established good practice for managing data security and emphasise the specific challenges presented by unstructured data in respect of identification and ownership.

The Aberdeen Group report [48] itself presents a number of interesting insights into the scale and scope of the problem. The report, published in June 2009, provides statistics of the behaviour of what the author refers to as “Best-in Class” performing firms.

The report offers some illuminating statistics. Respondents to the survey estimated that on average 40% of their data were in unstructured format, with a breakdown possible by file type:

File type	Percentage of respondents (n=159)
Instant messaging	46%
Custom / in-house	47%
Video files	49%
Audio files	53%
Still images	67%
Web pages (HTML, XML)	81%
Email	90%
MS Excel spreadsheets	91%
MS Powerpoint presentations	91%
MS Word documents	92%
Adobe Acrobat / PDF	93%

Table 3.2 Breakdown of unstructured data by file type

From these results can be noted the ubiquity of PDFs and of Microsoft Office documents and of mail and web page stores, and also the importance of multimedia files whose contents may be more difficult to ascertain systematically through standard file interrogation tools.

The report compares and contrasts the drivers for investing in protecting and managing unstructured data with the leading inhibitors for doing so. The main drivers are to protect the organisation and the brand, internal policies, industry standards and best practices, government regulations, the increased mobility of unstructured data, risks associated with external sources, and collaboration within the organisation. There is a clear shift in the perception of IT Security from being an obstacle to being an enabler.

The difficulty for organisations, and for IT Security groups within organisations, is in transferring these drivers into actions. Six leading inhibitors are identified:

Inhibitor	% of respondents
Not viewed as a priority relative to other current initiatives	30%
Budget is not available	30%
Responsibility and ownership are dispersed among different groups	28%
Lack of consistent policies for unstructured data	23%
Inability to discover / identify unstructured data	13%
Maturity of currently available solutions	11%

Table 3.3 Leading inhibitors in respect of security

These results show how difficult it can be to drive change within an organisation unless risks are properly appreciated and the rewards can be adequately demonstrated. There also needs to be cultural change in many organisations, with data ownership and identification being properly enforced and reinforced.

The research identifies a number of ways in which organisations are seeking to address the security issues and divides them into three categories, as follows:

Category	Representative solution
Monitoring/filtering	Email monitoring/filtering, web monitoring/filtering, database activity monitoring (DAM), data loss prevention (DLP)
Encryption	File/folder encryption, full disk encryption (FDE), USB drive encryption, email encryption, secure file transfer, encryption key management (EKM)
Management	Collaboration tools, enterprise content management (ECM), web access management (WAM), security information and event management (SIEM), data discovery tools, data classification tools, data governance tools, enterprise rights management (ERM)

Table 3.4 Leading solutions to security issues

It can be seen that a wide range of approaches and technologies are being adopted, with no one-size-fits-all solution being apparent. It is also interesting that so much attention is being paid to management tools and solutions, rather than the focus being more exclusively on technological approaches.

After identifying the companies that it considers to represent the Best-in-Class, the report discusses the most consistent activities undertaken within a PACE (Pressures, Actions, Capabilities, Enablers) framework to address the issue of managing and protecting unstructured data.

Pressures	Actions	Capabilities	Enablers
<ul style="list-style-type: none"> • Protect the organisation and its brand 	<ul style="list-style-type: none"> • Identify the unstructured data that exists throughout the organisation • Establish consistent policies related to unstructured data • Educate end-users about security, compliance, and management policies and practices for unstructured data • Enforce policies by implementing technologies / solutions that are being standardized enterprise-wide 	<ul style="list-style-type: none"> • Consistent policies for access and distribution • Consistent policies for actions on data • Standardized response for exceptions, security events, or incidents of non-compliance • Clear understanding of who has accountability for authorizing access • Responsible executive or team with overall ownership • Discovery and classification of data • Grant / modify / revoke user access privileges to hosting sites and specific data (e.g., files) within hosting sites • Monitor when and by whom sensitive data are created • Monitor access by privileged users • Assign access privileges based on job role; modify upon 	<ul style="list-style-type: none"> • Web monitoring / filtering (73%) • Email monitoring / filtering (65%) • Database activity monitoring (59%) • Data loss prevention (30%) • Secure file transfer (79%) • File / folder encryption (49%) • Full-disk encryption (49%) • Email encryption (46%) • Enterprise key management (43%) • USB drive encryption (36%) • Endpoint device / port controls (32%) • Web access management (71%) • Collaboration tools (63%) • Enterprise content management (47%) • Data discovery tools (38%) • Enterprise rights management (33%) • Data governance tools (30%) • Data classification tools (27%)

Pressures	Actions	Capabilities	Enablers
		change; revoke upon termination <ul style="list-style-type: none"> • Controls to monitor and verify that the requirements of internal policies and external regulations are being satisfied 	

Table 3.5 Measures to address security concerns

This method of structuring analysis allows a clear sequence to emerge from defining objectives, scoping high-level activities, determining lower-level levels through to identifying suitable tools and technologies. In this respect it is similar to the analysis leading to the design of control activities described in Chapter 2. It is also interesting to note how significant the high-level management aspects of the security management process are and how they emphasise the need for organisation and structure in the approach, as well as promoting non-technological elements of security management.

The report concludes with the list of eight steps to success highlighted in the press coverage of the report. It highlights these as an aide for decision-makers to ask the right questions about proposed solutions and implementation plans, and asking these questions both internally of colleagues and management and externally of service providers.

In a 2006 magazine article McAdams [49] considers the growing interest in “the new breed of data classification or information content management (ICM) offerings, which promise to help set policies and access controls on sensitive data buried in unruly, unstructured data sets. Vendors are positioning ICM storage software as an alternative to labor-intensive content management or metadata tools.” The level of complexity inherent in the process of classifying and categorising data makes the use of tools to at least partially automate the process very attractive. But such approaches should not be seen as a choice that will eliminate additional work for key staff, as expert users will still have to identify and classify datasets as a whole as well as quantities of individual data items. There should also be consideration of whether trying to apply structure to unstructured data is necessarily the most appropriate approach in terms of meeting business requirements. In some cases a more traditional content management approach might better meet requirements.

In an online article published in 2010, Lorenz [50] discusses the challenges in respect of securing unstructured data. He emphasises the difficulties that organisations often find in answering three straightforward questions: what data is sensitive to my organisation; where is this data stored; who has access to this data, and is the access profile appropriate? He argues that there are three primary dimensions of unstructured data risk and that “understanding and quantifying these three dimensions are the first steps towards executing a successful unstructured data protection program.” These dimensions are:

- Data sensitivity: impacted by the organisation’s regulatory environment, risk appetite and corporate culture;

- Sensitive data location: viewed in respect of both physical infrastructure and logical organisation; and
- Inappropriate access: defined as users and administrators having access to data without business or operational justification.

Partial solutions proposed by the author include adopting a logical and consistent approach to access control, building upon existing processes rather than developing new processes and procedures from scratch, and ensuring senior management involvement and commitment to any data management initiatives, driving change and enforcing compliance while demonstrating the advantages to be gained from the project.

In a commercial white paper published in 2010 the Aprigo Corporation [51] discusses the risks of data breaches and suggests methods for preventing them. The authors focus on the financial and other costs that can result from data breaches: “someone who knowingly violates the Health Insurance Privacy and Portability Act (HIPAA) prohibition against disclosing individually identifiable health information can be fined up to \$250,000 or imprisoned for up to 10 years – or both. The cost of a data breach goes far beyond measurable costs such as hefty fines, though, extending to the cost of:

- Litigation that is likely to ensue when private information is shared publicly
- Loss of competitive advantage when trade secrets are revealed
- Degraded reputation when sensitive company information is made public
- Compromised integrity when salary or HR files are leaked”

The risks are therefore significant, and organisations need to be able to manage and control their unstructured data. The authors highlight what they consider to be the six most significant obstacles to securing unstructured data:

- “Multiple potential points of failure. Just one terabyte of unstructured data can contain millions of data objects, or files. Organisations of all sizes face the epic challenge of protecting hundreds, thousands, millions, or even billions of data objects – every one of which, if not protected appropriately, carries the potential for a data breach.
- Effective access rights managed in multiple systems. Two distinct, hierarchical structures govern user access rights: the file system with the Access Control Lists (ACLs) and the directory services (Active Directory) with both ACLs and objects. Any change in file system ACLs results in a change in user access rights, as does any change in group membership in the directory. The challenge is in coordinating the two.
- No central authority or control over access. With structured databases, only the database administrator has access to actual data records. But with unstructured data, every user in the organisation is free to manipulate the data and the permissions for accessing it. For example, a user from the finance group can add an Access Control Entry (ACE) for another user from a different department, or even expose data altogether by adding the “everyone” group ACE.

- Distributed environments. Mid-size and large organisations typically operate distributed computing environments, in which there are multiple file servers in multiple locations. To complicate things further, today there are also likely to be not only physical servers, but also virtual servers, creating even more points of management.
- Lack of auditing capability. Being able to understand users' effective access rights is difficult enough when dealing with multiple files in distributed environments; being able to audit how permissions change at different points in time is impossible with today's file server systems. Ideally, one should have the ability to have multiple point-in-time snapshots of access rights, just as one would want to have multiple point-in-time backups of data – particular for folders and files containing sensitive data.
- Potential business disruption of remediation. While there are several methods available to organisations to control and secure sensitive data, implementing them within an organisation can be an expensive and disruptive proposition that may require significant IT investment in dedicated hardware resources, ongoing technology updates, and associated activities.”

These obstacles combine to make the task of securing data, and having the knowledge and confidence that data are secure, complex and challenging.

In response to these challenges, three methods are proposed for securing unstructured data. These are:

- Physical Access Controls. The first line of defense is physical security, or making sure that file servers and other data storage media are in locations where access is restricted to only the appropriate staff.
- Network Access Controls. File system access control lists can be used alongside directory service to ensure a “least privileged” access policy, or one in which users can access *only* the information and resources that their roles within the organisation require.
- Data Encryption. Encryption of data at rest and in transit gives a layer of comfort in respect of security, because if properly implemented, unauthorized access to datasets is not an immediate problem for security managers unless the unauthorized users also have access to the decryption keys.

These techniques present a solid and theoretically sound starting point to address the challenges, but are by their nature incomplete and high-level.

An online article in SC (Secure Computing) Magazine in 2008 [52] highlights the risks to data security presented by inappropriate access rights within an organisation. The article refers to research by Ponemon in which 89% of respondents “admitted that controlling access to unstructured data is a major challenge.” More interestingly, they also reports that “nearly 70% also felt that access to unstructured data by employees is very often unwarranted, a situation that they are unable to rectify because they do not have the means to monitor and control access.” This is a key concern for many organisations, especially if compliance regulations are in force. Organisations need to ensure that users only have the access rights to data that are commensurate with their roles and responsibilities, but to do so requires being able to identify unstructured data, know how and where these data are stored, and apply structured, practical and appropriate user access rights.

A commercial white paper published by Imperva [53] presents five key signs for organisations that their file data might be at risk. They ask five questions of management:

- “Who owns your file data?”
- Who is actually using your files?
- Who has the potential to access your files?
- Whose access rights should be revoked?
- How do you know when access rights or activity violate corporate policy?”

Although asked in this context in respect of file data, the same questions apply equally well to unstructured data and can serve prompt reflection from management responsible for unstructured datasets. The increased risk in respect of unstructured data is that management do not have sufficient control in place over the data, or even know enough about their location, nature and use, to be able to provide reliable and coherent answers to these questions.

In a strategy paper published by InformationWeek magazine in 2010, Ely [54] discusses the fundamentals of protecting unstructured data. He cites IDC research concluding that unstructured data quantities are increasing at a compound annual growth rate of 61%.

A graphical representation of the answers to a key question about means of managing data is also provided. This is interesting because it shows how approaches differ between organisations, and also how concepts of ownership and stewardship are inconsistently appreciated.

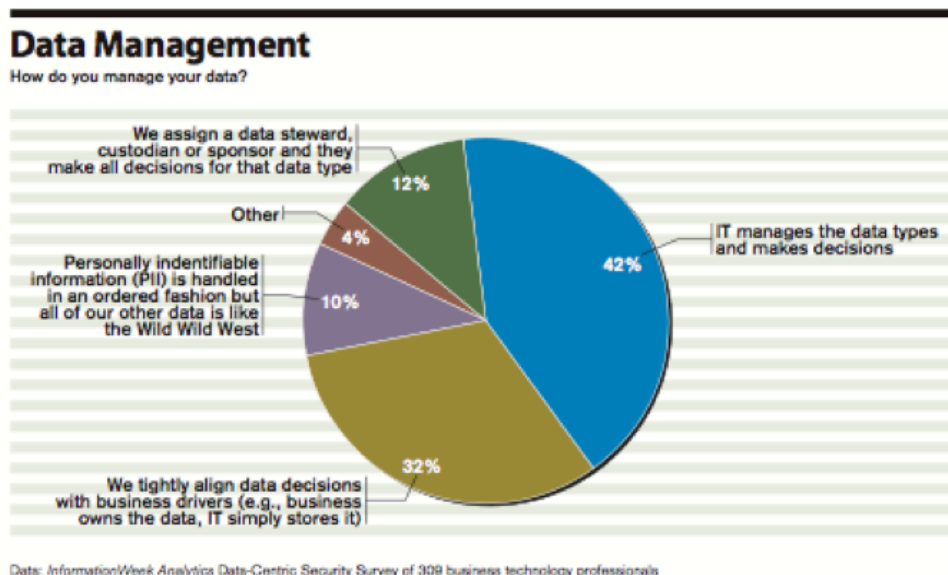


Figure 3.3 Data management approaches, cited by Ely (2010)

Ely discusses the need for multi-dimensional input into data management. “Representatives from IT, legal, compliance, HR, finance, business development and possibly other departments should be involved. If your organisation doesn’t conduct an annual risk assessment process that identifies new threats—including those stemming from data on the loose—this would be a good time to start.” This is in direct contrast to

the approach adopted by many organisations – 42% according to the survey results presented above – where data management is left to IT on its own.

Processes for identifying instances of unstructured data are discussed. Two approaches are presented: semi-automated and manual. “Begin searching file shares, laptops, connected storage devices— anywhere you can. Another approach is to ask users to review documents they own and identify those with sensitive data that needs to be protected or organised. This moves the burden from a small group of people and spreads it to a larger group, thus less effort per person. The only issue is getting people to actually do it. This process must be reinforced through user awareness of what is sensitive and risky data, what to do with this data, and whom to ask when in doubt.” This highlighted a key problem in the process, which is to train, encourage and motivate users. Such tasks will often be seen as an unnecessary addition to their workload that gets in the way of their daily tasks.

The difficulties involved in designing appropriate control measures over data are then discussed. First the application of strict and exception-free access controls over all data is considered. “Good for security? Sort of. Good for business? No. Our users need this data to perform their jobs, so making it hard to use doesn’t make sense. Or, perhaps, the user might not need the data for business, but it poses no risk. By applying extreme, overly broad controls we complicate the lives of the people who need to access and share data, and our own lives, since we must consistently update these controls.” A much better approach, Ely insists, is a measured one. “Once you identify where data is stored, the risk, and whether or not you want to protect it, establish priorities. Not all data is worth spending time on at this point—if it’s not highly valuable or sensitive, revisit it after you’ve dealt with your critical information.”

In common with many other writers on the subject, Ely then turns his attention to encryption strategies. Encryption is a cheap and widely-available technology, but its implementation varies widely from organisation to organisation, as the results of the survey indicate.

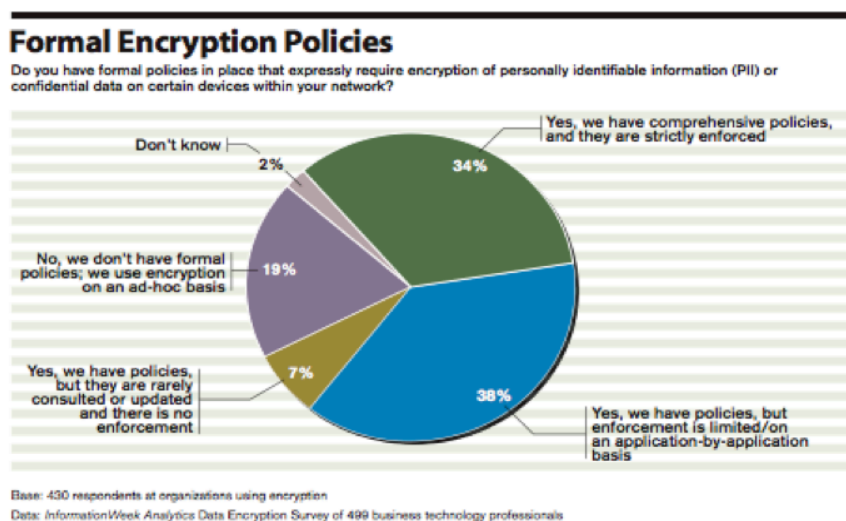


Figure 3.4 Use of formal encryption policies, cited by Ely (2010)

Only a third of corresponding organisations can claim comprehensive and wide-ranging encryption policies.

A final area discussed by Ely is the difficulty of protecting data once it has left the security perimeter. “Cloud services typically have their own security controls, but most offer no way for organisations to

interact with these controls to fine-tune them, or understand exactly where data resides. Additionally, we don't know what the cloud vendors' employees can do, or what data they can see. Could Google staffers log in to your account and read your documents? The methods currently available to IT do little to address finding sensitive data outside of the enterprise, for example, within Google Docs. If the accounts are owned by the organisation, the services may have APIs to allow IT to search documents, or they may not. The largest problem is the lack of consistency across third-party offerings and the ability to scale a single utility for searching and protecting each cloud service."

This shows the heart of the problem. Traditional methods for managing IT through internal controls reach their limits when confronted by the paradigm shift that is the transfer of services or data outside the scope of effective hands-on control.

3.6 Privacy in the Digital Age

3.6.1 Definition

Craig and Ludloff in their book "Privacy and Big Data" [55] define privacy in the digital age. They describe "Physical privacy – or freedom of intrusion into your physical person, possessions or space. Most countries have privacy laws that address unlawful search and seizures on your person or possessions. Informational privacy – your expectation of privacy when personal information is collected or, stored and shared in digital or some other format. Most countries have laws regarding the privacy of financial, medical and Internet information to some degree. Organisational privacy – government agencies, organisations, and businesses expect to be able to keep activities or secrets from being revealed to others. For example, companies may expect to keep trade secrets and governments may choose not to reveal security policies to prevent terrorism (such as "secrecy" that is codified in the US Patriot Act)" (p14). They then discuss the impact of the digital age on our expectations, what we consider private and may now be considered public: "Privacy of our communications, privacy of our behavior, privacy of our person" (p14-15).

A key distinction they make relates to context and expectations: privacy in the US is the right to be let alone (p16), while privacy in Europe is Honor and Dignity (p17), concluding that privacy is always viewed through some sort of prism (p19). They then address the question of privacy without borders: "In its truest sense, data has no borders" (p20), and draw attention to the clear conflict of legislations that the international nature of digital communications and transfers of information has made clear: "the recent admission made by Microsoft that data stored on its European servers can be handed over to American investigators without informing the individual in adherence with the US Patriot act. This is a violation of the EU's Data Protection Directive and Safe Harbor agreement with the US. In this case the Patriot Act trumps all other privacy legislation, regardless of where the data originated or where it resides" (p20). This issue is of supreme importance when considering how to control the security of data while ensuring compliance with all relevant regulations.

The reader's attention is then drawn to the key concepts of privacy versus security and safety (p36) and the notion that "Data never dies." The first is a manifestation of the age-old search for equilibrium between competing demands, while the second is a phenomenon the significance of which is becoming increasingly apparent as individuals and organisations are putting increasing amounts of information online, sometimes voluntarily through social networks and sometimes unknowingly through the use of online services. Once data are public, it is increasingly difficult to manage and delete them given the way that data can be shared

among multiple servers at multiple locations, backed up at will, and downloaded by anybody with appropriate access rights. The authors comment that “It is the regulations that prevent data collection without user consent that provide the true hope for a reasonable expectation of privacy” (p37), a statement that may sound a little naïve in the light of subsequent revelations about data handling practices, with organisations either unable or unwilling to provide clarity over their data management practices and various agencies requesting and obtaining access to data with considerable ease.

Attention is then turned to the practices of big data mining, with mention of the focus on pulling together information from multiple data sources. “The 2004 GAO report on government data mining found that more than one-fourth of all government data mining projects involved access to data from the private sector. The government has broad powers for doing so. It can access publicly available data on the same basis as any member of the public, it can contract for data, and it can exercise its unique power to issue subpoenas, search warrants, wiretap orders, National Security Letters, and FISA orders that require the product of personal data, usually in secret.” (citing [56]). The impacts of this are clearly set out: “Consumer fears over loss of privacy have been steadily rising and, unsurprisingly, are focused on the advertising industry. After all, they were the first to leverage technology and create a multi-billion dollar industry built on our personal data, and once it’s out there, it is pretty hard to control” (p72).

The key aspects to emerge from Craig and Ludloff are the definition of what is meant by privacy, which is a term used widely in both technical and popular discussions but which carries shades of nuance depending on context and environment, and the concept of the permanence of data. It is essential to bear in mind when considering the use and storage of big data that data online are not necessarily subject to expiry and cleansing, but that new entries will supplement rather than replace existing ones. Once data are available and can be consulted, corrections and deletions can only be enforced with great difficulty, if at all.

3.6.2 Responsibilities

In a white paper published in August 2013 [57], ISACA discuss the impact of the use of big data on organisations’ responsibilities in relation to privacy, in particular the privacy of personal data they may have gathered, stored, and analysed.

Big data is a phenomenon that can have significant effects on an organisation, improving decision making, shortening time to market, improving customer and service and increasing profits. The World Economic Forum describes the personal information derived from big data as “the new ‘oil’ – a valuable resource of the 21st century,” [58] and describes the analysis of these data as “the new engine of economic and social value creation.” [59]

Three major impacts of the arrival of big data are identified in the white paper.

The first concerns the requirements to comply with international privacy and data protection laws. “Currently, each region (European Union, USA, etc.), government and enterprise handles privacy and data protection in a different way. This geopolitical impact has forced enterprises to reconsider the way they handle and protect the privacy of individuals and the information collected about them and how enterprises implement their cloud-based big data solutions.”

The second concerns the way organisations deliver IT projects. “Most big data projects are technology- and data- intensive. The technology is complicated and the skills required to deliver are relatively scarce, which has resulted in project overruns and budget explosions.”

The third concerns the increased storage in disparate repositories of sensitive, personally identifiable data such as health records and credit card details. “The storage and analysis of such data have increased the pressure on organisations to comply with data and privacy regulations, such as the Payment Card Industry Data Security Standard (PCI DSS), the 1998 UK Data Protection Act and the US Health Insurance Portability and Accountability Act (HIPAA). A pragmatic approach is required to address these big data compliance requirements.”

A number of challenges are identified as resulting from dealing with such large quantities of data, daily and in real-time. These are:

- Complex technological evolution
- Data privacy and integrity
- Security of the data at rest and in motion
- Availability and data system resilience (IT infrastructure)
- Incident response to and management of big data breaches
- Governance, risk and compliance
- Identity and access management
- Skill shortage in this domain

Because of the inherent vulnerability of big data to data and privacy breaches, effective security preventing such breaches is an unavoidable requirement. This is a particular challenge, however, because of a number of complicating factors:

- Increasingly complex IT environment
- Massive growth of transaction data volumes
- Explosion of new types of interaction data, such as social media and device data
- Use of insecure java-based frameworks
- Insider and external threats
- Advanced persistent threat (APT)

It is therefore essential for organisations to develop, implement and monitor robust data privacy solutions that prevent data breaches and enforce data security. Such solutions should permit organisations to identify all sensitive data, ensure that sensitive data are identified and secured, demonstrate compliance

with all applicable laws and regulations, proactively monitor the data and IT environment, and react and respond faster to data or privacy breaches with incident management

Risk related to big data can be categorised as operational or information technology based. Operational risk includes external and internal factors that include geopolitical issues and time and delivery pressures within an organisation that encourage implementation and use before proper controls have been designed, implemented and evaluated. IT risk is a subset of business risk, the business risk related to the “use, ownership, operation, involvement, influence and adoption of IT within an enterprise.” Policies and procedures need to be in place to ensure that proper safeguards are not bypassed, ignored or forgotten as technologies and operating practices change.

The development of big data privacy strategies is pushing towards the establishment of norms for big data privacy. Among suggestions already made are the embedding of privacy into technology some that individuals can retain control of their own information, regardless of where it is stored [59], perhaps by means of metadata attached to data that encode the preferences of individuals in respect of how their data are used [60]. The Software and Information Industry Association is not in favour of potentially excessive legislation in respect of big data and privacy, recommending instead the organisations integrate privacy into their own big data policies. One aspect of this could be to anonymise consumer data as soon as possible after acquisition, removing personal identifiers. The SIIA and similar industry groups are in favour of policy being created through mutual exchange between policy makers, consumer advocates and other stakeholders [61].

The main messages emerging from the ISACA review are, unsurprisingly for an audit organisation, the need for frameworks, standards and consistency. These should apply throughout the data management lifecycle, and be based on the most fundamental principles of IT governance and security management, starting with effective risk management.

3.7 Toxic data

The concept of “toxic” data is alluded to by two sources consulted during this review: ISACA [57] and Verizon [46]. According to ISACA, toxic data are data the loss of which could be damaging to the enterprise. Examples of potentially toxic data could be:

- Private or custodial information such as credit card numbers, personally identifiable information such as Social Security numbers, and personal health information
- Strategic information such as intellectual property, business plans and product designs
- Information such as key performance indicators, sales figures, financial metrics and production metrics used to make critical decisions

It is important to note that both internally-generated data and data acquired from external sources can be considered toxic according to this definition.

According to Verizon, “toxic data... predominantly carries value for people outside the organisation. Examples of toxic data include all data that is regulated, including personally identifiable information. Customer data is naturally part of this category as it often is subject to privacy regulations. Leakage of this

data can result in regulatory fines and public relations issues.” Again, organisations need to take great care to protect and secure such data and prevent any leakage.

Summary

This definition is important because it helps to crystallise the impacts that the loss of certain types of data could have. This can be seen in concrete examples of recent data losses by high-profile organisations such as Adobe, Orange France, and UBS. Depending on the nature and validity of the data lost, the impact on users might be greater or lesser (Orange France, for example, reported not having leaked any client billing information, while UBS saw detailed client information copied and passed to the German tax authorities, with potentially serious consequences for the subjects of the data). The impacts on the data holders should not be forgotten, however. The loss of client confidence resulting from a high-profile data losses or thefts can damage a company’s standing and conceivably put its future in doubt.

3.8 Conclusion

The accumulation and acquisition of large datasets is not a new phenomenon, but over recent years it has taken on greater dimension and greater significance.

The reasons for this are, broadly speaking, threefold.

The first is the rapidly increasing capacity of data storage. Where space was for so long at a premium, both in respect of the traditional physical storage of information and records and then later in respect of electronic storage, rapid improvements in digital technologies have meant that far greater quantities of data and information can be stored in live or easy-to-access repositories for far longer. Of course, these opportunities bring with them an operational and compliance-driven obligation for efficient policies and procedures in respect of archiving, data quality and access management in order to ensure that control objectives are not forgotten and that the scope of control activities is sufficiently broad to include the changes to business operations and the organisation’s risk profile that are created by the expansion in the storage of data.

The second reason is the growing understanding of the volumes, nature, and potential uses of the different forms of unstructured and semi-structured data that have a tendency to accumulate within organisations. Individual organisations have been able to generate value from effective data management processes in the past, but it has only been in the last decade or so that the concepts and significance have become the subjects of widespread discussion and acceptance. Whereas in general terms, based on the author’s experience of auditing a large number of organisations of all sizes and types, many organisations were for many years unaware of the nature, quantities and potential significance of the data they owned that were located outside central data repositories, a number of factors have combined over recent years to change attitudes and drive and move towards managing such data. These factors include the possibility of competitive advantage as a positive driver, the requirements under various compliance regimes to enforce effective data management processes as a negative driver, and the always-present importance of cost management for an organisation. Rather than have quantities of data sitting around unused and unusable and consuming limited resources, it is clearly preferable from a purely financial perspective to manage the data effectively and extract value from them, or to delete the data and either decommission the systems and infrastructure elements or reuse them elsewhere in the organisation.

The third reason is related to the concept of Big Data. Organisations are increasingly becoming aware of the possibilities offered by deep data mining, both of their own data and of data acquired from other sources. In order to do this successfully, however, they need to develop, catalogue, clean and maintain their own datasets as a starting point. Data analysis can only ever be as good as the quality of the source data.

Several publications have already begun to address the impacts of the use of unstructured and big data and to identify and discuss the risks related to these activities. The risks are clear in respect of many key areas of corporate governance. From an overall perspective, once such activities form a part of operational activities, they fall into scope for overall compliance and regulatory questions, except where specifically excluded on the basis of materiality, for example. Then the major questions of privacy and security need to be addressed in detail. Owning, managing and utilizing data generate a significant responsibility on the part of the data owners towards the subjects of the data, who expect and require that their privacy and confidentiality are respected and ensured. This requires both careful consideration, and control over, the way data are used and to whom and in what form the data or summaries thereof are communicated. It also requires that the security of the data, their support systems and their platforms is carefully designed, implemented and managed. Only with a complete, coherent and monitored security framework can data owners and organisational managers hope to ensure that their responsibilities in respect of privacy and security are being met and that they are not being exposed to avoidable risks related to non-compliance.

It is impossible to discuss data storage and management, particularly in respect of big data, without considering the impact of outsourcing and cloud computing on the modern computing landscape. Such facilities increase the capacities for storage of every organisation and increase the possibilities for what they can do with all of the data that they have accumulated. This subject will begin to be addressed in the next chapter in the form of a literature review.

Chapter 4 Analysis of Existing Research II: Cloud Computing

4.1 Introduction

**“Was I deceiv'd, or did a sable cloud
Turn forth her silver lining on the night?”⁵**

The subject of the cloud has generated a large amount of literature over recent years, much of it dedicated to discussing and determining what cloud computing actually is.

As discussed previously, the cloud is a fundamental element in the use and development of information systems. Its role in supporting organisations in their storage, processing and management of large data sources is becoming increasingly significant, but as a new element in the information landscape for many organisations the specific risks it presents, as well as the opportunities it presents, need to be defined, clarified and understood.

In common with the approach adopted in Chapter 3, this review encompasses a range of sources of literature. Academic papers are considered, as are books aimed at IT and controls managers, and white papers published by controls organisations. This is in order to obtain a wide perspective on the issues and realities presented by cloud technologies, avoiding a purely academic (and thereby often naturally theoretical) view of the developing landscape and developing a more pragmatic and practical view. This is also consistent with the overall philosophy of this thesis, which is written very much from a risk management and internal controls perspective.

This chapter is divided into three major sections. Firstly the terms are defined and explained. Then the aspects of risks, governance, opportunities and benefits are outlined. Finally, the key questions of security and privacy are discussed. This structure suggests itself because of the realities of outsourced and cloud service provision: organisations will need to be expecting clear and demonstrable benefits from a move onto hosted solutions, while their decision will always, assuming awareness and competence, take the issues of security and privacy into account as perhaps the major factor after financial considerations.

4.2 Definition of terms

Adolph et al [62] describe “the new paradigm in which computing is offered as a utility by third parties whereby the user is billed only for consumption.” They continue with the definition of new computing paradigms: cloud computing, edge computing, grid computing and utility computing, many of which overlap in terms of the services offered and the nature of the technologies used. It is clear for the

⁵ John Milton, “Comus”, lines 221-222

authors that there are a number of strategic questions that potential users of cloud services need to answer before engaging any services: “The users of distributed systems need to consider legal aspects, questions of liability and data security, before outsourcing data and processes.” The motivations for using cloud services are clearly set out: “For small and medium-sized enterprises, the ability to outsource IT services and applications not only offers the potential to reduce overall costs, but also can lower the barriers to entry for many processing-intensive activities, since it eliminates the need for up-front capital investment and the necessity of maintaining dedicated infrastructure. Cloud providers gain an additional source of revenue and are able to commercialize their expertise in managing large data centres.” A list of utility and cloud providers (Amazon Web Services, Google App Engine, Salesforce.com, Azure Services Platform) is provided and the differences between the services and technologies discussed. The bigger picture in respect of the implications of using cloud services is also discussed: “The continued and successful deployment of computing as a utility presents other challenges, including issues of privacy, security, liability, access and regulation. Distributed computing paradigms operate across borders, and raise jurisdiction and law enforcement issues similarly to those of the Internet itself.” Among the key issues identified are reliability and liability, security, privacy and anonymity, and access and usage restrictions. The authors also cite research published by Gartner in which seven issues potential cloud customers should address are summarised: privileged user access; regulatory compliance; data location; data segregation; data recovery; investigative support; and long-term viability. Concerns are also expressed about portability and interoperability.

The ITU Telecommunication Standardization Bureau repository [63] is a summary of the activities of numerous organisations and consortia. This review distinguishes the differences in technologies, service offerings and attitudes to control that are evident when comparing different service providers, and confirms that “cloud computing will require rapid standardization in many fields such as security, interfaces, quality of service, portability and cloud interoperability.”

Gantz and Reinsel [64] describe “cloud computing solutions – both public and private and a combination of the two known as hybrid – provide enterprises with new levels of economies of scale, agility and flexibility compared with traditional IT environments. In the long term, this will be a key tool for dealing with the complexity of the digital universe. Cloud computing is enabling the consumption of IT as a service. Couple that with the “big data” phenomenon, and organisations increasingly will be motivated to consume IT as an external service versus internal infrastructure investments.” They comment that “while cloud computing accounts for less than 2% of IT spending today, IDC estimates that by 2015 nearly 20% of the information will be “touched” by cloud computing service providers – meaning that somewhere in a byte’s journey from originator to disposal it will be stored or processed in a cloud. Perhaps as much as 10% will be maintained in a cloud.” Issues and challenges are clearly set out: “The challenges for cloud adoption include: data preparation for conversion to cloud; integrated cloud/non-cloud management; service-level agreements and termination strategies; security, backup, archiving and disaster control strategies; intercountry data transfer and compliance; organisational politics. The latter is non-trivial. Many of the most successful virtualization projects succeed, in part, because CIOs have developed opt-in and opt-out strategies for internal departments that may be reluctant to share responsibility for what was once considered “their” information or data. Converting to cloud computing means changing the status quo – always a difficult task, even if there are good reasons to do so.”

The authors raise a number of questions in respect of the impact of the cloud on IT management practices. These include trust issues, and the need to classify information, providing an example of a possible

classification: “Privacy only – such as an email address on a YouTube upload; Compliance driven – such as emails that might be discoverable in litigation or subject to retention rules; Custodial – account information, a breach of which could lead to or aid in identity theft; Confidential – information the originator wants to protect, such as trade secrets, customer lists, confidential memos, etc; Lockdown – information requiring the highest security, such as financial transactions, personnel files, medical records, military intelligence, etc.” They emphasise a very significant point in respect of data security: “In 2010, 28% of the digital universe required some level of security. Note that this is information that *needs* security. It may not have it.”

Bertino et al [65] discuss identity issues and their fundamental importance to cloud services, where more traditional methods of verifying identity and retaining assurance that identity management processes are reliable are no longer as straightforward to apply. “Users have typically to establish their identity each time they use a new cloud service, usually by filling out an online form and providing sensitive personal information... therefore, the development of digital identity management (IdM for short) systems suitable for cloud computing is crucial. An important requirement is that users of cloud services must have control on which personal information is disclosed and how this information is used in order to minimise the risk of identity theft and fraud.”

In a case study written from the perspective of a service provider, Cooper et al [66] describe the set of requirements laid out for the development of the Yahoo! Cloud: multitenancy; elasticity; scalability; load and tenant balancing; availability; security; operability; metering; global and simple APIs. After discussing the challenges encountered by the project team during the initial phases of the project, the authors present a number of open questions faced by the Yahoo! Team: interacting with the cloud (user perspective); quality of service; automating operations; growth; privacy; capacity management. It is clear from this that the service providers themselves, far from presenting mature, stable and complete solutions for client needs, are developing solutions and addressing problems as they arise and developing management and control structures on the basis of experience.

Mather et al in their book “Cloud Security and Privacy” [67] begin by giving their definition of cloud computing: multitenancy (shared resources), massive scalability, elasticity, pay as you go, and self-provisioning of resources.

Multitenancy (shared resources) is contrasted with previous computing models. Instead of assuming dedicated resources, “cloud computing is based on a business model in which resources are shared (i.e., multiple users use the same resource) at the network level, host level, and application level.”

Massive scalability is the second key distinguishing factor: “Although organisations might have hundreds or thousands of systems, cloud computing provides the ability to scale to tens of thousands of systems, as well as the ability to massively scale bandwidth and storage space.”

The Elasticity of supply is noted: “Users can rapidly increase and decrease their computing resources as needed, as well as release resources for other uses when they are no longer required.”

The fourth element is defined as Pay as you go: “Users pay for only the resources they actually use and for only the time they require them.”

The final distinguishing element is given as the self-provisioning of resources: “Users self-provision resources, such as additional systems (processing capability, software, storage) and network resources” (p. 7-8).

The authors contrast the cloud with the traditional and well-established IaaS (Infrastructure as a Service) model of providing services to clients. IaaS is described: “In the traditional hosted application model, the vendor provides the entire infrastructure for a customer to run his applications. Often, this entails housing dedicated hardware that is purchased or leased for that specific application. The IaaS model also provides the infrastructure to run the applications, but the cloud computing approach makes it possible to offer a pay-per-use model and to scale the service depending on demand. From the IaaS provider’s perspective, it can build an infrastructure that handles the peaks and troughs of its customers’ demands and add new capacity as the overall demand increases. Similarly, in a hosted application model, the IaaS vendor can cover application hosting only, or can extend to other services (such as application support, application development, and enhancements) and can support the more comprehensive outsourcing of IT.”

It is important to compare and contrast this model with the newer utility and cloud paradigms, for there is overlap conceptually. The key distinction is made over the level of abstraction in the services and infrastructure paid for: “The IaaS model is similar to utility computing, in which the basic idea is to offer computing services in the same way as utilities. That is, you pay for the amount of processing power, disk space, and so on that you actually consume. IaaS is typically a service associated with cloud computing and refers to online services that abstract the user from the details of infrastructure, including physical computing resources, location, data partitioning, scaling, security, backup, and so on. In cloud computing, the provider is in complete control of the infrastructure. Utility computing users, conversely, seek a service that allows them to deploy, manage, and scale online services using the provider’s resources and pay for resources the customer consumes. However, the customer wants to be in control of the geographic location of the infrastructure and what runs on each server.” When considering what kind of solution to opt for, organisations should have a clear idea of exactly how distant they want to be from both strategic and operational decisions, knowing that there is a trade-off to be made between control and convenience and between retaining expertise and delegating technical know-how to providers.

The authors describe the features available within a typical IaaS system: “Scalability - The ability to scale infrastructure requirements, such as computing resources, memory, and storage (in near-real-time speeds) based on usage requirements; ... Pay as you go - The ability to purchase the exact amount of infrastructure required at any specific time; ... Best-of-breed technology and resources - Access to best-of-breed technology solutions and superior IT talent for a fraction of the cost” (p. 22) and then clarify the distinctions between types of clouds.

“Public clouds (or external clouds) describe cloud computing in the traditional mainstream sense, whereby resources are dynamically provisioned on a fine-grained, self-service basis over the Internet, via web applications or web services, from an off-site, third-party provider who shares resources and bills on a fine-grained, utility-computing basis. A public cloud is hosted, operated, and managed by a third-party vendor from one or more data centers. The service is offered to multiple customers (the cloud is offered to multiple tenants) over a common infrastructure. In a public cloud, security management and day-to-day operations are relegated to the third-party vendor, who is responsible for the public cloud service offering. Hence, the customer of the public cloud service offering has a low degree of control and oversight of the physical and logical security aspects of a private cloud.”

“Private clouds and internal clouds are terms used to describe offerings that emulate cloud computing on private networks. These (typically virtualization automation) products claim to deliver some benefits of cloud computing without the pitfalls, capitalizing on data security, corporate governance, and reliability concerns. Organisations must buy, build, and manage them and, as such, do not benefit from lower upfront capital costs and less hands-on management. The organisational customer for a private cloud is responsible for the operation of his private cloud. Private clouds differ from public clouds in that the network, computing, and storage infrastructure associated with private clouds is dedicated to a single organisation and is not shared with any other organisations (i.e., the cloud is dedicated to a single organisational tenant). As such, a variety of private cloud patterns have emerged:

- Dedicated - Private clouds hosted within a customer-owned data center or at a collocation facility, and operated by internal IT departments
- Community - Private clouds located at the premises of a third party; owned, managed, and operated by a vendor who is bound by custom SLAs and contractual clauses with security and compliance requirements
- Managed - Private cloud infrastructure owned by a customer and managed by a vendor

In general, in a private cloud operating model, the security management and day-to-day operation of hosts are relegated to internal IT or to a third party with contractual SLAs. By virtue of this direct governance model, a customer of a private cloud should have a high degree of control and oversight of the physical and logical security aspects of the private cloud infrastructure—both the hypervisor and the hosted virtualized OSs. With that high degree of control and transparency, it is easier for a customer to comply with established corporate security standards, policies, and regulatory compliance.”

“A hybrid cloud environment consisting of multiple internal and/or external providers is a possible deployment for organisations. With a hybrid cloud, organisations might run non-core applications in a public cloud, while maintaining core applications and sensitive data in-house in a private cloud” (p. 23-25). These distinctions and definitions are important because the cloud is not a single, coherent, cohesive entity. Paying for cloud services is not equivalent to purchasing a specific tangible asset or buying a piece of shrink-wrapped software.

The authors then outline certain key drivers for adopting the cloud, including lower IT costs, faster time to go live, and reduced complexity. They do note that “However, with cloud computing it is critical to understand how to integrate the cloud solution into existing enterprise architecture” (p26). Unless the entirety of the IT infrastructure and service provision is to be moved outside the organisation, integration will be a key consideration.

The discussion then turns to barriers to adopting the cloud, which they list as security, privacy, connectivity and open access, reliability, interoperability, independence from CSPs (Cloud Service Providers), economic value, IT governance, changes in the IT organisation, and political issues due to global boundaries (p31-34). These issues, which are critical to the success of any project to externalise part or all of the IT function, will be discussed in detail in Chapter 6.

Summary

The basic definition of the cloud is very straightforward but technologies and service offerings have rapidly developed and differentiated themselves, so the potential client has a number of options to consider. The arguments for adopting cloud technologies are prima facie very attractive, with organisations being able to buy storage and processing as a service, paid for on a usage basis, but already at an early stage in the decision process a number of questions are presented in respect of the management and security of the data and applications hosted and operated by third parties.

4.3 Risks and governance, opportunities and benefits

In another 2013 white paper [68] ISACA discusses the challenges facing board members in respect of cloud services.

“Board members are hearing more and more from their management teams about the noteworthy business benefits of cloud computing, such as:

- Cloud strategies make the enterprise more efficient and agile.
- Cloud computing allows delivered services to be more innovative and more competitive.
- Cloud computing reduces overall operating costs.

But how confident can boards be that management plans will achieve these benefits? Is there a way to know that, even if the benefits are real, increased operational risk will not outweigh those benefits?”

The paper highlights a significant difference between the cloud paradigm and traditional models of providing IT services. “The core goal of cloud computing is to turn enterprise computing into a fungible commodity. In the traditional model of IT, the enterprise acquires and maintains a portfolio of slowly depreciating technology assets that may or may not be employed efficiently. Conversely, the goal of cloud is to enable the enterprise to manage computing much like electricity, buying only what it uses (no more and no less).”

A number of benefits to the overall business are identified:

“By looking at cloud computing in this way, the board of directors can start to envision the possible benefits, for example:

- Shifting the funding of IT from large capital investments (legacy IT assets) to operational expenses.
- Reallocating IT resources to core business activities.
- Procuring applications that are easier and cheaper to implement, use and support.
- Increasing scalability and flexibility, enhancing the ability to respond to changing market conditions.
- Fostering innovation by shifting effort and resources from implementation projects to final product development.”

In terms of the technologies used in the cloud, many of the components that make up the cloud computing model are not new. Key elements such as high-speed Internet access, server virtualization, new software development approaches, and advances in high-capacity storage have existed for several years. What is new, however, is the impact outside the IT landscape: it brings with it cultural changes and specific changes in domains such as customer service, change management, and vendor management. Management needs to ensure that the cloud is incorporated into activities in such a way that it enhances and enables rather than disrupts negatively.

Governance needs to be reconsidered from first principles. "Like any investment, cloud must be governed; however, some cloud computing characteristics (virtualization, agility, flexibility, faster deployment and minimal initial investment) may require additional governance considerations to ensure that benefits are realized within acceptable levels of risk." With cloud services frequently provided on a metered basis, organisations can exploit the flexibility offered by the cloud infrastructure to quickly and easily increase or decrease the service level. Such changes, however, can easily bypass the standard procurement processes that form a key part of proper governance. A Forrester report [69] estimates that for every cloud project that is tracked centrally and formally by IT functions, there are between three and six that escape proper tracking. The downside to such flexibility, concludes ISACA, is that activities "could also result in the bypassing of expense authorization, change control processes, information protection controls and other oversight processes."

The potential consequences of this are straightforward. "Bypassing established governance processes and failing to inform others within the enterprise about cloud computing initiatives may result in the enterprise assuming unknown risk and, thereby, increasing potential exposure."

ISACA conclude by identifying five questions to which management should have firm and reliable answers if they are to be certain that stakeholder needs have been considered and that risk and resource utilisation are being optimised. These are:

- i. Do management teams have a plan for cloud computing, having weighed value and opportunity costs?

This involves setting the risks involved in any new venture against the potential benefits to be drawn from leveraging new technologies and opportunities, including improving products and services, increasing productivity, and containing costs.

- ii. How do current cloud plans support the enterprise's mission?

This is fundamental: there should be a clear and traceable link to the enterprise strategy so that the added value expected from adopting the cloud can be defined, accepted and subsequently measured.

- iii. Have executive teams systematically evaluated organisational readiness?

It is essential to minimise pressure points arising from such factors as:

- Cloud computing implementations conflicting with enterprise culture
- Skills necessary to support cloud solutions not being available
- Cloud-related processes conflicting with established processes

- Organisational structures not maximising cloud effectiveness or efficiency.
- iv. Have management teams considered what existing investments might be lost in their cloud planning?

The adoption of the cloud could lead to necessary changes in current infrastructure and resources, forcing a premature end to the use or development of systems, the end of the use of a certain technology, and the need to retrain or replace staff whose skills and experience will no longer be relevant in the long-term.

- v. Do management teams have strategies to measure and track the value of cloud return vs. risk?

It is essential that proper reporting mechanisms are designed and implemented so that value and risk aligned with enterprise goals can be measured.

Summary

This analysis can usefully be summarised in the form of a traditional SWOT (Strengths, Weaknesses, Opportunities, Threats) analysis. This helps to compare and contrast the possibilities offered by cloud services with the consequences that they bring.

Strengths	Weaknesses
<ul style="list-style-type: none"> • Unlocking funding from capital investments to operational expenses • Increasing scalability and flexibility • Access to newer and more efficient technologies • Access to technical specialists in the fields of operations and security 	<ul style="list-style-type: none"> • Ongoing cashflow commitment • Loss of hands-on control of security and operational matters • Loss of visibility over change management and security management processes
Opportunities	Threats
<ul style="list-style-type: none"> • Reallocation of resources from IT to more obviously revenue-generating activities • Ability to respond more rapidly through services on demand • Greater impetus for innovation 	<ul style="list-style-type: none"> • Potential lock-in to a single supplier or service provider • Potential clashes with internal culture or control environment • Perceived or real loss of control over systems, data and activities, with impact on the ability to attest to compliance.

Table 4.1 SWOT analysis in respect of the cloud

This analysis demonstrates that the decision to migrate onto cloud services is not one to be taken lightly by management. The short- and medium-term advantages of such a move are certainly attractive, especially from a financial and strategic perspective. The longer-term consequences require careful consideration, however, particularly in respect of exerting control over activities and being able to ensure compliance with regulations and legal requirements. This specific topic is discussed in detail in Chapter 8.

The five questions posed by ISACA provide an efficient way of determining whether the wider consequences of a move onto the cloud have been fully considered by management. In such projects it is easy to get side-tracked by either the financial or the technical aspects, without viewing the impacts as a whole. It is also essential to determine how the success of the project, including quality of performance and returns on investment, is to be measured, and indeed to perform this measurement on an appropriately regular basis.

4.4 Security considerations

Security considerations are a fundamental element of any decision to adopt cloud services and are the subject of a lot of detailed analysis in the technical and professional press. The reason for such attention is straightforward: as soon as data or systems are hosted or operated outside an organisation's security perimeter, the amount of direct assurance management can obtain over the security of those data or systems is immediately reduced because of the decrease in direct control and visibility. This topic is discussed in detail in Chapter 6.

In 2012 ISACA published a lengthy white paper on Security Considerations for Cloud Computing [70], containing definitions of the characteristics of the cloud, details of the security risks and threats related to operating in the cloud, and checklists for decision making in respect of the adoption of cloud services.

The paper works from the NIST definition of cloud computing as "a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction" [71]. Five essential characteristics of cloud computing are identified:

- On-demand self-service—Computing capabilities can be provisioned without human interaction from the service provider
- Broad network access—Computing capabilities are available over the network and can be accessed by diverse client platforms
- Resource pooling—Computer resources are pooled to support a multitenant model
- Rapid elasticity—Resources can scale up or down rapidly and in some cases automatically in response to business demands
- Measured service—Resource utilization can be optimized by leveraging charge-per-use capabilities.

The paper argues that security and data privacy concerns are often seen as critical barriers to the adoption of cloud services and that although service level agreements and contracts are essential, the fundamental element in service provision is trust. "There can never be sufficient controls and agreements to mitigate all concerns if trust is a missing factor in the client-supplier relationship."

This is a critical remark in the context of an approach to control and assurance that is based upon a framework of risk management leading to robust internal controls. In the same way as the kind of documented and verified internal control framework over financial reporting as required under Sarbanes-Oxley can be undermined by deliberate dishonesty and collaboration on the part of a few highly placed

individuals, internal controls over the security and privacy of data can be undermined if there is no fundamental trust between service providers and their clients. This needs to be borne in mind when considering the merits of any proposed solution that relies upon robust internal controls (see Chapter 10).

User management should constantly be asking themselves how they can rely on a cloud service provider to protect their data. The answer to this will depend on a number of factors:

- The possibility of auditing and verifying internal controls
- The financial position and market recognition of the cloud service provider
- The existence of certification or other recognition of the service provider by standards authorities
- The availability and credibility of disaster recovery plans (DRPs), business continuity plans (BCPs) and reliable backup procedures across all locations where data will be stored
- The quality and coherence of the users’ own data management and classification frameworks, policies and procedures, applying to both before any data are transferred to the service provider and once it is being stored and processed there
- The overall relationship with the service provider, including contracts and service level agreements, communications processes, and the clarity of roles and responsibilities.

4.4.1 Visibility of third-party activities

The lack of visibility of what happens within the service provider’s organisation is highlighted as a critical risk factor. The people, the processes and the technology supporting the service are largely invisible to the data owners; this lack of visibility, also known as “abstraction”, is aggravated as increasing layers of service are provided. Thus IaaS (infrastructure only) provides the lowest abstraction level, PaaS a greater level of abstraction, and SaaS the greatest. ISACA illustrate this in a convenient graphic.

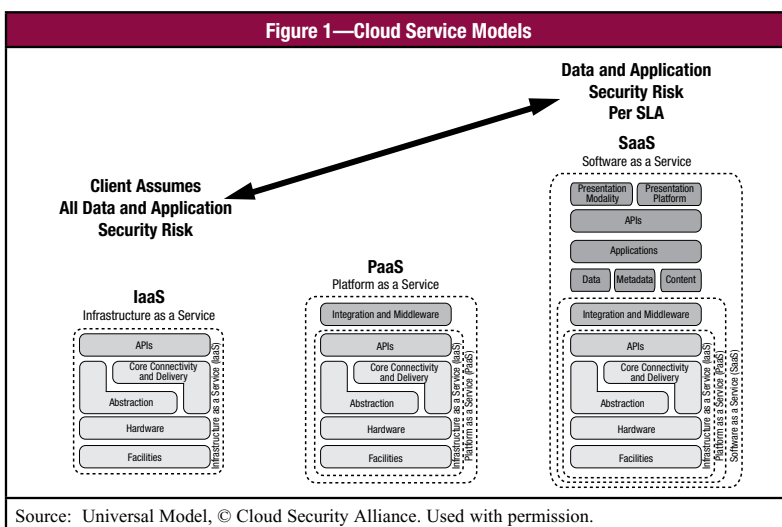


Figure 4.1 Cloud Service Models, showing levels of abstraction, from ISACA p. 14

When considering moving information assets – data, applications and processes – to the cloud, it is necessary to identify clearly which information assets are to be transferred and what kinds of risk events might affect them. Four major types of risk events are identified:

- Unavailability – the asset is unavailable and cannot be used by the organisation, because of accidental failure, intentional attack, or legal restriction
- Loss – the asset is lost or destroyed, accidentally or intentionally
- Theft – the asset is intentionally stolen and is in the possession of a third party. This covers both copying and the physical theft of media leading to data loss
- Disclosure – the asset is released to unauthorised parties or the public, deliberately or by accident.

It is important to note that data might be targeted or affected directly, or affected as a result of an attack on, or problem with, an application system or element of infrastructure. For example, an effective DDoS attack on the service provider might cause a loss of availability to critical data for a period without the attackers knowing, or indeed caring, that those particular data were being hosted by that service provider. Accordingly protective measures should be designed with the whole range of possible scenarios in mind.

Management should also be aware of the indirect effects of costs and financial issues on the security and availability of assets. Although not generally considered to be a direct risk to information assets, financial considerations can be important, for example in the case of an organisation that because of cashflow issues or inefficient internal procedures neglects to pay its provider for services rendered. The provider could, perhaps, respond by blocking access to data until the situation is rectified. The data are therefore unavailable and there is no technical means of preventing or rectifying this.

The key concern for many users when considering migrating onto cloud services is privacy. This concern is typically expressed when an information asset consists of sensitive or personal data. The issue is, however, wider than this, and concerns the concepts of the privacy of data *within* the information asset and the privacy of data *outside* the information asset.

A useful illustrative example of a scenario in which both concepts of data privacy are illustrated is presented in the ISACA paper. In this example, an enterprise that has migrated to a cloud service provider possesses a database of customers and sends emails to these customers to advertise new products. Both the database and the email content are considered sensitive information assets that must be kept private, and have appropriate measures (such as encryption, e-signatures, data access management) to protect them. However, the cloud service provider or an intruder can use the network logs to trace the destination of the emails and can therefore rebuild the database, thus compromising asset privacy.

In respect of the *first type of privacy* (the privacy of data within information assets), the primary concern is to ensure that the information asset is not disclosed unnecessarily within the normal operational processing framework. Such assets should be identified through proper data classification prior to migration and then subjected to appropriate measures to protect against disclosure.

The *second type of privacy* (privacy of data outside information assets) is more complex and more difficult to ensure because it involves the collection, retention and processing of data that are not part of the information assets of the enterprise. Service providers often collect such data for benign and entirely

legitimate purposes (such as troubleshooting and incident analysis) or for legal reasons (such as data retention policies). It can thus often be very challenging to prevent disclosure or theft.

It should be pointed out that this specific problem is not unique to the cloud environment; it applied to all situations where an organisation is not in complete control of the infrastructure on which its data and applications are hosted.

This illustration is useful because it demonstrates how even everyday activities can lead to risks of the disclosure of data and breaches of privacy without the data owners and operators necessarily being aware of those risks.

Mather et al [67] also address key privacy concerns in the cloud. This is a large subject that frequently sees an overlap between security and privacy. A number of considerations are raised, including access, compliance, storage, retention, destruction, audit and monitoring, and privacy breaches. It is noted at the end of the section that “Many of these concerns are not specific to personal information, but to all types of information and a broader set of compliance requirements” (p. 149-150).

The next discussion concerns the need for changes to privacy risk management and compliance in relation to cloud computing. Mention is made of the Collection Limitation Principle, which “specifies that collection of personal data should be limited to the minimum amount of data required for the purpose for which it is collected. Any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject.” It is noted that “in the privacy arena, lack of specifics on data collection with providers creates misunderstandings down the road. For instance, one global outsourcer said, “Clients come in expecting the right things in security, but the wrong things in privacy. They are expecting best practices, but they don’t know what they are.” There are comprehensive security frameworks and standards (such as the ISO 27000 series, NIST guidelines, etc.), and organisations know how to implement them. There is no universally adopted privacy standard—instead, there are conflicting laws, regulations, and views on what privacy is and what it requires from organisations to protect it. Many organisations want to do what they perceive to be “the right thing”; however, their perception may be different from the law. As a result, there may be different expectations regarding what privacy means between the organisation and the CSP, and no agreed best practices” (p.151). This summarises a significant issue for many organisations: knowing what data they hold and what they are allowed to do in respect of it. In a context where compliance is becoming ever more stringent and complicated, privacy is becoming a minefield.

Further principles that are discussed by Mather et al include the Use Limitation Principle, the Security Principle, the Retention and Destruction Principle, the Transfer Principle, and the Accountability Principle. These principles set out how organisations should behave and what points should concern them when managing data security, and all require a great deal of analysis and awareness and well-defined and robust policies and procedures for addressing the issues raised.

In line with other initiatives that will have significant impacts on the IT infrastructure of an organisation, it is essential to perform an appropriate risk assessment when migrating to the cloud. Only through detailed risk assessment can relevant risks be identified and appropriate measures be designed. Such an assessment needs to be tailored to the precise circumstances of the organisation, the nature of cloud services being considered, and the nature of the data being considered for transfer. In respect of IaaS and the storage of data, the following risk factors are identified:

Risk factor	Description
Scalability and elasticity	Risk is reduced because cloud technologies are by nature scalable and flexible
DRP and backup	Risk is reduced because service providers should have, as a matter of good practice, effective disaster recovery and backup procedures. Users do need to ensure, however, that their own RPO and RTO requirements can be met
Patch management	Risk is reduced because patches and other updates are typically handled centrally by a dedicated team, responding rapidly to requirements to apply changes
Cross-border legal requirements	Risk is increased because cloud service providers are often transborder entities and with different countries having different requirements for compliance, there is often the potential for violation of regulations when storing, processing or transmitting data across the service provider's infrastructure
Multi-tenancy and isolation failure	Risk is increased because of shared infrastructure, especially if good maintenance policies for securely deleting data, for example, are not followed
Lack of visibility over security measures	Risk is increased because the users cannot see in real-time either the detailed nature of the internal controls in place or the effectiveness of those controls
Physical security	Risk is increased because physical resources are shared and individuals may have physical access to infrastructure elements – and thus information assets – who have no business or operational requirement to do so
Data disposal	Risk is increased because the data owners do not control the process for disposing of data, nor the processing for replacing, recycling or upgrading infrastructure elements such as disks
Offshoring infrastructure	Risk is increased because the attack surface area is expanded. With more points to attack, the chances of an attacker finding a weak spot are increased

Table 4.2 Cloud storage risk factors

It should be noted that these are specific to the IaaS offering and that management should also take into consideration common and more general risk factors, such as external hacking, mobile computing vulnerabilities, and business impact due to provider inability. It is also important to note that not all the risk factors identified are actually increased by a move onto cloud services. The scope of the risk assessment exercise covers all relevant factors, regardless of whether their impact on risk levels can subsequently be determined to be positive, negative, or indeed neutral.

The risk factors specific to each deployment model are also outlined and can be summarised in the table below.

Deployment Model	Risk-decreasing factor	Risk-increasing factor
Public cloud	<i>Public reputation</i> – service providers are aware of the need to maintain a good reputation by providing high-quality,	<i>Full sharing of the cloud (data pooling)</i> – tenants share infrastructure but do not necessarily have common interests and

Deployment Model	Risk-decreasing factor	Risk-increasing factor
	secure services.	concerns for security. <i>Collateral damage</i> – if one tenant, or the provider itself, is attacked, there could be impacts on other tenants.
Community cloud	<i>Same group of entities</i> – the “trust” component lowers the risk relative to a public cloud.	<i>Sharing of the cloud</i> – different entities may have different security requirements or measures in place.
	<i>Dedicated community access</i> – access can be configured and restricted to authorised community users only	
Private cloud	<i>Can be built on-site</i> – allowing greater control over physical security and location issues.	<i>Application compatibility</i> – older or customised software may not run well in virtual environments without direct access to resources.
	<i>Performance</i> – being inside the security perimeter and on a local network, transfer rates are increased as there are fewer nodes to cross.	<i>Investments required</i> – creating shared infrastructure requires investment in material and human resources, which runs counter to the common cost-cutting motivation for moving onto the cloud.
		<i>IT skills required</i> – specific cloud skills are required alongside traditional IT expertise.
Hybrid cloud	See above.	See above, also: <i>cloud interdependency</i> – strict identity controls and strong credentials will be needed to allow one cloud to have access to another.

Table 4.3 Risk factors specific to cloud deployment models

For each of the risk factors outlined above, management should identify the relevant threats and vulnerabilities and define appropriate mitigating actions.

Krutz and Vines in their book “Cloud Security” [72] define the concerns involved in cloud computing. These are numerous: “Leakage and unauthorized access of data among virtual machines running on the same server; Failure of a cloud provider to properly handle and protect sensitive information; Release of critical and sensitive data to law enforcement or government agencies without the approval and/or knowledge of the client; Ability to meet compliance and regulatory requirements; System crashes and failures that make the cloud service unavailable for extended periods of time; Hackers breaking into client applications hosted on the cloud and acquiring and distributing sensitive information; The robustness of the security protections instituted by the cloud provider; The degree of interoperability available so that a client can easily move applications among different cloud providers and avoid “lock-in”” (p xxiii-xiv).

None of these concerns are entirely new to IT management, particularly if various types of outsourcing have already been used within an organisation. What is significant, however, is the extent to which the visibility of internal controls and management procedures is limited, or even impossible, and the extent to which cloud customers are reliant upon the service providers to inform them of any events of incidents.

The authors refer to the central principles of data security and two of the three elements that make up the information security triad. "Cloud users should also be concerned about the continued availability of their data over long periods of time and whether or not a cloud provider might surreptitiously exploit sensitive data for its own gain." (p xxiv). Data availability is a pre-requisite for many modern businesses, and the consequences of a loss of systems or data for any but the briefest period can be catastrophic. Well-managed organisations will have determined the maximum time that they can function without their key information systems, and the amount of data that they can afford to lose, and developed contingency plans based around those requirements. Data confidentiality is also critical, both from a competitive advantage perspective and in respect of compliance. Sensitive data need to be stored securely and protected from inappropriate access.

Some historical perspective is provided in order to demonstrate that many of the concepts are not new. "Cloud computing evokes different perceptions in different people. To some, it refers to accessing software and storing data in the "cloud" representation of the Internet or a network and using associated services. To others, it is seen as nothing new, but just a modernization of the time-sharing model that was widely employed in the 1960s before the advent of relatively lower-cost computing platforms. These developments eventually evolved to the client/server model and to the personal computer, which placed large amounts of computing power at people's desktops and spelled the demise of time-sharing systems." (p 1). There is a sense that trends in information systems tend to go in circles, with periodic shifts from centralized systems to local (desktop) processing back to centralized processing, and then to very small and portable devices in permanent contact with large centralized systems again.

Key definitions are provided, referring to familiar and reliable sources. "In an October, 2009 presentation titled "Effectively and Securely Using the Cloud Computing Paradigm," by Peter Mell and Tim Grance of the National Institute of Standards and Technology (NIST) Information Technology Laboratory, cloud computing is defined as follows: "Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable and reliable computing resources (e.g., networks, servers, storage, applications, services) that can be rapidly provisioned and released with minimal consumer management effort or service provider interaction." This cloud model is composed of five essential characteristics, three service models, and four deployment models. The five essential characteristics are as follows: On-demand self-service; Ubiquitous network access; Resource pooling; Location independence; Rapid elasticity; Measured service" (p2).

4.4.2 Open standards and consistency

The key questions of open standards and supplier consistency need to be considered because of the importance of cloud users not being tied into one unique provider because of proprietary platforms or methodologies. "In 2009, the Open Cloud Manifesto was developed by a group of organisations including IBM, Intel, and Google to propose practices for use in the provision of cloud computing services. In the "Open Cloud Manifesto" (www.opencloudmanifesto.org), cloud computing is defined with a set of characteristics and value propositions.

The characteristics outlined in the manifesto are as follows:

- The ability to scale and provision computing power dynamically in a cost-efficient way;
- The ability of the consumer (end user, organisation, or IT staff) to make the most of that power without having to manage the underlying complexity of the technology;
- The cloud architecture itself can be private (hosted within an organisation’s firewall) or public (hosted on the Internet)” (p3).

This manifesto thus provides a wide definition of applicable cloud services and a broad scope for activities to fit under the cloud umbrella while conforming to the basic principles of openness and consistency.

The manifesto also contains a number of value propositions, as follows:

Value proposition	Description
Scalability on demand	All organisations have to deal with changes in their environment. The ability of cloud computing solutions to scale up and down is a major benefit. If an organisation has periods of time during which their computing resource needs are much higher or lower than normal, cloud technologies (both private and public) can potentially deal with those changes more quickly and effectively than an organisation running its IT in-house
Streamlining the data centre	An organisation of any size will have a substantial investment in its data centre. That includes buying and maintaining the hardware and software, providing the facilities in which the hardware is housed, and hiring the personnel who keep the data centre running. An organisation can make its data centre more responsive and efficient by taking advantage of cloud technologies internally or by offloading workload onto a service provider
Improving business processes	The cloud provides an infrastructure for improving business processes. An organisation and its suppliers and partners can share data and applications in the cloud, enabling everyone involved to focus on the business process instead of the infrastructure that hosts it
Minimizing startup costs	For new startups, organisations in emerging markets, or even advanced technology groups in larger organisations, cloud computing can greatly reduce initial costs. The new organisation starts with a ready-made and available infrastructure

Table 4.4 Value propositions arising from the Open Cloud Manifesto

Some prominent figures in the information systems and security field do hold dissenting views in respect of the cloud. These include a video blog published by Oracle CEO Larry Ellison in which he bluntly dismisses the whole idea: “What the hell is cloud computing? ... When I read these articles on cloud computing, it is

pure idiocy. ... Some say it is a using a computer that is out there. ... The people that are writing this are insane. . . . When is this idiocy going to stop?” while the prominent information security expert Bruce Schneier has also expressed reservations, as in his online newsletter on 4 June 2009 [73]: “This year’s overhyped IT concept is cloud computing.... But, hype aside, cloud computing is nothing new. It’s the modern version of the timesharing model from the 1960s, which was eventually killed by the rise of the personal computer. It’s what Hotmail and Gmail have been doing all these years, and it’s social networking sites, remote backup companies, and remote email filtering companies such as MessageLabs. Any IT outsourcing - network infrastructure, security monitoring, remote hosting—is a form of cloud computing.”

There is much to be said for such analyses when placed into the correct context. The history of computing and information systems in cyclical and many advances are the restatement or the further development of principles or services that have long been established, often applying new technologies to improve or transform existing services. Cloud computing for both organisations and individuals could not be feasible until reliable access to high-speed broadband networks was possible, especially over wifi, but the core principles of webmail services, for example, predated fast consumer networks by many years.

The realistic impacts of cloud computing should be considered, with the understanding that it will not cause a revolutionary change in the way all information services are provided to organisations. In a 2009 article in Information Week magazine [74] Russ Daniels of Hewlett Packard stated: “Virtually every enterprise will operate in hybrid mode,” with some of its operations on the premises and some in the cloud. He did not see cloud computing as being in any way a replacement for the data center. “The idea that we’re going to one day throw a switch and move everything out to one of a small number of external data centers, located next to a low-cost power source, is nonsensical. It’s not going to happen. Cloud computing is not the end of IT.” This analysis is conceptually very reasonable, except in the cases of small or medium-sized businesses who historically would have outsourced a very great deal, if not all, of their information systems to specialist providers in order to avoid the complications of having in-house IT resources managing in-house systems.

The importance of technological influences is discussed, as these are the key drivers behind the growth of cloud services alongside purely financial considerations. The most significant influences are as follows:

- Universal connectivity: it is no longer necessary to design and implement expensive dedicated networks on specific sites when high-debit broadband access is a global commodity and technologies such as virtual private networks mean that secure and efficient connectivity is possible from any location;
- Commoditization: rather than having IT services within an organisation as cost centre with complex costing arrangements, cloud services allow each aspect of the service offering to be available and priced as a separate commodity. Disk space, processor time, network resources and even support can be ordered as required and billed independently;
- Excess capacity: organisations can develop their activities without needing to plan internally for excess capacity that may or may not be required in the short- and medium-terms;
- Open-source software: the move towards open-source software and open standards has brought with it the portability of environments and a partial end to being limited to using particular combinations of infrastructure elements;

- Virtualization: rapid advances in virtualization have freed organisations from being tied both to specific hardware and operating system combinations to run their applications and to having separate machines for each environment or each instance of an environment.

These technical considerations are tempting to many organisations, especially given the increasing difficulty of recruiting skilled and experienced resources to manage and operate complex and proprietary systems and the desire to move away from limiting proprietary hardware and operating systems. If there is a legitimate business case to replace systems and infrastructure with new technologies and platforms, why not take one step further and allow a specialist service provider to manage operational aspects of information processing?

4.4.3 The benefits of outsourcing

Other key elements that are discussed in the context of the motivation to use cloud services are the benefits of outsourcing. These are defined as:

- Resiliency: dedicated sites that host multiple platforms and have specialist teams in place to monitor performance and capacity issues are likely to provide an acceptable level of system resilience;
- Disaster recovery: as part of the outsourcing contract, disaster recovery planning will (or should) be addressed. The service provider should have the experience and resources to be able to meet contingency planning requirements;
- Economy of scale: outsourcing providers who operate large infrastructures and essentially rent out part of their resources to multiple clients should be able to provide the services more cheaply than an organisation trying to meet its own information systems requirements and not using all of its infrastructure at maximum efficiency;
- Security: once security requirements have been defined and access levels designed, the outsourcing provider should be able to manage logical access centrally and efficiently. In respect of physical security, outsourced systems move many of the risks from the customer to the service provider, which should be able to manage them by implementing good practices around its data centres;
- Automatic and continuous hardware and software upgrades: service providers have the resources to identify necessary upgrades, updates and patches and to test them before implementation;
- Libraries of applications: outsourcing can provide access to up-to-date, well maintained and thoroughly tested libraries and applications and layers of middleware;
- Performance monitoring and management: good outsourcing providers will have demonstrable experience and expertise in systems monitoring and management, and as part of a package of services this may be cheaper than retaining the same service in-house, even if the skills are available in the marketplace;
- Rapid capability expansion: if an appropriate contract can be agreed, extra resources can be provided and put into service more rapidly through an outsourcing provider that has excess

capacity ready for such client demands than through in-house IT that might require significant lead-times for both technical installations and for the internal ordering process to be completed and necessary financial approvals to be obtained;

- Rapid capability downsizing: in the same way, moving towards lower capacity can more easily be effected by leveraging the flexibility of a specialist provider than by remodeling internal provisions and needing to address subsequent issues of redundant or outdated infrastructure;
- Significant energy savings: *ceteris paribus* a well-managed and structured service provider will be more efficient hosting infrastructure and systems for a number of clients than each client operating its own data centre at far lower efficiency.

These factors also appear attractive to potential customers. It is important to consider these critically, however, as there is an emphasis on the conditional behind many of these points. Using a cloud service provider *should* indeed allow these advantages to be realised, but a great deal will always depend, as has always been the case with traditional outsourcing and service provision, on the contract terms that are agreed, on the quality of the service provided, and the ability of management to monitor this.

The proposed advantages in respect of the specific area of security are also subject to significant caveats. If appropriate frameworks and procedures can be developed, there are clearly advantages to be gained from having specialists manage systems and process. But against this must be set the disadvantages of having no direct control over accesses and access rights, and well as the risks introduced by having systems and data hosted on shared platforms with the systems and data of other organisations. These questions are discussed in detail in Chapter 6.

The decision to outsource services or move onto a cloud platform is not one that should be taken in isolation as a purely technical move. “An important activity that should accompany an outsourcing decision is the development and implementation of a long-term plan to improve business processes and the corresponding employment of computing platforms. This planning process should evaluate the impact of outsourcing on the organisation” (p26). Management should be certain that the decision makes strategic, financial and technical sense, especially given how difficult it can be, technically and logistically, to bring services back in-house once they have been outsourced for any length of time. Particular attention needs to be paid to the impact of such a move on human resources, the retention and maintenance of appropriate and effective skill-sets and experience being a traditional issue in IT management. If services are outsourced and staff move on, it may be difficult to replace them at some point in the future, particularly if the skills and experience required relate to older and less common technologies.

4.4.4 Legal and regulatory issues

The legal issues related to outsourcing require consideration. “Service-level agreements (SLAs) - an SLA should be entered into with mutual understanding and commitment from both the customer and the cloud provider. The service provider should have a clear understanding of the customer’s expectations and concerns.” It is essential that the SLA is complete and detailed and documents exactly what the requirements, responsibilities and expectations are from both parties, in order to ensure that services of the expected quality are provided and paid for and to avoid situations in which control activities are neglected because of confusion over whose responsibility they were. In my experience, the following elements are typically included in an SLA:

- **Application security:** it needs to be clear how security is going to be designed and implemented, and who is going to be responsible for ongoing management and monitoring. In such situations the customer prefers to manage user access rights and activities, simply requiring the service provider to host the applications, while in others it is preferred to let the service provider manage all practical access issues, responding to appropriate requests from the customer to create, modify or delete user accounts, and monitoring application access attempts and patterns;
- **Intellectual property protection:** when providing data or system hosting services to clients, service providers by definition and by necessity have wide-ranging access to the software and data in use. It should be clear that any intellectual property belonging to the client remains the client's property and that the service provider does not have the right to copy it, reproduce it, or use it elsewhere;
- **Termination:** the terms and conditions in respect of the termination of the contract should be clear, covering notice periods, any financial penalties, and how data transfer and deletions will be managed and monitored;
- **Compliance requirements:** compliance requirements that apply to the data owners in respect of systems and data will also apply to any service providers who store or process data on their behalf, and thus it is essential that the details of such compliance requirements are documented and that it is clearly understood where operational responsibility for demonstrating and reporting compliance will lie;
- **Customer responsibilities:** even if operational responsibilities for managing storage or processing are transferred to a third party, the data owners retain a certain number of responsibilities. These include informing the service provider of changes that they are making or would like to have made, providing the final confirmations and approvals for compliance issues, and ensuring that disaster recovery and business continuity plans do work in a timely and coordinated manner;
- **Performance tracking:** both parties to an outsourcing agreement should agree on how performance will be measured, which key performance indicators (KPIs) will be monitored, and how performance will be reported. This could take the form of regular reporting of statistics and periodic meetings with client relationship and technical staff;
- **Problem resolution:** both parties need to agree on how problems will be identified, communicated and allocated for resolution, as well on how the final resolutions will be recorded and agreed upon;
- **Lead time for implementation:** any transfer of systems or data, and any significant modification to the hosting environment or material being hosted, will require a period of design, implementation and testing before being finalized and put into full production. Both parties should be in agreement over the lengths of the periods that are appropriate for each kind of modification.

It is clear from this that any SLA that provides both parties with the framework and clarity that they require will be a substantial document requiring input from management, technical management and legal counsel on both sides.

In order to manage the risks related to non-compliance, it is essential to address the regulatory issues related to data processing, in particular in respect of public privacy laws, and the contracts in place

between provider and customer should specify the legal requirements and relevant legislation to which the provider should adhere.

Customer privacy takes on a particular level significance in the context of data on cloud platforms because customer or other sensitive or valuable information resident on a provider's platforms can be subject to aggregation and data mining. IT is therefore necessary for providers to provide guarantees that personally identifiable information will be protected and not used for gain. A related issue concerns how long the providers will retain such data and how such data will be securely deleted: contract terms should clearly specify the periods for data retention and the confirmations to be provided that data have been deleted.

It is also important, for compliance reasons and to avoid potentially complicated legal consequences, to consider and clarify the selection of jurisdiction – the contract usually specifies that the jurisdiction should be that of the customer – and the impact of any export issues. The cloud service provider will need to ensure that it complies with relevant export control laws that, for example, prohibit some types of encryption software being stored or transmitted outside of a country. This kind of situation might arise if, as commonly is the case, for the purposes of resilience and customer service, a cloud provider has data centres located outside a nation's boundaries.

4.4.5 Service management processes

In order to attract and retain discerning and professional customers, cloud providers need to be able to demonstrate an appropriate level of IT Service Management processes and functions. These include:

- Fault tolerance: providing acceptable availability of necessary resources through redundancy and flexibility;
- Visibility: monitoring performance and making adjustments and improvements as required;
- Security: intrinsically linked to risk management processes and compliance requirements;
- Control: managing resources, physical and virtual, across shared platforms;
- Automation: employing reliable and efficient technologies to maintain performance and security and reduce the costs of managing environments;
- Performance: meeting the defined and essential requirements of customers and optimising the efficiency and utilisation of resources.

Cloud customers are also required to implement and operate effective IT service management in order to be able to evaluate the quality, effectiveness, and cost of contract services, as well as review operational reports and financial aspects of the service provision. In particular, they need to develop procedures to follow up with the cloud provider in order to determine whether specified compliance requirements are being met and SLAs are being satisfied. This is a reason why cloud and other outsourcing services do not automatically present as much of a cost saving as initial estimations and budgets might suggest. In order to manage the contract, the customer needs to maintain the in-house competence and resources to be able to monitor the provision of services, the quality of service and the compliance with contract terms. Depending on the nature and complexity of the services being provided, and the scale of the contract, such monitoring can involve the employment of several full-time equivalents.

4.4.6 The overall benefits of cloud computing

Krutz and Vines outline the expected benefits of cloud computing, the key motivations for potential customers to consider making such a strategic move. Among the benefits are:

- The means to move part (or all) of IT from a capital expenditure environment to an operational expenditure environment, thereby freeing up capital for other, possibly revenue-generating, activities;
- The ability to rapidly deploy innovative business and research applications in a cost-effective manner, assuming that the service provider can ensure such services rather than being more geared towards a more static hosting model;
- The use of virtualization to detach business services from the underlying execution infrastructure, although this is not a unique selling point of the cloud. Many organisations are operating their own virtualization clusters if they have the expertise to do so, a trend driven both by a desire for the efficient use of resources and by the recognition that it is unrealistic and strategically inappropriate to try to maintain instances of specific outdated and resource-intensive infrastructure, particularly if only one legacy application is running upon it. The decision to replace legacy applications and can be difficult and fraught with complications, and so migrating them onto easier to manage virtual platforms can be a useful and sensible intermediate step;
- Disaster recovery and business continuity capabilities are, at least in theory, a core and well-managed part of the cloud service offering, although of course the customer does retain great responsibility for the overall operation and verification of the quality of contingency planning activities;
- The ability of the cloud provider to apply security safeguards more effectively and efficiently in a centralized environment, as part of their core competencies;
- The ability to select among a variety of cloud suppliers that provide reliable scalable services, metered billing, and advanced development resources, meaning that customers can shop around for the terms and conditions that suit them best, rather than operating their own infrastructure with its often significant fixed costs, and rather than being tied to one single supplier;
- The availability of scalable infrastructure that can rapidly provision and de-allocate substantial resources on an as-needed basis, as opposed to making long-term investment decisions on the purchase or leasing of infrastructure and being obliged to work within those constraints throughout a budgetary or operational period.

They conclude that “the major benefits of the cloud paradigm can be distilled to its inherent flexibility and resiliency, the potential for reducing costs, availability of very large amounts of centralized data storage, means to rapidly deploy computing resources, and scalability.”

This is a reasonable summary that considers the big picture for an organisation that is considering such a strategic move. It is very much targeted at larger organisations that are aware of potential future changes

in their requirements and are anticipating the need to be able to respond to such changes; smaller and more stable organisations would be less likely to identify all of these elements as key criteria and would instead be more concentrated on the cost and reliability aspects. Those two elements on their own, however, may well be sufficient to encourage management to decide to contract for cloud services.

4.4.7 Security aspects

In order to reduce the critical risks of data leakage and exposure, it is essential for prospective cloud customers to consider all aspects of security in relation to the use of cloud services. The significance of the CIA triad is paramount, in particular *the first element, confidentiality*, which is viewed as applying to numerous types of data and information. These include:

- Intellectual property (IP), which includes inventions, designs, and artistic, musical, and literary works;
- Covert channels, which are unauthorized and unintended communication paths that enable the exchange of information. These can be accomplished through timing of messages or the inappropriate use of storage mechanisms;
- Traffic analysis, which allows confidentiality to be breached by analysing the volume, rate, source, and destination of message traffic. This can even be performed to an extent if messages are encrypted, because increased message activity and high bursts of traffic can indicate that a major event is occurring. Good practice if traffic analysis is recognized to be a feasible and significant risk and countermeasures are deemed appropriate can include maintaining a near-constant rate of message traffic across network connections and disguising the source and destination locations of the traffic;
- Encryption, encoding messages so that they cannot be read by an unauthorized entity, even if intercepted;
- Inference, which is the ability to use and correlate information protected at one level of security to uncover information that is protected at a higher security level.

The *second element* in the triad is *integrity*. This requires that three principles are met:

- Modifications are not made to data by unauthorized personnel or processes;
- Unauthorized modifications are not made to data by authorized personnel or processes; and
- The data are internally and externally consistent, meaning that (as far as can be controlled by internal processes in respect of data management) internal information is consistent both among all sub-entities and with the real-world, external situation.

The *third element* in the triad is *availability*. In this context, it is the process of ensuring reliable and timely access to current and appropriate cloud data or cloud computing resources by the appropriate personnel. Availability is the guarantee that systems are functioning properly when needed, and the concept can additionally be extended to consider the availability and proper functioning of relevant security services of

the cloud system are in working order. As an example, a denial-of-service attack is a threat against availability.

A number of additional cloud information security objectives can be identified. The three properties that need to be demonstrated in order for information to be considered secure are as follows:

- **Dependability:** software that executes predictably and operates correctly under a variety of conditions, including when under attack or running on a malicious host;
- **Trustworthiness:** software that contains a minimum number of vulnerabilities or no vulnerabilities or weaknesses that could sabotage the software's dependability. It must also be resistant to malicious logic;
- **Survivability (Resilience):** software that is resistant to or tolerant of attacks and has the ability to recover as quickly as possible with as little harm as possible.

Krutz and Vines also mention seven complementary principles that support information assurance. These are confidentiality, integrity, availability, authentication, authorization, auditing, and accountability. This is significant because it conforms to the discussion of information security objectives set out in Chapter 6.

In addition, a number of more detailed and organisation-specific principles and elements of good practice can be identified that will necessarily assist in ensuring effective security. These are:

- Understand the underlying structure of the network
 - By creating application and network diagrams of the security architecture, security managers will better understand the underlying cloud structure better and thus be better equipped to achieve security control objectives.
 - The network diagram should cover how the data travels through the network (data in motion), where data are stored (data at rest), and who is using data and how (data in use). It will also be necessary to understand how cloud vendors are interacting with applications, because of the security implications of such activities.
- Implement traditional best practice security solutions
 - Security best practices should be enacted at both the traditional IT and virtual cloud layers.
- Employ virtualization best practices (if relevant)
 - Virtualized systems also require their own sets of best practices, above and beyond traditional physical IT security best practices. This is likely to be an area for which specialised external advice and support is necessary.
- Prevent data loss with backups
 - Plentiful and reliable backups are the bedrock of all successful computer operations. Redundancy and availability of backups are key objectives. Care should be taken to clarify at an early stage the split of responsibilities between the organisation and the cloud service provider: some cloud vendors provide backup services or data export processes, enabling

companies to create their own backups, while others require customers to use a custom solution or provide their own third-party application.

- It should also be ensured that, as far as is possible and technically advisable, the backup stream traversing the cloud is encrypted, and that the physical backup location and media transfer and storage are also secure.
- Regular restore tests should be scheduled, performed and reviewed.
- Monitor and audit
 - Monitoring should be continuous and audit very frequent. It is essential that the cloud vendor can guarantee that it can monitor who has accessed customer data, and if need be justify such accesses.

In summary, a number of key objectives can be identified in respect of information security. Each needs to be addressed to a satisfactory level in order for management to be comfortable over the overall quality of their security measures.

Confidentiality	Integrity	Availability	Authentication	Authorization
Auditing	Accountability	Dependability	Trustworthiness	Survivability

Table 4.5 Key objectives in respect of information security

Mather et al also discuss data security and storage. They write: “In today’s world of (network-, host-, and application-level) infrastructure security, data security becomes more important when using cloud computing at all “levels”: infrastructure-as-a-service (IaaS), platform-as-a-service (PaaS), and software-as-a-service (SaaS)” (p61). Clear distinctions are drawn between the characteristics of, and risks relating to, various aspects of data handling, including:

- Data-in-transit
- Data-at-rest
- Processing of data, including multitenancy
- Data lineage
- Data provenance
- Data remanence

They clearly state that “As with other aspects of cloud computing and security, not all of these data security facets are of equal importance in all topologies (e.g., the use of a public cloud versus a private cloud, or non-sensitive data versus sensitive data)” (p 61).

An interesting addition to the discussion of security matters comes with an analysis of the place and nature of data security mitigation. They explain that in practice it will be very difficult, if not impossible, to secure data in the cloud sufficiently to meet the requirements of particular users or customers. The only viable approach is therefore to identify data that are too sensitive and keep them in-house. “If prospective customers of cloud computing services expect that data security will serve as compensating controls for possibly weakened infrastructure security, since part of a customer’s infrastructure security moves beyond its control and a provider’s infrastructure security may (for many enterprises) or may not (for small to medium-size businesses, or SMBs) be less robust than expectations, you will be disappointed. Although data-in-transit can and should be encrypted, any use of that data in the cloud, beyond simple storage, requires that it be decrypted. Therefore, it is almost certain that in the cloud, data will be unencrypted.” This is the heart of the problem. At some point, data will have to be unencrypted in order to be used.

Mather et al also emphasize the importance of rigorous security management in the cloud. “As a customer of the cloud, you should start with the exercise of understanding the trust boundary of your services in the cloud. You should understand all the layers you own, touch, or interface with in the cloud service—network, host, application, database, storage, and web services including identity services. You also need to understand the scope of IT system management and monitoring responsibilities that fall on your shoulders, including access, change, configuration, patch, and vulnerability management. Although you may be transferring some of the operational responsibilities to the provider, the level of responsibilities will vary and will depend on a variety of factors, including the service delivery model (SPI), provider service-level agreement (SLA), and provider-specific capabilities to support the extension of your internal security management processes and tools” (p. 109-110).

This is a very important principle for client management to retain. Using third-party services is not a mechanism for abdicating any responsibilities at all in respect of security management.

4.4.8 Risk management

Before taking a firm decision to move onto the cloud, it is essential to perform a proper risk analysis in order to determine whether the benefits of cloud computing outweigh the risks that it introduces to the organisation.

The risk management process, which consists of three sub-processes, risk assessment, risk mitigation, and evaluation and assessment, is designed to minimise losses by reducing risk to an acceptable level. Risks need to be identified and analysed in terms of the likelihood of their arising and their possible consequences. They should be considered in the light of the value of the assets affected and the costs of mitigating actions.

Very early in the requirements process, cloud security should be approached from a high, non-technical, business-driving level in the organisation. Any decision to transform an organisation’s IT infrastructure should be based on and driven by business objectives, not technical objectives. Commonly four aspects of business objectives are considered:

- Business: the vendor integration challenges; the importance of data portability; the choice of data leave in-house

- Financial: the choice to build or buy; the costs of data loss; the constituent parts of a disaster recovery plan
- Legal and regulatory: key regulations including national data protection laws and international regulations
- Technical: issues such as authentication, data integrity, data flow, and privacy assessment.

Not only should these elements feature in the preliminary stages of defining objectives, but they should be referred to throughout any project that is undertaken to ensure that the scope had been adhered to. The proposed workplan in Chapter 10 returns to these points in order to ensure completeness.

4.4.9 Cloud security design principles

In order to implement effective security, attention has to be paid to the key principle of cloud security design principles. Drawing on Krutz and Vines, these can be summarised as:

Principle	Description
Least Privilege	This principle maintains that an individual, process, or other type of entity should be given the minimum privileges and resources for the minimum period of time required to complete a task, thereby reducing the opportunity for unauthorized access to sensitive information through excessive or inappropriate access rights
Separation of Duties	This principle requires that completion of a specified sensitive activity or access to sensitive objects is dependent on the satisfaction of more than one condition. Hence the initialization and authorization of a significant transaction would require signatures of more than one individual, or the arming of a weapons system would require two individuals with different keys. Separation of duties makes it more difficult for individuals to perform unauthorized activities on their own, obliging them towards the acquisition of the access credentials of others or towards collusion among entities in order to compromise the system
Defense in Depth	This principle is commonly implemented in the form of the application of multiple layers of protection wherein a subsequent layer will provide protection if a previous layer is breached. An everyday example of this is the standard configuration of corporate laptops in which, for example, the entire hard disk is encrypted and requires a valid password on boot, then there is an operating system password, and then there are individual passwords for applications and services
Fail Safe	This principle maintains that if a cloud system fails it should fail to a state in which the security of the system and its data are not compromised. This might mean configuring a system to default to a state in which a user or process is denied access to the system
Economy of Mechanism	This principle promotes the simple and comprehensible design and implementation of protection mechanisms, so that as far as possible unintended access paths do not exist or at least can be readily identified and eliminated
Complete Mediation	Under this principle, every request by a subject to access an object in a computer system must undergo a valid and effective authorization procedure, even during system initialization, shut-down, restart, or maintenance

Principle	Description
Open Design	This principle is not without dissenting views, but maintains that the most robust security emerges from mechanisms that are published and thus available for review and study, because weaknesses can be identified and remediated, or if that is impossible the mechanisms can be replaced. A part of the security community, such as the US National Security Agency (NSA), maintains that secrecy around all parts of a security system is a more reliable approach
Least Common Mechanism	According to this principle, a minimum number of protection mechanisms should be common to multiple users. This avoids shared access paths being used as sources of unauthorized information exchange
Psychological Acceptability	This principle refers to the ease of use and intuitiveness of the user interfaces that control and interact with access control mechanisms. Essentially, users must be able to understand the user interface and use it without having to interpret complex instructions, as otherwise there is commonly resistance to the measures and a tendency to try to bypass or ignore them by adopting behaviours contrary to good security practice such as writing passwords down or reusing them across systems
Weakest Link	It is a fundamental of information security that the security chain is only as strong as its weakest link. Security managers should therefore aware of where the weakest links in their chains and layers of defence are located and continually strive to improve them
Leveraging Existing Components	This principle states that rather than trying to implement an entirely new set of security mechanisms to accompany a change in infrastructure or systems, it will often be beneficial to review the state and settings of the existing security mechanisms and ensure that they are operating at their optimum design points or can be improved to reach an appropriate level

Table 4.6 Cloud security design principles

A number of aspects of the 33 NIST Security Principles [75] are particularly relevant to the cloud environment. These are:

Number	Description
1	Establish a sound security policy as the “foundation” for design
2	Treat security as an integral part of the overall system design
3	Clearly delineate the physical and logical security boundaries governed by associated security policies
6	Assume that external systems are insecure
7	Identify potential trade-offs between reducing risk and increased costs and decreases in other aspects of operational effectiveness
16	Implement layered security; ensure there is no single point of vulnerability
20	Isolate public access systems from mission-critical resources (e.g., data, processes, etc.)

Number	Description
21	Use boundary mechanisms to separate computing systems and network infrastructures
25	Minimize the system elements to be trusted
26	Implement least privilege
32	Authenticate users and processes to ensure appropriate access control decisions both within and across domains
33	Use unique identities to ensure accountability

Table 4.7 Relevant NIST security principles

There is clearly a close correspondence between these sets of principles of cloud security design and taken together they provide a solid basis for developing effective security structures, focusing on the areas of balancing risk and benefit and of separating systems and infrastructures as far as is possible.

4.4.10 Risk management

In order to manage security effectively, both good practice and experience insist that it is essential to recognise threats to infrastructure, data and access control and to evaluate how applicable and significant these are to the organisation. Common threats that apply to both the cloud environment and to traditional infrastructure include the following:

- Eavesdropping – this can include a wide range of means of gathering information such as data scavenging, traffic or trend analysis, social engineering, economic or political espionage, sniffing, dumpster diving, keystroke monitoring, and shoulder surfing. Experience indicates that eavesdropping is a primary cause of the failure of confidentiality;
- Fraud – fraud within and around information systems can take the form of collusion, falsified transactions, data manipulation, and other altering of data integrity for gain;
- Theft – this can take the form of the theft of information or trade secrets for profit or unauthorized disclosure, and the physical theft of hardware or software. When assessing the security chain in order to identify the weakest link, care should be taken not to forget the intrinsic weaknesses of material, particularly storage media, transported outside the recognised security perimeter. Stealing a USB key full of files is likely to be far easier than breaking into remote systems;
- Sabotage – this can include denial-of-service (DoS) and distributed denial of-service (DDoS) attacks, production delays, and data integrity sabotage;
- External attack - examples of these can include malicious cracking, scanning, and probing to gain infrastructure information, and the insertion of a malicious code or virus into a website or an application programme.

It is also essential to be able to identify risks that are more specific to cloud service providers so that appropriate mitigating actions can be taken, if necessary. These include the following:

- Complexity of configuration – the creating and operation of virtual systems add more layers of complexity to networks and systems, thereby increasing the possibility of suboptimal configuration or the introduction of unanticipated vulnerabilities;
- Privilege escalation – an attacker may seek to his or her system privileges by leveraging a virtual machine using a lower level of access rights to then attack a virtual machine requiring higher levels of security privileges through the hypervisor;
- Inactive virtual machines – it is not uncommon for inactive (dormant) virtual machines to remain on a server. Such virtual machines could store sensitive data, access to which might not be detected because monitoring dormant virtual machines is very difficult;
- Segregation of duties – virtual systems face similar problems to traditional physical systems in respect of the proper definition of user access roles and responsibilities. This is aggravated because virtual machines provide access to many types of components from many directions;
- Poor access controls – the virtual machine’s hypervisor is responsible for all aspects of virtual machine configuration and management, facilitating hardware virtualization and mediating all hardware access. This therefore representing a key attack vector into virtual machines through the single point of access. As a result of poorly designed access control systems, deficient patching, and lack of monitoring, the hypervisor can thus expose a trusted network to attack.

Summary

The number and range of risks to be considered when beginning to design security frameworks is large and the risk assessment process is accordingly complex and time-consuming. It is an unavoidable cornerstone of the whole security management process, however, because security measures must correspond to genuine and relevant risks. This was touched upon in Chapter 2 and will be addressed in greater detail in Chapter 6.

4.5 Security management processes

It is essential for corporate management to address the numerous security challenges that cloud computing presents: “Security managers must be able to determine what detective and preventative controls exist to clearly define the security posture of the organisation. Although proper security controls must be implemented based on asset, threat, and vulnerability risk assessment matrices, and are contingent upon the level of data protection needed, some general management processes will be required regardless of the nature of the organisation’s business.”

The management processes include the following:

- Security policy implementation;
- Computer intrusion detection and response;
- Virtualization security management

Within the security policy, a hierarchy should be respected from top to bottom, ensuring that the overall objectives can be clearly and easily linked to the detailed procedures implemented to address specific risk elements. A typical security policy hierarchy will consist of the following elements:

- Senior Management Statement of Policy – describing the rationale and the background to the policies, alongside overall objectives;
- General Organisational Policies – describing how security should be organised and overseen;
- Functional Policies – setting out the specific policies related to each area of activity and each environment;
- Mandatory Standards and Baselines – defining the reference materials and standards upon which measures are based;
- Recommended Guidelines – describing additional reference sources that are deemed appropriate for consultation;
- Detailed Procedures – describing the specific control activities to be used to respond to individual risks.

4.5.1 Security management standards

Mather et al include a useful discussion of security management standards, with the Information Technology Infrastructure Library (ITIL) and the ISO 27001/27002 standards set out and compared. The authors set out their analysis of the key management process disciplines across the ITIL and ISO frameworks that are related to security management in the cloud, identifying a set of relevant processes as their recommended security management focus areas for securing services in the cloud. These are:

- Availability management (ITIL)
- Access control (ISO/IEC 27002, ITIL)
- Vulnerability management (ISO/IEC 27002)
- Patch management (ITIL)
- Configuration management (ITIL)
- Incident response (ISO/IEC 27002)
- System use and access monitoring (ISO/IEC 27002) (p. 112-113)

4.5.2 Security-focused internal controls

Security controls in the cloud are similar in nature to those implemented in and around local-managed and hosted information systems. “The objective of cloud security controls is to reduce vulnerabilities to a tolerable level and minimize the effects of an attack. To achieve this, an organisation must determine what impact an attack might have, and the likelihood of loss. Examples of loss are compromise of sensitive information, financial embezzlement, loss of reputation, and physical destruction of resources. The process of analysing various threat scenarios and producing a representative value for the estimated potential loss is known as a risk analysis.”

Internal controls function as countermeasures for vulnerabilities, responding to clearly identified risks. Typically internal control activities are categorized into one of the following four types:

- Deterrent controls, which are designed to reduce the likelihood of a deliberate attack;
- Preventative controls, which protect vulnerabilities and attempt to prevent attacks or errors from being made;
- Corrective controls, which aim to reduce the effect of an attack by repairing the damage;
- Detective controls, which discover attacks, either in real-time or after the event, and are designed to trigger preventative actions or allow the scoping and application of corrective measures. Detective controls can warn of violations or attempted violations of security policy.

When comparing these types of control to the definitions used by such internal controls bodies such as ISACA (see Chapter 8), it is interesting to note that ISACA do not consider that deterrent controls fall within the scope of control activities in an audit and control framework. There are also contradictory views within the controls community about whether corrective controls are indeed internal control activities in the same sense as preventative and detective controls: a compelling argument can be made that these are remediation activities performed as part of a process either triggered by detective controls or performed as part of general administration and housekeeping procedures.

4.6 Information classification

When considering transferring information outside the standard security perimeter, information classification becomes even more vital. According to Krutz and Vines, there are several good reasons to classify information. “Not all data has the same value to an organisation. For example, some data is more valuable to upper management, because it aids them in making strategic long-range or short-range business direction decisions. Some data, such as trade secrets, formulas, and new product information, is so valuable that its loss could create a significant problem for the enterprise in the marketplace — either by creating public embarrassment or by causing a lack of credibility.” This refers, although not by name, to the concept of Toxic Data described above.

“For these reasons, it is obvious that information classification has a higher, enterprise-level benefit. Information stored in a cloud environment can have an impact on a business globally, not just on the business unit or line operation levels. Its primary purpose is to enhance confidentiality, integrity, and availability... and minimize risks to the information. In addition, by focusing the protection mechanisms and controls on the information areas that most need it, you achieve a more efficient cost- to-benefit ratio.”

The major benefits of information classification can be set out clearly.

- It demonstrates an organisation’s commitment to protecting security and privacy
- It helps identify which information is the most sensitive or vital to an organisation.
- It supports the fundamental principles of confidentiality, integrity, and availability in respect of data.
- It helps identify which levels and measures of protection apply to which information.

In addition, information classification might be required for regulatory, compliance, or legal reasons.

The basic principle is that the information that an organisation processes must be classified according to the organisation's sensitivity to its loss or disclosure. The information system owner – not corporate management, not IT management, is responsible for defining the sensitivity level of the data.

There exist numerous frameworks for classifying information. For illustration, we can consider a series of typical classification terms that are used in the private sector and are applicable to cloud data.

- Public data - information that is similar to unclassified information; all of a company's information that does not fit into any of the next categories can be considered effectively public. It should not be disclosed according to policy, but any disclosure should not have any seriously adverse impact on the organisation, its employees, and/or its customers.
- Sensitive data - information that requires special precautions to protect it from a loss of confidentiality and a loss of integrity due to an unauthorized alteration;
- Private data - personal information that is intended for use within the organisation, unauthorized disclosure of which could seriously and adversely impact the organisation and/or its employees. Typically salary levels and medical information are considered to fall into this category;
- Confidential data – the most sensitive business information, intended strictly for use within the organisation. Unauthorized disclosure of such information, such as information about new product development, trade secrets, and merger negotiations, could seriously and adversely impact the organisation, its stockholders, its business partners, and/or its customers.

Each organisation can choose to use a high, medium, or low classification scheme based upon its CIA needs and whether it requires high, medium, or low protective controls. For example, a system and its information may require a high degree of integrity and availability, yet have no need for confidentiality. This is a key principle of both data classification and wider security management: measures should not be blindly followed and adopted because that is simply what the guidance says; rather, measures should be selected based on the specific circumstances and requirements of the organisation.

It should be clearly understood that “the designated owners of information are responsible for determining data classification levels, subject to executive management review.” This is consistent with other aspects of security management, such as the definition of user access rights to applications and data.

There are a number of criteria that may be used to determine the classification of an information object. These include:

- Value – the value of the information to the organisation. This is the most commonly used criterion for classifying data; if the information is valuable to an organisation or to its competitors, then it needs to be classified;
- Age – a case can be made for lowering the classification of information as the information's value decreases over time. A public example of this might be United Kingdom census information, which can be published (and indeed is actively made available) one hundred years after it has been collected, after a century in which it is subject to strict confidentiality measures;

- Useful life – information can have a limited lifespan due to its nature or can be made obsolete due to the creation or acquisition of new information, substantial changes in the company, or other reasons. In such circumstances, the information can often be declassified;
- Personal association – if information is personally associated with specific individuals or is specifically addressed by a privacy law, it might need to be classified.

The links between classification and security and audit are clear. Effective classification leads to efficient and effective security measures, while clear and effective security measures allow efficient and conclusive auditing procedures to be performed, giving management, auditors and regulators the comfort they require over the quality of information security.

The procedures required for information classification are reasonably straightforward. In summary they can be reduced to seven key steps, namely:

1. Identify the appropriate administrator and data custodian. The data custodian is responsible for protecting the information, running backups, and performing data restoration.
2. Specify the criteria for classifying and labelling the information.
3. Classify the data by its owner, who is subject to review by a supervisor.
4. Specify and document any exceptions to the classification policy.
5. Specify the controls that will be applied to each classification level.
6. Specify the termination procedures for declassifying the information or for transferring custody of the information to another entity.
7. Create an enterprise awareness program about the classification controls.”

These steps ensure that appropriate resources are in place and that responsibilities are clearly defined. They also ensure that the project is visible to, and understood by, users throughout the organisation. Experience shows that such awareness building and information sharing is essential for the success of a project, especially if there will be any reliance on users to update or modify any information classifications. It is also essential that they know how to react when confronted by new or unexpected types or sources of information.

4.7 User account management

An extremely important topic for consideration is access control in the cloud, with a number of key questions highlighted for attention by Mather et al. These are:

- Who should have access to what resource? (Assignment of entitlements to users)
- Why should the user have access to the resource? (Assignment of entitlements based on the user’s job functions and responsibilities)
- How should you access the resource? (What authentication method and strength are required prior to granting access to the resource)

- Who has access to what resource? (Auditing and reporting to verify entitlement assignments) (p. 124)

Practically these concerns are identical to those relevant to in-house systems security, but with the added complication of users wanting or requiring access who are not vetted employees working within the existing security perimeter.

The creation of user accounts and associated access rights is commonly reasonably well-controlled within organisations, usually for the simple reason that new staff cannot begin to work until access to key systems has been granted. The deletion of accounts is in practice far less well controlled and monitored and this weakness can create significant vulnerabilities in the security framework. This is because the use of a genuine – if no longer needed – user account to gain access to programmes and data is unlikely to attract the attention of security officers or security software in the way that an attempt to break into systems remotely would do.

Management should ensure that the employee termination is appropriately designed and implemented in order to respect overall policies with regard to information security. “It is important to understand the impact of employee terminations on the integrity of information stored in a cloud environment. This issue applies to employees of the cloud client as well as the cloud provider.”

There are two types of terminations, friendly and unfriendly, and both require specific actions. Friendly terminations are easier to manage and can be accomplished by implementing a standard set of procedures, procedures that can also be applied to staff transfers where changes to roles and responsibilities and thus to access rights are anticipated. These procedures will usually include the following steps:

- The removal or deletion of access rights, computer accounts, authentication tokens such a cryptographic tags;
- A briefing for the leaver on continuing responsibilities in respect of confidentiality and privacy;
- The return of corporate information systems equipment and material, such as laptops, remote access tags, smartphones and external storage devices.

Unfriendly terminations will require each of these steps, but will also often require additional advance planning in order to restrict access to systems and data from the moment that the decision to terminate is taken or communicated, for example, or to secure equipment before it can be damaged.

In both cases it will always be necessary to consider the continued availability of data. Ideally policies and procedures will be in place to ensure that data are never stored on the local drives of laptops but are always kept on corporate servers and thereby both available to colleagues and management with appropriate credentials and subject to standard corporate backup and recovery procedures. As a matter of course employees should be instructed whether or not to “clean up” their PC before leaving.

4.8 Security awareness, training and education

An essential element of effective security management consists of security awareness, training and education. “Security awareness is often overlooked as an element affecting cloud security architecture because most of a security practitioner’s time is spent on controls, intrusion detection, risk assessment, and

proactively or reactively administering security. Employees must understand how their actions, even seemingly insignificant actions, can greatly impact the overall security position of an organisation.”

It is essential that employees of both the cloud client and the cloud provider are fully and practically aware of the need to secure information and protect the information assets of an enterprise. Very specifically, operators need ongoing training – at least as a refresher - in the skills that are required to fulfil their job functions securely, while security practitioners need updated and relevant training to implement and maintain necessary security controls.

More generally, all employees need education and frequent refreshers in the basic concepts of security and its benefits to an organisation. If the three pillars of security awareness training - awareness, training, and education – are sufficiently robust, improvements will be seen in the behaviour and attitudes of staff and through a quantifiably significant improvement in the organisation’s security, detectable in reduced numbers of support calls and incidents related to ignorance of, or departure from, good security practice.

These three elements, computer security awareness, training, and education improve security by doing the following:

- Improving understanding of the need to protect system resources, infrastructure and data;
- Developing skills and knowledge so that computer users can perform their jobs more securely and more efficiently, without exposing systems or data to security-related risks;
- Developing the detailed knowledge necessary for the design, implementation and operation of well-structured and appropriate security programs.

4.9 Identity management

Another fundamental principle to be addressed properly is identity management. According to Krutz and Vines, “Identification and authentication are the keystones of most access control systems. *Identification* is the act of a user professing an identity to a system, usually in the form of a username or user logon ID to the system. Identification establishes user accountability for the actions on the system. User IDs should be unique and not shared among different individuals. In many large organisations, user IDs follow set standards, such as first initial followed by last name, and so on. In order to enhance security and reduce the amount of information available to an attacker, an ID should not reflect the user’s job title or function.”

The process and methods of authentication are deserving of more detailed consideration, given its crucial importance is granting or refusing access to systems and data. *Authentication* is verification that the user’s claimed identity is valid, and it is usually implemented through a user identifier and password at login.

Authentication can require more than one factor of identification. The different factors can be categorised into three, possibly four, types:

- Type 1 - something you know, such as a personal identification number (PIN) or password;
- Type 2 - something you have, such as an ATM card or smart card or a code sent by SMS to a mobile phone;

- Type 3 - something you are (physically), such as a fingerprint or retina scan.

For some authorities there is a Type 4 - something you do, such as the speed or rhythm of typing a set phrase or accuracy in using a mouse. Other authorities consider this to be a subset of Type 3 - something you are.

Two-factor authentication, which is commonly implemented for many sensitive everyday activities, requires two of the three factors to be used in the authentication process. For example, logging into e-banking systems usually requires following the user ID with a set password (something you know) and an additional session-limited access code generated by an electronic tag or sent by the bank via SMS to a pre-registered mobile phone (something you have).

Implementing effective identity management is no trivial matter. “Realizing effective identity management requires a high-level corporate commitment and dedication of sufficient resources to accomplish the task.”

There are a number of phases in a typical project for implementing identity management. These include the following:

- Defining the strategy in respect of identity management;
- Determining the cost-benefit ratios of the various mechanisms available and choosing the most suitable;
- Establishing a database of identities and credentials;
- Managing user access rights;
- Communicating and enforcing security policy;
- Developing the capability to create and modify accounts in a controlled and monitored way;
- Setting up monitoring of accesses – and access attempts – to resources;
- Installing procedure for removing access rights;
- Providing training in proper procedures.

Mather et al also consider Identity and Access Management (IAM). The challenges related to this field are identified as “managing access for diverse user populations (employees, contractors, partners, etc.) accessing internal and externally hosted services. IT is constantly challenged to rapidly provision appropriate access to the users whose roles and responsibilities often change for business reasons.” The authors also discuss the impact on access management of the turnover of users within the organisation and of the absence and inconsistency of access policies for information. The limitations of possible solutions are discussed, with mention of the weaknesses of technology solutions and the long timescales that are often encountered before results are apparent (p. 76-77).

4.10 The importance of auditing and testing

There are a number of ways in which management can gain comfort over the effectiveness of the security measures in place over their systems and data. One particular method is highlighted by Krutz and

Vines because of its attractiveness in a cloud context: unlike other audit methods, there is no requirement for physical access to documentation and material in order to perform a penetration test. A second advantage is that it replicates the likely means and techniques employed by genuine external attackers, thus giving an accurate picture of vulnerabilities and how they might be exploited.

The importance of audit and compliance is emphasized by Mather et al [67], taking care to specify that “audit and compliance refers to the internal and external processes that an organisation implements to:

- Identify the requirements with which it must abide—whether those requirements are driven by business objectives, laws and regulations, customer contracts, internal corporate policies and standards, or other factors
- Put into practice policies, procedures, processes, and systems to satisfy such requirements
- Monitor or check whether such policies, procedures, and processes are consistently followed” (p. 167)

It is further explained that “Cloud service providers (CSPs) are challenged to establish, monitor, and demonstrate ongoing compliance with a set of controls that meets their customers’ business and regulatory requirements. Maintaining separate compliance efforts for different regulations or standards is not sustainable. A practical approach to audit and compliance in the cloud includes a coordinated combination of internal policy compliance, regulatory compliance, and external auditing” (p. 167).

4.10.1 Governance, Risk and Compliance

A concept that has become increasingly popular over the last decade is Governance, Risk and Compliance and this is presented and explained. The key components are the risk assessment, the identification and documentation of key controls, the monitoring and testing process, the reporting of performance and findings, and the process of continuous improvement.

A number of external compliance drivers have an impact on the provision and quality of audit services; these include in particular the Sarbanes-Oxley Act of 2002 (which has relevance to cloud computing services), PCI DSS where service providers store or process cardholder data, and HIPAA in respect of individually identifiable health information.

Mather et al explain the distinctions between the Internal and External Audit Perspectives, mentioning the difference in focus and objectives of the two processes (p. 195), and the various commonly-used audit frameworks such as SAS 70 – ISAE 3402, SysTrust, WebTrust, and ISO 27001 are compared and contrasted (p. 196).

A number of governance factors to consider when using cloud computing are described, and the key processes identified. These include:

- Managing identity
- Provisioning access
- Defining data storage requirements

- Managing key management
- Monitoring and managing service levels
- Monitoring and maintaining availability
- Providing assurance on internal controls
- Providing secure connectivity
- Providing for data governance
- Managing for problem management and incident response
- Developing, maintaining, and, when necessary, executing a business continuity program (p. 235)

4.10.2 Penetration testing

The importance and even necessity of penetration testing to gain comfort over the quality of security measures is emphasised by Krutz and Vines. “A penetration test is a security testing methodology that gives the tester insight into the strength of the target’s network security by simulating an attack from a malicious source. The process involves an active analysis of the cloud system for any potential vulnerabilities that may result from poor or improper system configuration, known and/or unknown hardware or software flaws, or operational weaknesses in process or technical countermeasures. This analysis is carried out from the position of a potential attacker, and can involve active exploitation of security vulnerabilities. Any security issues that are found are presented to the system owner together with an assessment of their impact, and often with a proposal for mitigation or a technical solution.”

A detailed and professional fully-scoped penetration test is a component of a full security audit. Although targeted tests can always highlight areas of weakness, they are more efficient and effective when performed as a part of a structured review of security that considers how security is appreciated, designed and implemented throughout an organisation. Revealing technical weaknesses is interesting but meaningful recommendations, which are the key output of audits, can only be made with an understanding of the whole environment, where the most significant weaknesses are, and what the priorities should be for making improvements.

Generally speaking a security audit will consist of three levels of activities, performed sequentially. The first element will always be a high-level assessment, which involves and top-down look at the organisation’s policies, procedures, standards, and guidelines. This will be hands-off and involve consideration of consistency with standards and good practice, rather than involving access to the systems under review. The next element involves more information gathering, from document review, interviews with staff, and some preliminary scanning of the network. This enables the auditor to continue to identify high-level weaknesses and to determine which aspects of security should be tested in detail and which systems and infrastructure elements might be vulnerable to attack. The third element is the penetration test itself, where tools and expertise are deployed in order to try to gain access to systems or data. In contrast to the previous phases, which are performed very much from a management perspective, the penetration test is more adversarial and puts the auditor in the position of hacker, seeing what is visible and what is possible.

There are several approaches to penetration testing: “black box”, where the attackers have no prior knowledge of the infrastructure to be tested and have to make discoveries and deductions as they go about their work; “grey box”, where they have partial knowledge; and “white box”, where they have full knowledge. (Interestingly, the same classifications can be applied to software testing, particularly at the functional and user acceptance testing stages). Typically a penetration test will involve three phases: Preparation, in which conditions, objectives and terms of reference are defined and issues relating to data protection and legal protection for the tester are clarified and formalised; Performance, during which the tester performs the attacks and looks to exploit any vulnerabilities found; and Reporting, in which the findings of the tests and recommendations for improvement are communicated to the client. Two aspects of the process are particularly important: firstly that the legal and compliance impacts of the testing are appropriately addressed in order to avoid issues arising from the success of any tests, the attacker and the data owner potentially being in breach of confidentiality and disclosure requirements if sensitive data are accessed; and secondly that the findings are understood by management and appropriate remedial actions undertaken, if necessary. Commissioning an audit is only half of the audit and control process – management needs to be able to respond to the findings and make improvements where indicated.

A number of factors have come together in the cloud environment that make penetration testing an essential part of audit and control activities. These include:

- The increased ease of use of information systems for end-users coupled with the increased complexity of computing resources, resulting in greater demand for access to resources and services without necessarily an accompanying improvement in administration skills and resources;
- The skill level required to execute a hacker exploit has steadily decreased, meaning that attacks are more widespread;
- Increased high-speed connectivity means that more potential attackers are in a position to launch attacks;
- The size and complexity of the network environment has mushroomed, making effective security management more difficult;
- The number of network and cloud-based applications has increased, providing more targets and greater scope for attacks;
- The detrimental impact of a security breach on corporate assets and goodwill is greater than ever, as compliance requirements become more stringent and third-parties are better informed about and more critical of organisations having been attacked.

Summary

Penetration testing is a key part of a wide-ranging security testing approach but its value and practicality should not be overstated. From the perspective of the degree and scope of comfort that it provided, there are two important limitations. Firstly, it only applies to the elements of the infrastructure and applications that have actually been tested. Testing of perimeter controls can give some overall comfort about the quality of perimeter defences, but testing of the security of one application can only give comfort as to the security of that application. Secondly, the results of each test only relate to the means and methods used during that test. If the testers use old and outdated tools or exploits that might already have been patched,

it is possible that they will not identify weaknesses that could be exploited by more up-to-date techniques. Penetration testing runs in parallel to technical hacking activities, where being completely up-to-date and equipped with the latest tools and information are essential for both attackers and defenders.

From an organisational perspective too there are limitations to the effectiveness of penetration testing. Such tests need to be carried out with the approval of the service provider, and this might not necessarily be fully forthcoming. It is possible that the service provider is sensitive enough about their security measures not to want them tested by any third party. There is also a risk management aspect in the part of the service provider to consider. They may have security measures in place that would block or interrupt hacking attempts but also block access to legitimate external users, at least temporarily. They might also be aware that a successful attack, even a white hat one, could lead to access issues for other users as items of communications equipment or network elements were manipulated or shut off. As a result of this, getting appropriate approvals to carry out penetration tests might not be straightforward.

4.11 Contingency planning

Business continuity planning (BCP) and disaster recovery planning (DRP) are fundamental aspects of systems and business management, regardless of whether systems and data are maintained in-house or outsourced to a service provider. These activities “involve the preparation, testing, and updating of the actions required to protect critical business processes from the effects of major system and network failures. From the cloud perspective, these important business processes are heavily dependent on cloud-based applications and software robustness and security. BCP comprises scoping and initiating the planning, conducting a business impact assessment (BIA), and developing the plan. DRP includes developing the DRP processes, testing the plan, and implementing the disaster recovery procedures.”

The resources and efforts required in respect of effective planning should not be underestimated. “Designing, developing, and implementing a quality and effective BCP and DRP is a major undertaking, involving many person-hours and, in many instances, high hardware or software costs. These efforts and costs are worthwhile and necessary, but they impact a large number of organisational resources.”

Two separate scenarios should be considered when discussing BCP and DRP in the context of the cloud. One is where an organisation has already outsourced part or all of its activities and infrastructure to the cloud and needs effective planning in respect of this, to ensure that operating requirements are met, to ensure that the activities being undertaken by both parties are coordinated and synchronised, and to receive sufficient assurance that the service provider has implemented efficient and effective measures. The other is where an organisation operates its own systems and infrastructure in-house but has contracted with a cloud service provider to provide services if necessary in the event of a loss of service. Here again management needs to ensure that the measures in place conform to the objectives set for recovery activities and correspond to any measures by operated locally.

A BCP includes a business impact assessment (BIA), which, itself, contains a vulnerability assessment. A BIA is a process used to help business units understand the impact of a disruptive event by evaluating the potential and probable effects of incidents leading to interruption or loss of service. A vulnerability assessment is similar to a risk assessment in that it contains both a quantitative (financial) section and a qualitative (operational) section, but differs in being focused on providing information that is used solely for the business continuity plan or disaster recovery plan.

There are a number of benefits of using the cloud for BCP/DRP. “Adopting a cloud strategy for BCP/DRP offers significant benefits without large amounts of capital and human resource investments. Effective cloud-based BCP/DRP requires planning, preparation, and selecting the cloud provider that best meets an organisation’s needs. A critical issue is the stability and viability of the vendor. The vendor should have the financial, technical, and organisational resources to ensure it will be around for both the short term and the long term. In addition, in order for cloud BCP/DRP to reach its full potential, standardization across a variety of architectures has to evolve.” The basic motivations for choosing a cloud service are then costs and delegation, transferring the burden of the strategic and operational decisions to a specialist organisation. Management should not forget, however, that ultimate responsibility for ensuring the adequacy of contingency planning rests entirely with them, and that they have a crucial ongoing role in verifying the readiness of the service provider to respond to requests to activate contingency plans and in informing the service provider of any changes in policies, procedures, activities or requirements that might require corresponding changes in contingency planning or in the technologies used.

Krutz and Vines list a number of factors that should be included in the design of cloud-based IT systems that meet the requirements of a BCP and DRP. These include:

- Secure access from remote locations
- A distributed architecture with no single point of failure
- Integral redundancy of applications and information
- Geographical dispersion

These factors do of course apply to non-cloud contingency planning, above all in the situations where an organisation chooses to create one or more warm or hot sites to be used for remote processing.

Summary

Contingency planning is a key element of information systems management regardless of the specific choices made in respect of platforms and environments. The impact of cloud services on contingency planning should not be neglected because of both the opportunities they bring and the risks that are also created. Typically it is during crisis situations that internal controls are at their weakest and management should ensure that their contingency planning is stress tested and reliably robust in order to ensure both adequate continuity of service and minimal opportunity for data loss or leakage.

4.12 Conclusion

It is clear from consideration of the authorities discussed above that use of cloud services is far from being an easy option for organisations, one by which they can pay a service provider to assume the responsibility for their data or systems and concentrate on their own daily operational activities.

Responsibility for the security and privacy aspects of data storage and utilisation remains uncontestedly with the data owners. The fact of outsourcing data to a third party for management, storage and exploitation does not remove this responsibility.

Exerting sufficient control to be able to discharge these responsibilities is not trivial. Without the provision of suitable means of communication with the service provider and the possibility to obtain rapid, current and valid information about the status of data, the uses to which they are being put, and how and by whom the data are being accessed, the management of any organisation contracting outsourced services will be incapable of obtaining the information they require to be able to confirm that there are or have been no significant security breaches. This is in marked contrast to the situation when data are stored within an organisation's security perimeter and all necessary information can be provided and all necessary safeguards put into place, at least in theory.

Control measures must necessarily be driven by an evaluation of risks and the identification and assessment of these risks is clearly, on the basis of the opinions of the experts cited above, not a simple or rapid process. In particular management will need to become familiar with a number of aspects of remote management and reliance upon third parties in order to develop a clear and useful picture of the risk environment.

The security challenges facing organisations in respect of cloud services are many. They need to be confident that there are no design or operational deficiencies in respect of infrastructure security, data security and storage, identity and access management, and overall security management. Each of these areas requires experience and expertise to identify the specific risks, evaluate their importance, decide upon a strategy (avoid, mitigate, transfer, accept), and design and implement appropriate internal controls. Some of these controls will be operated by and reported on by the service provider (see Chapter 6) but others will be operated and evaluated by the organisation itself and great care will need to be taken to ensure that the combination of controls chosen does correspond to the defined approach to addressing the identified risks.

Management also need to be aware of the implications of moving onto cloud platforms. The initial move requires careful planning, professional project management, and an accurate understanding of current, relevant risks. It is important for management not to be blinded to the long-term consequences of such a strategic decision though. Firstly, such a commitment requires them to plan for regular reviews of the operating environment and its appropriateness for their needs. Secondly, and perhaps more importantly, they may find themselves locked into a relationship with the service provider that has no clearly defined exit strategy. Moving away from such a core business supplier, particularly once systems have become bedded down and start to undergo even an unconscious and positively intended customisation towards the service provider's infrastructure and practices, may prove to be a complicated and costly process, perhaps even prohibitively so.

Chapter 5 Identification of Needs, Trends and Awareness: User and Market Surveys

5.1 Introduction

“Le savant n'est pas l'homme qui fournit de vraies réponses; c'est celui qui pose les vraies questions.”⁶

The review of relevant literature presented in Chapter 3 and 4 gives a great deal of theoretical background to the application and use of unstructured data, big data, and cloud technologies. The objectives, risks and benefits arising from the use of these technologies are discussed in detail, with reference to standards, requirements and obligations, but apart from a small number of sources there is little information on the reality of the adoption of these techniques and services by organisations and businesses.

In order to gain practical insights into the reality of the use of unstructured data and the take up of cloud services, it was therefore necessary to obtain some concrete and up-to-date data on the activities of organisations, identifying their choices and the reasoning behind those choices. The most effective way of doing so is to take a snapshot of the marketplace by means of a standardized and easy to complete survey.

In this chapter I therefore present the results of two surveys of the use of the cloud.

The first survey I carried out within Switzerland in 2011 and 2012, asking organisations about their current and planned use of unstructured data and cloud services. The second survey was performed and published by the Cloud Security Alliance and ISACA (formerly known as the Information Systems Audit and Control Association) in 2012 and concentrates on more on approaches to using cloud services. This survey provides, I believe, a useful counterpoint to my own limited research because of its scope and focus: many of the questions and themes are complementary, particularly with the underlying focus on governance and internal controls, but because it was targeted at corporations based in the United States there is no overlap in the datasets. The results of the ISACA survey can therefore be seen as supporting and extending the results of my own survey.

The choice of the ISACA survey as a support to my own was driven by two factors. The first was the depth and clarity of the details presented in its publication: many surveys have appeared in the technical and business-oriented press, but these tend to be fairly superficial and the detailed findings are rarely available. The second was that as an audit and controls organisation (of which I have been a full member since 1996), the focus of ISACA's work closely parallels my own objectives and interests. Questions of risk management, control and compliance underlie both surveys.

⁶ Claude Lévi-Strauss, “Le cru et le cuit”, p. 15

The objectives behind this approach were to identify and understand the attitudes of organisations in respect of the subjects mentioned above, to confirm the relevance of the results arising from the literature review presented in Chapters 3 and 4, and to ensure that any conclusions drawn and proposals made in Chapters 6 to 10 correspond to real-world situations and respond to genuine needs and risks.

5.2 The situation in Switzerland : original research into current trends

In order to gain an independent perspective on how the questions of cloud facilities and unstructured data were affecting organisations in Switzerland, a survey was carried out in 2011 and 2012.

5.2.1 Background to the survey

50 questionnaires were sent out to contacts in 43 organisations across Switzerland, using the lead author's business experience to identify correspondents across a range of industries and sectors of activity who would be likely to respond. Completed questionnaires were received from 34 organisations, with three sending two responses and two sending three. The organisations that declined to respond did so on the grounds of confidentiality, not wishing to divulge details of their IT strategy or approach to security to a third party.




Figure 5.1 Responses received by organisation (population of 43)

The organisations were selected in order to provide a wide cross-section of the range of IT environments and attitudes to the management of data and the use of new technologies. Of the organisations that did respond, five were publicly owned, eighteen privately, and the remaining eleven were public administrations or NGOs. The breakdown by organisation size also shows variety: eight with less than fifty employees; ten with between fifty-one and one hundred; five with between one hundred and five hundred; and eleven with more than five hundred.




Figure 5.2 Ownership of responding entities (population of 34)




Figure 5.3 Size of responding entities in terms of number of employees (population of 34)

The rationale behind sending multiple questionnaires to the same organisation was to attempt to discover whether there would be cases of poor communication of strategy or developments within those organisations. It is the author's experience of large corporations in particular that there can be differences in the understanding of the overall strategy, the firm objectives, the initiatives undertaken and the impact on users between top management, middle management and IT management, for example.

5.2.2 The Survey

The questions in the survey were split into two sections, dealing with data security within the organisation and with data storage in the cloud, as follows:

Section A

Q1	Is there an overall policy concerning data management, security and retention?
Q2	Is there awareness at the level of senior management and/or security management of the concept of “unstructured data”?
Q3	Has there been any assessment of whether the organisation should be concerned, from security or efficiency perspectives, about the existence of unstructured data?
Q4	Has the organisation established any guidelines on the classification of data according to their sensitivity, age and relevance?
Q5	Has there been any structured attempt to identify and quantify the data stored around the organisation outside centralized databases?
Q5A	If so, was this a manual process?
Q6	Does the organisation have policies on the extraction, use and storage of data by end-users?
Q6A	If so, is compliance with these policies monitored and enforced, and how?
Q7	Are there policies and/or restrictions in place of the use of removable storage media for file transfer or storage?
Q7A	If so, is compliance with these policies monitored and enforced and how?
Q8	Does the organisation have any mechanisms for determining whether there have been breaches of security or confidentiality in respect of its sensitive data?

Section B

Q9	Is the organisation using, or planning to use, cloud computing services for the storage of data?
Q10	If yes, have policies and guidelines been drawn up for the nature of data that can be stored in this way?
Q11	If cloud computing is being used, is the organisation’s approach based on the centralized management, monitoring and retrieval of data, or do departments and/or individuals retain responsibility for their data?

Table 5.1 Survey questions

These questions were designed to establish how much management planning and consideration of the importance of internal controls had been performed in any relevant projects. These questions were asked from a typical information systems auditor perspective, in which there is an expectation that in respect of IT-related projects, and indeed in respect of all projects and all daily operations within an organisation,

there will be compliance with good practice with regard to risk management and internal controls and that there will exist a clear and coherent statement of control objectives, relevant risks and related control activities, and of the links between the items in each of these three categories. This is a distinguishing factor in comparison to other surveys, which often focus on technical choices and financial considerations, and forms part of the underlying theme of this thesis, which is the importance of risk management and internal controls in strategic choice and the implementation of changes to technology and operating practices.

5.2.3 Survey results

The detailed results of the survey are presented in Appendix B; here I present the results in a summary form that supports the following analysis.

Section A				
	Yes		No	
	Yes	No	Yes %	No %
Q1	32	9	78%	22%
Q2	22	19	54%	46%
Q3	17	24	41%	59%
Q4	9	32	22%	78%
Q5	6	35	15%	85%
Q5A	6	0	100%	0%
Q6	14	27	34%	66%
Q6A	5	9	36%	64%
Q7	12	29	29%	71%
Q7A	8	4	67%	33%
Q8	4	37	10%	90%
Section B				
Q9	39	2	95%	5%
Q10	6	33	15%	85%
Q11	7	32	18%	82%

Table 5.2 Summary survey results

5.2.4 Interpretation and analysis of the results

The results demonstrated two major trends in respect of unstructured data. The first was that although 78% of organisations reported having designed and implemented an overall policy concerning data management security and retention, the exceptions being overwhelmingly small organisations with informal internal control structures, there was little systematic follow-up in terms of the management of unstructured data or in terms of managing and monitoring the use of data by users. 54% reported that senior management were aware of the issue of unstructured data; but only 41% reported that a risk analysis had been carried out to evaluate their exposure; 22% reported that guidelines for the classification of data had been developed and published; and only 15% reported having carried out a structured attempt to identify and quantify the data stored outside centralized databases. In each case it was a large multinational company that had undertaken such an initiative; interestingly, each one reported that the process had been largely manual.




Figure 5.4 Selected responses to Section A (population of 41)

From the perspectives of internal controls and good corporate management, many of these numbers are worryingly low. That more than three quarters of organisations report the existence of an overall policy in respect of data management is a reasonable starting point, but the lower numbers in respect of detailed data management indicate that few organisations had progressed further than the big picture, high-level aspects of data management at the time this survey was performed.

In respect of the management of user activities related to data handling, 34% reported having policies on the extraction, use and storage of data by end users, and only 36% of these were able to report that

compliance with these policies was monitored and enforced. Only 29% of organisations reported having policies or procedures in place concerning the use of removable storage media for file transfer or storage, and 67% of these were able to describe compliance mechanisms. Finally, only 10% of organisations were able to describe the existence of mechanisms for determining whether breaches of security or confidentiality in respect of sensitive data had occurred.

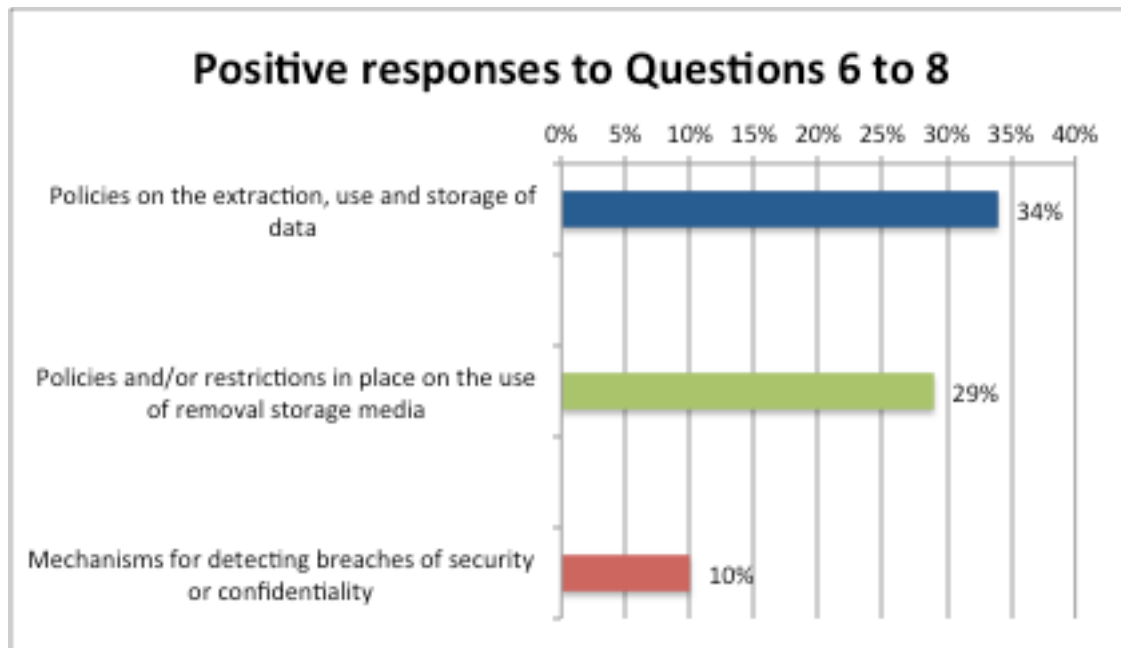


Figure 5.5 Positive responses to questions 6 to 8 (population of 41)

These rather low numbers suggest that organisations will have a significant amount of work to do in collating, cleaning and preparing data. The responses suggest an overwhelming absence of effective internal control to date over data usage, and a lack of certainty over the nature, quality and integrity of the datasets in use around organisations, factors that will automatically increase both the amount of scoping work that needs to be carried out and the quantity of detailed cleaning and tidying of data.

In the cases where more than one response was received from an organisation, it was noted that end users were less aware of the existence of strategies and policies than senior management, a situation that is consistent with the lead author's long experience of auditing large organisations.

In respect of the use of cloud computing facilities for the storage of data, 95% of the responses reported that the organisation was using or was planning to use such facilities. The reasons most frequently given were cost management, flexibility, and a desire to streamline IT activities to concentrate on more value-added activities in-house. Of these organisations, however, only 15% had drawn up policies and guidelines concerning the nature of data that could be stored in this way, and 18% were adopting an approach based on the centralized management, monitoring and retrieval of data rather than leaving it to departments or individuals to manage.

Once again, the wide absence of overall policies and guidelines and the tendency as reported to adopt potentially uncoordinated approaches to project scope and management runs counter to established good

practice in respect of internal controls. Without clear and enforced structure, inconsistency is likely to become a significant barrier to success. In addition, delegating down to departments or individuals increases the risks of decisions being taken without sufficient skills, experience or perspective.

Overall it was possible to draw a clear distinction between large and small organisations and publicly and privately owned ones in respect of their approach to structured and formalized control environments. It was also possible to identify with high accuracy the responses received from publicly-owned corporations and those from organisations subject to other strict and demanding controls requirements such as Sarbanes-Oxley or local industry or environment-specific regimes. It was also possible to identify organisations that had significant internal audit or internal controls functions, or that had been alerted to the risks involved in going through a process of discovery in the context of a legal dispute.

5.2.5 Summary evaluations

The tentative conclusions that can be drawn from this very specific and targeted survey are the following.

5.2.5.1 *A structured approach to data management and security*

Firstly, the importance of having a structured and documented approach to data management and security is widely understood among the organisations surveyed, with a particularly positive attitude towards risk management and compliance from larger organisations and those subject to definite compliance regimes because of their ownership or industry. How this understanding actually translates into positive measures designed to ensure compliance is another question, however, and IT security managers in particular reported seeing greater enthusiasm for establishing policies within their organisations than for implementing and complying with the necessary procedures. There was also a question of priorities and resources raised by smaller organisations that did not feel that such policies corresponded to or were a part of their core daily activities.

The significance of compliance needs to be underlined because it is a core element of internal control as an unavoidable objective and requirement for many organisations that will shape and motivate many of their decisions in respect of security.

5.2.5.2 *Recognition of unstructured data*

Secondly, the concept of unstructured data and the particular challenges posed by such data is reasonably widely understood, but there has been little activity outside large publicly-owned corporations to address the issue in any systematic way. In general IT departments had a good grasp of the nature and impact of the matter, and the subject was frequently raised by internal and external auditors and by legal advisers, but it was rarely considered to be a subject of great priority by senior management.

5.2.5.3 *Monitoring of the performance of controls*

Thirdly, even if thought has been given within some organisations to managing employee access to and extractions of sensitive data, in the majority of cases monitoring is weak and compliance cannot be ensured. Generally speaking, users who have access to data stored in centralized databases tend to have the ability and the opportunity, and frequently the encouragement, to extract those data and use them for analytical or reporting purposes. Once the data have been extracted, access controls around them are

usually weaker, often being restricted to network or workstation access controls, and even these restrictions reach their limits once data are copied onto portable devices such as USB keys.

The significance of this widespread weakness with its impact on demonstrable compliance is again not to be underestimated. As will be discussed in subsequent chapters, it is not sufficient simply to have designed detailed policies and procedures in respect of a certain activity; it is necessary to be able to demonstrate that these policies and procedures have systematically and consistently been followed.

5.2.5.4 Detection of security breaches

Fourthly, very few organisations are in a position of being able to detect reliably whether their security has been breached or their data compromised, even when all their systems and data are hosted and managed internally. Indeed, in respect of both security breaches and employee misuse of data, it was reported that incidents were typically identified either by chance or on the basis of information received, rather than on the basis of regular and reliable compliance measures. Of course, in practice being able to design, implement and monitor the operation of such control activities is frequently non-trivial and requires competence, resources and careful planning. Whether organisations would be able to apply such controls to outsourced data, or be satisfied by the monitoring and reporting services provided by their cloud service provider, is the same question with an added layer of complexity.

5.2.5.5 Cloud computing as a financially attractive option

Fifthly, the idea of cloud computing as a financially attractive option for outsourcing a number of traditional IT activities including data storage is widespread and there is a great deal of enthusiasm for it across a wide range of organisations, but this enthusiasm is not yet being widely and systematically backed up by detailed risk assessments and careful consideration of the approaches needed to identify, classify, manage and monitor the data being transferred into the cloud.

Although not explicitly mentioned in the questionnaire, several correspondents drew a parallel of sorts with the concept of Bring Your Own Device (BYOD), a concept popular within some communities and organisations that is moving away from the centralized control and provision of devices by the IT department of an organisation towards a model where users are permitted to provide, and be responsible for their own devices, be these computers, smartphones, or peripherals. For these correspondents, there are similarities in philosophy between the loosening of corporate control over the devices that are connected to internal networks and the (perhaps only symbolic) loosening of hands-on control over systems and data that is perceived when moving onto cloud services. Similar cost-based and easy-access arguments are advanced to support the two initiatives.

Of course, BYOD is far from being a model of service fulfilment that is acceptable to every organisation. The objections to this are many: security is frequently raised as the first reason not to adopt such a strategy, given the inability to control the contents or configuration of machines connected to an internal network and the threat of virus infection, for example; a second objection concerns the consistency and compatibility of application software, which is not guaranteed across platforms and versions; a third concerns support, as in-house IT typically finds it far easier to support a standard hardware configuration and a standard software image. It is a not unreasonable generalization to state that organisations that are more concerned with demonstrable security and compliance are less amenable to the application of BYOD to their user bases.

5.2.6 Lessons learned

The comments provided alongside the answers also provided useful information. In particular, several respondents referred to the different types of cloud that are beginning to exist: in certain industries such as financial services it would be unthinkable to use cloud services in which confidential data might be stored outside Switzerland or for which it would be difficult to obtain adequate audit comfort over key concerns, but the use of some kind of industry-specific Swiss cloud, perhaps set up as a joint venture, with appropriate controls and safeguards in place, might be conceivable.

Several respondents also flagged up the importance of the proper management of backup media, which of course need to be subject to the same policies and procedures for data management and security as live datasets. From the perspective of the management who will retain the responsibility for ensuring the availability and reliability of system and data backups for good practice and going concern reasons, and for the auditors who will be verifying this, it will be a challenge to identify exactly which data are backed up where, how this is managed from a security and availability perspective, and what the timelines, sequences and interdependencies would be for restoring part or all of a missing dataset.

Within all businesses there is constant pressure to reduce costs and cloud computing could be seen as an effective method of managing and reducing costs, particularly in the short-term. If there are no significant in-house IT systems, a case will always be made for reducing to a bare minimum, or even eliminating entirely, the IT function, thereby reducing staff costs alongside the operational costs of monitoring and maintaining systems.

The decision to choose cloud computing services is not one to be taken lightly. For individuals, the use of personal services in the cloud (such as Facebook, gmail and Dropbox) is, or rather should be, a matter of a calculated assessment of risks and benefits, for the potential negative impacts of breaches in security, for example, can be significant. For businesses and other organisations, the same concerns apply but on a larger scale. For all organisations that have a responsibility to keep their, and others', data secure and confidential, the decision can only be taken after a detailed analysis of how they will obtain, and continue to obtain, the necessary comfort that this is the case. If they cannot build into their own procedures, into enforceable contract terms, and into audit plans, the means of confirming the confidentiality, integrity and availability of their systems and data, they should not consider externalizing it.

In addition, organisations should not forget the immediate costs and efforts involved in moving onto the cloud. Datastores need to be identified, classified, cleaned up and archived, and serious technical and operational decisions are needed to determine what data will sit where. This will frequently be a project of a significant size requiring expertise, resources and input from a number of people across the business who understand the business, the systems, the data, and their use.

Experience of traditional outsourcing suggests that it is very easy for organisations to overestimate the cost savings generated by a move towards service providers and to underestimate the amount of internal competence and dedicated management required to make a success of such initiatives. As long as the organisation relies on its data and retains responsibility for all aspects of its business from a regulatory perspective, it will need to ensure that its management of the relationship with its service providers and its access to critical operational information are both adequate and appropriate. Typically this will require retaining or recruiting skilled, experienced and reasonably senior staff to liaise with and monitor the performance of the service provider.

The long-term consequences of opting for a cloud solution also need to be examined. Once on the cloud, systems and data are likely to stay there, and feasibly with the same provider. What begins as a simple and cost-effective solution to a small problem could develop into a long-term strategic commitment with little scope for alteration. Such dependence can have an impact on internal control frameworks and quality because when given no other option, clients may find themselves obliged to follow technical and procedural decisions taken by the service provider and thus adapt their internal control frameworks accordingly. Should the service provider not be able to provide the level of comfort necessary, nor the information that the client requires in order to conclude on the effective operation of internal controls, what course of action is open to the client?

5.2.7 Further research issues suggested by the results and the limitations of the survey

It will be necessary to monitor the development of the use of cloud services for data storage from a controls perspective in order to see how such use develops, whether it becomes widespread and whether it does become a factor in assessing the quality of internal controls. Perhaps practical and effective operating procedures will be developed and become standard very rapidly, so that the subject does not become a major concern for controls managers, or perhaps the take up of the technologies will not be great, because of perceived controls issues or questions of cost or access.

It will also be interesting to study the impact of the uptake of such services on audit opinions and compliance reports. This will surely be a major driver in the development and the use of these services: if organisations find themselves subject to adverse comments from regulators or external auditors, that will necessarily cause a slow-down in adoption. On the other hand, if clean reports are issued and no concerns are raised, take up will only be encouraged.

5.2.8 Overall conclusion from the survey

In common with all IT-related projects, initiatives in respect of data collation and accumulation and in respect of data migration and transfer require a great deal of planning, strategic awareness and effective controls in order to ensure that the key security objectives defined by the organisation continue to be met. Both managing unstructured data and managing the migration onto, and ongoing monitoring of, cloud services, present countless opportunities for the loss of the confidentiality, integrity and availability of data, to cite once again just the three most famous security objectives.

The use of large datasets for competitive advantage is highly tempting for many organisations from a tactical perspective, while the use of cloud services is attractive for a number of reasons, financial, operational and strategic. The senior management of organisations tempted by such initiatives should be aware, however, that neither type of project can be successfully completed overnight, and that they should be prepared to provide the necessary resources, guidance, oversight and supervision to ensure that the advantages obtained through the initiatives are not outweighed by decreased security, increased costs, or reduced comfort from internal controls.

5.2.9 Summary

This survey, limited in scope but completed by a representative cross-section of organisations within Switzerland, demonstrates the reality of the understanding and adoption of unstructured data and cloud

technologies within the country. The highest-level concepts are well-known but initiatives to exploit them are still in their infancy, with a great deal of hesitation discernable within organisations. This is partly explained by natural caution, partly by concerns of controls, and partly by questions of budget and priority. The work demonstrates that there is thus a gap between theory and practice, with opportunities recognized but not yet seized in any widespread manner.

This survey and its accompanying discussion as set out above has been considered of sufficient technical interest and originality to form the basis for a paper presented at the WAINA 2013 conference in Barcelona (see Appendix D) and subsequently a chapter in a technical book published by Springer Verlag in 2014 (see Appendix C).

5.3 Cloud Security Alliance/ISACA Survey 2012

The above survey does have the limitation of being geographically limited in scope. Its findings are supported on an international level, however, by the results of a larger survey carried out in the United States.

5.3.1 Background to the survey

In the 2012 Cloud Computing Market Maturity Study Results [76], the Cloud Security Alliance and ISACA jointly present the findings from a survey of 252 participants representing cloud users, providers, consultants and integrators.

5.3.2 Relevance of this survey

This survey does not address the subject of unstructured data specifically, but is very much focused on the security aspects of modern corporate computing, with particular focus on the use of the cloud.

This overlaps particularly with Section B of my survey and then extends it to dig down into the perceptions of the need for, and means of ensuring, security in a corporation, with particular attention paid to the differing perceptions of importance and significance typical of different service lines within an organisation. This again is consistent with my approach.

5.3.3 Key findings and executive summary

The key findings are that “the cloud market has not yet reached a level of maturity that will support this scenario. Instead, the survey participants believe that platform and infrastructure service offerings are still in the infancy stage of maturity, while software service offerings are just emerging from infancy and are in the early stages of market growth. The respondents estimate that it will take approximately three years for cloud platform and infrastructure services to be firmly placed within the growth stage, and at least two years for software services to reach that stage” (p. 6). This is important because “it is within the growth stage of maturity that a clear understanding of what cloud computing is can be established in the market, empowering users to appreciate how it can be leveraged to provide value, what role changes in terms of accountability and responsibility need to evolve, and how enterprises can leverage cloud to benefit from supplier-user relationships” (p. 6).

These findings are important because they are consistent with the findings from the Swiss survey, which were that the cloud platform remains too immature and that, perhaps equally importantly, organisations’

strategies remain too immature to adopt and benefit from the proposed advantages of cloud services. Opportunities are recognised and the potential is seen, but it is recognised that there need to be advances and developments on the parts of both service providers and potential clients in order for mutually beneficial services to be provided and utilised.

It is reported that “Participants in the CSA/ISACA study are optimistic about the future of cloud computing, but they are concerned that viewing cloud purely as a technology rather than as a business issue constrains cloud market maturity” (p. 6). Although technology management understands the value of the cloud, cloud risks are seen as technological issues rather than business issues. The lack of appreciation of the benefits and risks to business from the adoption of cloud technologies is preventing the cloud from transitioning from a technological choice to a business resource. “Participants are confident that cloud is helping enterprises become more virtualized and distributed; reduce IT cost; and optimize the use of IT resources in supporting business units, in particular by bringing new applications to market more quickly. They are less confident that cloud is driving business innovation, increasing customer satisfaction, opening new markets and increasing revenues” (p. 6). These are legitimate and widespread hesitations about using the cloud, although critical questions could be asked about how realistic such expectations are. Considering the four items mentioned above, it is not unreasonable to expect that judiciously deployed and managed cloud services could assist with business innovation, given that changes can be made and new technologies or platforms deployed as rapidly as the service provider can do so. Increased customer satisfaction should also follow from increased stability and improved responsiveness of systems, if those advantages are indeed obtained from moving onto cloud services. But neither opening new markets nor increasing revenues seem to have any direct linkage with a migration onto the cloud, which from an operating perspective is rather more a cost management issue than a revenue generating one.

Key advantages cited are availability, business continuity and disaster recovery. Participants are also “highly confident that cloud service performance, system outages and problem resolution—typical technology issues—are currently being addressed.” By contrast, they are “less optimistic that issues such as provider longevity, an understanding of data ownership and custodian responsibilities, legal issues, contract lock-in, and exit strategies are being addressed to the same extent. They are also concerned that government regulations are not keeping pace with market changes—a significant limitation, given the weight of regulatory compliance on enterprises” (p. 6).

A key finding of the survey concerns the factors reported as critical in making cloud decisions. Cloud users “were asked to identify (from a supplied list) the factors that are most important to their enterprises in making cloud computing decisions. The factors listed included aspects that contribute to business enablement (e.g., elasticity of offerings, reliability and availability, quality of service, and the ability to move more quickly into markets) and financial factors (e.g., cost reduction, the ability to pay per use, opportunities for higher ROI, and the ability to turn capital expenses [CapEX] into operational expenses [OpEX]). Participants were also asked about the importance of reducing their environmental footprint as a factor that influenced their cloud decisions” (p. 10).

Business enablers were viewed as far more important than purely financial considerations in making cloud decisions, with the reduction of environmental footprints far behind. It is interesting that within the business enablers, the factors relating to the reliability and availability and the quality of service were reported to be the most important, demonstrating the importance of these elements to corporate

management and their confidence in the cloud service providers to guarantee an appropriate level of service.

5.3.4 Specific findings of interest

The tables presenting the detailed results relating to each topic discussed in this section can be found in Appendix A.

5.3.4.1 *Cost reduction and business enabling*

From a financial perspective, cost reduction is the key driver behind the adoption of cloud services. Interestingly, the next two highest scores relate to elements that are not purely financial: higher ROI is based on the financial cost of investing, but Quicker Time to Market is very much an operational issue and one that is closely related to business efficiency and flexibility (Figure 2 of the report).

5.3.4.2 *The influence of business groups on innovation*

Another revealing finding emerged from the analysis of the business groups exerting influence on cloud innovation within organisations. It is clear from the results that although the adoption of new approaches to infrastructure and new strategies for supporting the business should ideally come from senior management, it is in reality IT management and staff who are driving change (Figure 7 of the report).

5.3.4.3 *Positive and negative influences on cloud adoption and innovation*

A further key finding emerges from the analysis of the most important positive and negative influences on cloud adoption and innovation.

Influences will vary depending on the perceived maturity of the service or product being offered. “While they are in their infancy, disruptive technologies may move more slowly toward adoption because of the unbalanced weight of perceived negative over positive factors and their impact on expected benefits. With maturity comes greater visibility of the positive aspects of the technology. Potential problems and limitations are minimized, solutions are found, and benefits take on greater importance in decision making” (p. 16).

Among the key motivating factors are the ability to better manage costs, improve customer service, meet internal and external demands, and leverage technology to enter new markets more quickly.

These can be overwhelmed, however, by the factors that discourage take-up. These include security and data ownership (both related to the ability to protect information assets) and factors related to legal issues, contracts and regulatory compliance, while the fifth-ranked item on the list, information assurance, is also very significant because it reflects management’s need to obtain comfort over the overall security of data, a task often made difficult by the potential lack of transparency over control activities performed in respect of cloud services (Table 4 of the report).

5.3.4.4 *Negative influences: users and providers*

When looking in detail at the perceived negative influences on cloud adoption and innovation, the survey revealed interesting differences between the viewpoints of cloud users and cloud providers. “Users evidence a more cautionary perspective than providers, as illustrated by a comparison of the mean scores.

Although the scores of the users (overall mean score 3.76) and the providers (overall mean score 3.51) are not widely divergent, they do reflect a difference in perspective. Most notable in this regard is the result relating to disaster recovery and business continuity; users reported a more negative perspective, ranking it as the eighth most limiting factor as compared to the providers, who ranked it tenth. Users also saw contract lock-in as being a more limiting factor. Not surprisingly, providers saw this as a less significant negative influence on cloud adoption” (p. 17). These differences between users and providers are interesting because they demonstrate the different impact on each of any problems that might arise: the users depend absolutely on effective contingency planning in order to continue their operations, while for the service providers it is less of an immediate concern. Clearly providing a professional and reliable service is part of their core objectives and the basis of ongoing and future contracts, but problems with continuity procedures will have less immediate impact upon them (Table 6 of the report).

5.3.4.5 Risk components

In respect of risk, the report states that “there is not so much agreement that cloud risk should be addressed as a business risk and even less that cloud should be addressed as part of the enterprise risk program” (p. 21). This is an unexpected finding and from an audit and controls perspective is a clear weakness in corporate behaviour because treating cloud risk solely as a technology risk means that the wider value that cloud activities can bring to an organisation may be missed and opportunities for real innovation may be spurned. Across organisations there is a need for flexible thinking about the value the cloud can bring and what the risks will be in relation to introducing new methods of managing and using information. As a general point, senior management need to ensure that they have a thorough understanding of the cloud, can articulate a vision of the benefits it can bring, and communicate the key messages about risk down into business units and technology functions (Table 11 of the report).

5.3.4.6 Roles and responsibilities

A key element of internal controls is the clear definition and acceptance of roles and responsibilities for the performance and monitoring of controls. Without clarity and effective operation, confusion can occur and controls be weakened. The survey indicated, however, that this remains a perceived area of weakness in the cloud environment, with confusion between suppliers and user and within user organisations (Table 12 of the report).

5.3.4.7 Service level agreements and user requirements

As discussed in Chapter 4, service level agreements are an essential part of the management and control frameworks. Without clearly defined agreements on the nature, quality and quantity of services to be provided, the user organisation cannot be certain that its needs are being met, at least in theory. In addition, it is necessary for there to be regular and thorough monitoring of adherence to the service level agreements, so that errors or omissions can be identified and rectifications undertaken, and so that the users do not find themselves exposed to unnecessary risk.

The results of the survey demonstrate that these questions remain, however, a source of real weakness for both users and suppliers. In respect of service level agreements themselves, there is slightly greater confidence that service level requirements are clearly defined than that they are properly incorporated into contracts, which would be consistent with the experience of CIOs and IT organisations in dealing with technical specifications with third parties (Table 13 of the report).

5.3.4.8 *Performance monitoring*

Of even greater concern are the results in respect of the monitoring of performance against service level agreements. Both parties to the contract need to carry this out: the service provider to ensure that they are operating effectively and not exposing themselves to financial risk for breach of contract; and the users to ensure that they are aware of any potential problems. The survey indicates that this is not performed with any consistency, however (Table 14 of the report).

5.3.4.9 *Cloud computing for the resolution of problems*

A key motivation for adopting a strategy involving outsourcing or the use of cloud services is the expectation that it will improve the quality of service in respect of the resolution of problems. The survey enquired about this aspect of cloud services. “A series of 27 problem statements was presented to the 252 survey takers, who were asked to rank each statement on a scale of 0 to 5 according to how confident they were that the problem described is currently being resolved. The overall level of confidence in problem resolution was reflected in a mean score of only 2.38. Of the three components that came together in the confidence barometer—service, strategy and confidence in problem resolution, all were equally low. Only strategy confidence—an expression of aspirations related to strategy, alignment, customer value and opportunity realization— received higher confidence ratings” (p. 26)

This is revealing because of the insight it gives into the realities of cloud service provision. Users are tempted by cloud services for a number of reasons, but clearly there is low confidence that the use of cloud services will actually provide effective solutions to a number of challenges that the organisations face. Users believe that cloud services and cloud service providers will assist them in accomplishing their overall objectives and in responding to customer requirements and new opportunities, but there is little confidence that cloud services will help them in addressing the more prosaic challenges posed by routine operations.

The report provides further analysis of these responses. “The 27 problem statements were grouped into eight separate categories for presentation to survey participants. Each category combined two or more related measures providing greater insight into specific problems often associated with cloud computing. Of the eight categories, not a single grouping of cloud-related problems received a score that would indicate strong confidence that problems are being resolved. None received a score indicating even a moderate level of confidence that problems are being solved. For the eight categories the scores ranged from 2.89 for cloud continuity and availability—the highest rating—to the lowest mean score of 2.06 for regulatory and legislative problems” (p. 26).

These results are rather damning in respect of both the level of expectations held by users in relation to their cloud services, and of the attitudes in place among management that mean that services are being contracted for and paid for with such low expectations. It would be revealing to be able to put these results into a more specific context and understand the motivations behind the use of these services. It could be speculated, in spite of the insistence upon the importance of business enablers in the formal responses, that financial motivations are a significant driver, with organisations freeing up capital and committing to fixed budgets for the provision of IT services by outsourcing IT activities. A more worrying explanation, or at least partial explanation, would be management do not believe that in-house IT services would be capable of providing services of any greater quality or flexibility.

The ranking of these results should not come as any surprise. The major motivation for outsourcing infrastructure and processing has historically been to shift the burden of ensuring the continuity of service and a high quality of service onto an organisation that is dedicated to such matters and has the resources, infrastructure and competence to provide a high quality of service. Given that this is the core activity of cloud service providers, it is unsurprising that users have relatively high confidence that these areas are being addressed (Table 23 of the report).

5.3.4.10 Continuity and availability

This confidence is only high in relation to other factors, however, and can be further broken down to demonstrate a difference in confidence levels between the issues of availability, disaster recovery and business continuity (here grouped together although strictly speaking related but separate activities), and resolving system outages. It is clear that users have more confidence in the progress being made towards ensuring constant availability than in the progress being made towards resolving problems that do occur (Table 24 of the report).

5.3.4.11 Service performance

The detailed breakdown in respect of performance also throws up some interesting results. There is reasonable confidence that service performance issues will be addressed, but there is a striking lack of confidence that measures are effectively being taken to address issues related to performance metrics and standards and the availability of expert support. This latter point is particularly striking given that the provision of expert support is a key selling point for service providers, especially in respect of infrastructures or technologies where there is little expertise available in the marketplace or where the costs of maintaining an internal competence centre are dissuasive for all but the largest organisations (Table 25 of the report).

5.3.4.12 Security and assurance

Issues related to the security of data and assurance over such security are key concerns in respect of cloud computing. The survey reveals that users continue to have grave concerns over how several aspects of this domain are being addressed, especially those related to data ownership and custodianship and international data privacy. (Table 32 of the report)

When viewed in the specific context of data security, structured and unstructured, this demonstrates that cloud services providers still have a great deal of work to do to provide services of a sufficient level of security to satisfy their users.

5.3.4.13 Security and assurance – differences between providers and users

It is informative to drill further down into the data and differentiate between the perceptions of users and service providers, and those of different disciplines within the users.

When comparing the confidence levels of users and providers, there is consistency in the rankings (out of the 27 problems under consideration) for each of the five security and assurance components except for the specific component related to information security. There users have more confidence than providers that security issues are being addressed, which might be due to greater knowledge of those within the data centre of how robust procedures are and how well they are followed (Table 33 of the report).

5.3.4.14 Security and assurance – differences between disciplines within user organisations

A similar difference in rankings can be identified between disciplines within user organisations. “In every one of the security and assurance items, except for testing and assurance, security personnel are less confident that issues are currently being addressed. The difference in ranking among business, information security and IT participants is especially striking for the concerns for multi-tenancy item. The information security personnel who participated in the study displayed considerably less confidence than business and IT professionals that those problems are being solved” (p. 30). It can be argued that this is due to the relative closeness of the different units to each of the activities. Security will usually be far closer than the other units to the testing and assurance processes related to the cloud, and will accordingly have a better understanding of the progress actually being made. Technology staff will have a slightly different perspective, understanding multi-tenancy and more business-oriented issues, while the business units are likely to have little involvement in technical discussions and decisions and perhaps also be affected by other sources of comment and information (Table 34 of the report).

5.3.4.15 Regulation and legislation

The area where the least confidence in progress is expressed is that of regulation and legislation. This is a critical issue for users that need to meet regulatory requirements, a facet of data management that can incur significant costs and be subject to frequent and onerous change. Interestingly, users have low confidence in their ability (and that of their service providers) to make progress in respect of attaining and maintaining compliance, and have even less confidence in government regulations keeping pace with the development of technology and adopting approaches that enable the wider adoption of cloud services rather than limiting them (Table 38 of the report).

5.3.5 Overall conclusions from the survey

Overall this survey demonstrates that the cloud market remains a long way from maturity and that there are many key managerial and operational aspects of the use of cloud technologies that are causing users to hesitate. Without the confidence that progress will be made in addressing the weaknesses that are commonly recognised, progress towards more widespread take-up of cloud services and the provision of mature, robust and responsive services will be slow.

Such confidence will require action from all parties. Service providers will need to understand the precise causes of hesitation on the part of their existing and potential customers and demonstrate that measures have been put in place to address those concerns in a structured, consistent and demonstrable way. At the same time, clients will need to carefully define their objectives and requirements and identify what kind and level of comfort they will require in order to feel reassured that their concerns have been met. Overall this is a domain that will require clear and open dialogue between all parties.

It is revealing to note the differences in attitudes between service providers and users, with users demonstrating a clear level of scepticism about the benefits to be obtained from using cloud services. This should not come as any real surprise, as it is essential for service providers to sell their services and be confident about both the benefits to be gained and the ease of overcoming any problems encountered; at the same time, it is clear that the service providers will need to find convincing and credible means of informing clients of the benefits of their services.

It is also revealing to consider the difference in attitudes within organisations, with technical and security groups having different perceptions of, and confidence in, the ability of cloud service providers to address issues of security and reliability than their colleagues in business units.

This can be explained, perhaps a little cynically, by the ability of service providers to convince a non-technical audience of the quality of their offerings, and by the lack of detailed appreciation of a non-technical audience of sophisticated questions related to security and reliability. But regardless of the reasons for such a difference in attitudes, this finding underlines the need for any cloud project to seek input and opinions from a range of people within the client organisation, including business unit leaders, finance managers, technical staff familiar with the technologies and their applications, and legal staff. Only in this way can measured, informed and responsible decisions be taken.

5.4 Comparison of the two surveys and their results

A number of common themes can be identified in the results of the two surveys, even allowing for their different scopes, sizes and objectives.

The major findings as discussed above can be summarized in the following table.

Swiss Survey	ISACA Survey
Importance of having a structured and documented approach to data management and security	Cost reduction and business enabling
Awareness of unstructured data, benefits and risks	Influence of business groups on innovation
Weakness of security monitoring in respect of data access	Positive and negative influences on cloud adoption and innovation
Difficulty of recognizing or being made aware of security breaches	Negative influences as viewed by users and providers
Absence of detailed risk assessments and maturity analyses in respect of the cloud	Risk components
Requirement to subject all elements of data storage to proper management	Roles and responsibilities
Pressure to reduce costs as a driver towards using the cloud and reducing in-house IT services	Service level agreements and user requirements
Need for full consideration of risks and benefits before adopting a cloud solution	Performance monitoring
The need to full anticipate the costs, direct and indirect, of transfers	Cloud computing for the resolution of problems
The possibility of being locked in to a single provider	Continuity and availability
	Service performance

Swiss Survey	ISACA Survey
	Security and assurance
	Regulation and legislation

Table 5.3 Major findings of the two surveys

The first and most significant point is the lack of maturity that can be demonstrated, both within client organisations and service providers. Organisations have been slow to develop objectives, strategies and policies for preparing systems and data and getting into a state of readiness for a migration onto cloud services. Similarly service providers have been slow in providing the kind of robust frameworks that demanding users wish to see to be certain that their systems and data will be appropriately secured and reliable when moved beyond their direct control. That there has been so much activity in this domain, with providers offering services and customers contracting for these services, shows that there is a great deal of optimism about cloud services being a significant and reliable aspect of the enterprise computing landscape in the medium and long terms, even if from an audit perspective the absence of high-level control objectives is a concern.

Both surveys raise questions about the ability of cloud technologies to provide solutions to a range of IT or management problems. There will always be a tendency for the providers of such services to present their products in the most favourable and flattering light, but potential clients, and particularly the more technically-aware groups within those potential clients, are at the moment unconvinced by the descriptions of the services offered and the capacity to modify and adapt them to respond to changing requirements.

Concerns over security remain paramount, with different disciplines within organisations having different perceptions of the changes that moving to a cloud provider will have on the quality of the security around their systems and data. Security specialists are the most sceptical, understandably, with technical staff close behind. Business users are the most confident about the quality of the security provided by service providers: this may stem from a lack of detailed knowledge of the types of measures in place and how these might differ in design and operation between in-house security and service provider security; it may also stem from a lack of understanding of the specific and technical risks involved in moving systems and data outside the organisation’s own security perimeter. In particular, concerns are raised about the assurance of being aware of or informed about security breaches in an appropriate timeframe; when activities are performed at arm’s length or by a proxy, management lose an element of hands-on operational control that is particularly valued in the field of security management.

The question of costs, savings and efficiencies is inevitably a key factor reported in both surveys. There is an underlying feeling reported in the comments to the Swiss survey that the quoted costs of cloud services do not necessarily present an accurate picture of the overall costs. To the amounts to be paid to the service provider for the baseline service have to be added internal costs for monitoring of service and contract management, as well as the costs incurred in reperforming risk analyses and designing and implementing updated internal control procedures.

5.5 Overall summary of the surveys and their results

The results of these surveys demonstrate that a move onto cloud services is neither straightforward nor a panacea. This is largely recognised by cloud client organisations, who in general retain a healthy scepticism over the potential benefits of such a move, but at the same time the Swiss survey indicated that, among smaller organisations at least, there has been little internal control-focused activity driven by the realization that detailed planning and risk management are required.

Given the popularity of the use of cloud services, at least as a short- or medium-term objective for many organisations, it is therefore essential that the risks and benefits are clearly identified and understood by the senior management of organisations considering such a move. A key part of this will involve the analysis of options and impacts by appropriate specialists, and the effective communication of their findings to relevant members of senior management.

These findings are referred to in the analysis of the trends, constraints and limitations related to the use of unstructured data and cloud services presented in Chapters 6 to 9, and also form part of the framework presented in Chapter 10.

Chapter 6 Identification of Control Constraints: Management Theory and Security

6.1 Introduction

“If you can't control your peanut butter, you can't expect to control your life.”⁷

It is clear that while there are many advantages and benefits to be drawn from the use of new technologies including cloud services, there are a number of constraints and stumbling blocks to be considered by prudent and compliance-minded management, especially where multiple datasets containing sensitive and unstructured data are being employed.

Over the next four chapters I will discuss the key issues that are generally applicable, before focusing on the specifically Swiss situation. This approach is necessary because there are a number of general principles that apply globally, and which are systematically discussed in the largely United States-focused literature, but there are also specificities to the Swiss context of which businesses and organisations operating in Switzerland need to be aware.

The issues can conveniently be broken down into theoretical and security-based, taxonomies and models, audit and compliance, and technological elements; each topic is the subject of a separate chapter.

In this chapter I will discuss generic, high-level risks and requirements in respect of the adoption of cloud services and the transfer of data to a service provider, consider the key aspects of information security, and discuss the concepts of privacy and of the lifecycle of data. These topics are overwhelmingly important in their own right as fundamentals of internal control and corporate governance, while they are specifically important to the type of data management and migration project discussed throughout this thesis.

6.2 Generic risks and requirements

Mather et al [67] distinguish a number of barriers to adopting the cloud that need to be considered and evaluated so that appropriate decisions can be taken. Their list is reasonably complete; and can be applied more generally to the wider adoption of new technologies or practices; these are generic risks that affect the ensemble of operational IT activities. The key elements are commented below.

6.2.1 Security

The first of these barriers concerns *security* in its widest sense. The use of cloud technologies has security implications at many levels, including network, server, data transmission, backup, database and application

⁷ Bill Watterson, “The Authoritative Calvin And Hobbes”

level, and it needs to be demonstrated that even the weakest link in the chain is sufficiently strong. Management, users and those to whom data related will all be concerned about the quality of the security measures.

6.2.1.1 *Privacy*

It remains to be seen whether sufficient controls can be implemented around clouds computing to address widespread concerns about *privacy*. Legislation and restrictions have become increasingly strict in respect of the privacy of individuals' information, and with the degree of distance provided by any kind of outsourced services, coupled with the lack of hands-on control by data owners, there is likely to be an increasing battle to obtain the comfort necessary to conclude that privacy is indeed being respected.

6.2.1.2 *Connectivity*

As soon as any kind of offsite processing or data storage solutions are considered, the question of *connectivity* and consistent access takes on significant proportions. Remembering that one of the three standard security criteria is that of availability, organisations wishing to use offsite services and intent on achieving the security objectives need to be confident that connectivity concerns have been addressed.

From the perspective of individuals, too, connectivity is becoming an ever more important factor of personal computing. The evolution of personal computers and other devices has seen a movement from small and expensive disk storage that had to be carefully managed (the author's first home PC, purchased in 1994, came with 4 megabytes of RAM and a then enormous 170 megabyte hard disk), to the situation in 2014 where 750 gigabytes or 1 terabyte of disk space is not at all unusual in a consumer machine and smartphones commonly offer 64 gigabytes of fixed storage. This increased capacity, together with the development of encoding technologies permitting the easy and compressed storage of audiovisual and multimedia files, allowed users to encode and store huge quantities of music and video on their PCs. The arrival of both the cloud and widespread high-speed broadband connectivity, however, are pulling users in yet another direction. Users of such services as Apple's iTunes and Amazon's music services are encouraged to keep their music in the cloud so that it is accessible from any device, while hardware manufacturers such as Acer and Samsung who produce the Chromebook range of laptops for Google offer machines that typically have very small (by today's standards) fixed storage capacity (often a 16 gigabyte Solid State Drive (SSD), chosen for its rapid performance compared with traditional mechanical drives), accompanied by offers of 100 gigabytes of cloud storage. The argument for the consumer is that there is minimal boot-up time, minimal maintenance, and effortless integration of the most popular applications with the cloud services and the pre-installed "Chrome OS" operating system. What is necessary, of course, is that in order to access all of the cloud data, users will need a suitably rapid and reliable Internet connection, both domestically, which is often assured, and while travelling, which is less reliable.

6.2.1.3 *Reliability*

A further concern is *reliability*. Very often enterprise platforms and applications are critical to the business and thus their performance must meet certain reliability criteria. Service providers must be able to demonstrate a solid track record of reliability in order to attract and retain customers, while both service providers and customers will need to design and implement, on the basis of a measured risk assessment, appropriate backup and recovery plans. The relative differential cost of contingency planning compared to the costs of outsourcing or retaining services in-house can be used as a straightforward indicator of the

attractiveness of any outsourcing or cloud offer: the more the organisation feels it needs to spend on contingency planning, the less likely it will be to accept an offer of service provision.

It is important in this context to consider the specific requirements and expectations that an organisation might have in respect of using cloud services, and also to underline the various levels of responsibility held by the service provider. Not every organisation will require, say, 99.99% uptime for all of its systems, and given that the higher the level of reliability sought the higher the cost will be, it would be uneconomic to demand such quality of service. Management should therefore perform a realistic and detailed analysis of their requirements in order to rank applications and platforms in order of criticality and to attribute to them a realistic idea of their required reliability. To take two extreme cases, there will clearly be a difference in requirements between an online ordering system upon which a business depends for all of its incoming orders and hence revenue, and say a production data archive system they may be consulted perhaps once a week for the purposes of statistical analysis.

It should also be recognized that service providers, no matter how well managed and well resourced, cannot ensure an absolute continuity of service. At the smallest, most mechanical scale, all hardware components can fail, and large data centres replace hard disks, network cards and motherboards on a daily basis. Data centres are not immune from the loss of network connectivity, either internally or externally, nor can they guarantee a faultless electrical supply. Then there are larger scale external factors such as fires, floods, earthquakes and tornadoes, to take a few as examples, which cannot be prevented. What is essential is that the service providers have appropriate measures in place to limit the lack of service, using preventive measures such as careful monitoring of performance in order to identify deteriorations in hardware and perform preventive maintenance, using monitoring systems to detect failures or anomalies and activate the necessary resources to investigate and remediate the situation, and taking preventive strategic decisions to build in redundancy to connectivity and power supply installations, or provide backup network connections and local power generators, for example, in order to be able to respond quickly. Service providers need to be able to demonstrate that outages will be as infrequent as they can manage and will be properly managed; customers need to confirm that the level of service proposed is sufficient for their needs.

6.2.1.4 *Interoperability*

Given the wide variations in systems, platforms, processes and formats that permeate the technical world, another key consideration is that of *interoperability*. It is essential for the adoption and development of cloud services that the interoperability and portability of information between private clouds and public clouds and between standard cloud facilities and existing industry-wide standards used by potential clients can be guaranteed. In some respects this can be seen as a continuation of two initiatives affecting two different levels of technologies: it required a great deal of work to update standards and create true interoperability between the ways that files were structured and handled on different platforms, while interoperability between different character coding conventions and character sets can still be problematic (consider, for example, the difficulties still encountered today with the handling of accented characters in some train and airline booking systems); more recently, many large organisations have invested heavily in medium- and long-term projects to standardize their data, processes, platforms and systems through the use of a single, all-encompassing Enterprise Resource Planning (ERP) system such as SAP or Oracle Financials. Such projects are typically costly and of long duration, in my experience of large manufacturing companies often running into the millions of dollars over two or three years, and require a great deal of

planning, investment, process re-design, testing and user training. That they are undertaken with such frequency shows how highly the expected improvements and benefits are valued.

Even a decade ago it was not unusual to see multinational manufacturing companies, for example, using a variety of applications sitting on different databases and operating systems (particularly in production environments it was common to see OS/400, flavours of UNIX, versions of Windows, and Novell networks all in use), with a resulting maze of interfaces, automated, semi-automated and manual, transferring data from one system to another, often performing conversions of format or structure on the way, and controlled – if controlled – by a complex sequence of reviews, reconciliations, analysis of exception reports, and user oversight. All of these activities were necessary to ensure the creation and maintenance of consistent and useful data, both master and transactional, and to be able to generate reliable and consistent consolidated information.

Such projects are complicated to manage internally when the organisation has full control over the platforms and datasets in use. In a cloud environment, it will be even more difficult to ensure that the business changes as necessary to reflect changes in the technology or to address the opposite case, where changes to the business make changes to the IT infrastructure necessary. Organisations will need to be confident that their service providers are capable of the flexibility and speed of change to be able to respond to such requirements.

6.2.1.5 *Independence from providers*

A further key consideration is the fundamental principle of retaining *independence from cloud service providers* (CSPs), a topic mentioned in Chapters 4 and 5. The world of traditional outsourcing contains stories, often anecdotal because the organisations at the receiving end do not wish to publicise the nature of their difficulties, of organisations having decided to outsource part or all of their activities, IT or business-related, and discovering a little later that it was at best very difficult, and at worst impossible, to bring the activities back in house or to transfer them to another service provider. Often the skills needed to operate services in-house would no longer be widely available in the market place; this is typically seen in the domains of mainframe operations where the technology has been superseded in many environments and there is no new generation of experts. In other cases the technologies themselves are no longer offered by service providers. The two solutions that exist are then to continue with the existing service provider, which may be costly and not necessarily a long-term strategy, or to scope, select and implement alternative platforms and applications, using more modern and more widely-supported technologies. This, of course, can be more easily said than done, especially in environments such as transactional banking systems or operational airline systems, where any changeover needs to be very carefully managed and there needs to be a guarantee of interruption-free system from the replacement systems.

The problems can also occur on a less fundamental scale, for example if a service provider is incapable of providing the necessary rapidity in evolution of offerings and support. There will not always be a convenient alternative service provider to which activities can be switched, or perhaps there will be compatibility problems with the way data are stored or processed, or maybe there will be financial penalties.

As a general recommendation, whatever the context, organisations should be well aware of the nature of the problems that could arise in respect of being too closely dependent on the services provided by one individual supplier.

6.2.1.6 *Economic value*

Management also have the over-riding responsibility to consider the *economic value* of the strategic decisions they take. There is an obvious argument to be made that using cloud services will be financially beneficial, paying only for the services that are used and reducing the upfront capital investment involved in hosting IT solutions, as well as reducing the in-house headcount and associated costs. This argument does need, however, to take into account all the costs and benefits associated with the cloud, short- and long-term, obvious and hidden. Economies of scale may only begin to make a difference at a certain level of commitment and this may not correspond to the organisation's requirements. Costs that might not be immediately apparent when considering cloud services include ongoing support, both internal and external if this is deemed necessary, platform and application modification, disaster recovery, and data loss insurance, while careful attention needs to be paid to the details of agreements for payment based on usage: it is an old truism that data will expand to fill the space available to it, and management will need to be very well informed about both current requirements and the prospects for future growth and development of data services and processing.

6.2.1.7 *IT Governance*

While considering management responsibilities, it is essential not to ignore the fundamental requirement for effective *IT Governance*. This covers a wide range of activities, and has been defined by the ITGI as "the responsibility of the Board of Directors and executive management. It is an integral part of enterprise governance and consists of the leadership and organisational structures and processes that ensure that the organisation's IT sustains and extends the organisation's strategy and objectives" [77] and by van Grembergen as "the organisational capacity exercised by the Board, executive management and IT management to control the formulation and implementation of IT strategy and in this way ensure the fusion of business and IT" [78]. The key concept is that IT activities should be aligned with business objectives, including strategies for business direction and development and financial planning. Economic value as discussed above is a key part of this, but management also needs to consider the short-term and long-term impacts of the decisions they take. Moving to cloud services may change the financial position in the short-term or palliate an absence of internal technical resources, but it needs to be ensured that these are appropriate and supportable decisions in the long-term.

6.2.1.8 *Changes in the IT Organisation*

Management also need to be aware of the *changes in the IT organisation* that will be provoked by a move towards cloud computing. There are two main fields in which such changes will be noticed. Firstly, there will be the acquisition of new skill-sets, both technical and managerial, necessary to deploy the technology and to exploit it efficiently. This will often be accompanied by a loss of some existing skill-sets that are no longer required in-house. Secondly, there will be the effects of changing technologies on the role of IT. Where design and development in the service of solving problems were once a purely IT responsibility, roles have already changed with users doing more analysis and desktop development, and the adoption of cloud technologies and the exploitation of large datastores will only contribute further to that trend. The role of IT will logically move more towards facilitating and providing the necessary level of internal controls to address the risks to the business that application development presents.

6.2.1.9 Political issues

The final area of concern in this context is that of *political issues due to global boundaries*. In traditional outsourcing, clients knew which data centre belonging to their service provider would be the main host of their systems and data and could reliably know where their data was being stored for production and backup and also where it was transiting. This is not the case in the cloud computing world, where there is potentially great variability in terms of where the physical data resides, where processing takes place, and from where the data are accessed. As a result of this variability, data may be subject to different privacy rules and regulations, rules and regulations that are perpetually affected by multi-jurisdictional political considerations.

These considerations can have clear impacts on the political stability and cooperation needed for the cloud to survive. When a state enacts legislation in respect of the security of and access to data, this can have an impact on the use and development of the cloud because other states, or large organisations, might choose to avoid any possibility of their data being stored or processed on machines in that territory and therefore subject to local law. An example of this would be the USA Patriot Act, the introduction of which persuaded Canada to cease using computers in the global networks that are operating on US soil, purely to protect the privacy and confidentiality of Canadian data. Such concerns about access to and manipulation of data are increasingly important as more and more potentially sensitive or confidential data are leaving in-house data stores.

6.2.1.10 Criminality

One other issue that should be included in this discussion, although rarely mentioned in the controls-based literature, is that of *criminality*. Much of the discussion of risks and responses is based around the prevention and detection of errors and omissions, of faults in processing and problems with the integrity and accuracy of information flows. This disregards a clearly large, although to date difficult to quantify, source of problems for organisations that are storing and processing data.

Reliable statistics in respect of cybercrime are difficult to obtain. There are several reasons for this, each significant in its own right but which also combine to make the world of cybercrime largely hidden and immeasurable. Prominent among these factors are:

Reason	Description
Definition of cybercrime	There are no standard international definitions of cybercrime. Divergences appear for two main reasons: differences over whether a particular criminal act should be considered as cybercrime or not (often hinging on whether the target or the means of the crime is more significant); and differences over whether an act is in fact criminal. Such differences have given rise to so-called digital paradises in which certain activities can be performed or hosted because legal in that jurisdiction while illegal elsewhere. Typically this applies to activities such as hosting files for downloading or providing gambling facilities, but can also apply to activities such as hosting websites dedicated to subjects such as Holocaust denial or certain kinds of pornography that are illegal in some countries but not in others.
Reporting of cybercrime	The subject of cybercrime is so sensitive for many individuals and organisations that there is an understandable unwillingness on their part to report having fallen victim. This can be in order to preserve their credibility, in the case of an organisation that is

Reason	Description
	supposed, or claims, to have sophisticated security measures in place; to preserve analyst and market confidence by refusing to make public financial losses caused by cybercrime; out of a sense of shame on the part of individuals for having fallen victim; or for fear of further attacks.
Recognition of cybercrime	One specific facet of cybercrime with comparison to traditional crime is that its effects, and even the fact of having been committed, can pass unnoticed by the victim. Financial crime such as the shaving of small sums from bank accounts or the manipulation of invoice details may not be recognised if small enough to slip through reconciliations and approval processes. Other activities such as the pirating of accounts to send spam or to gain access to systems may again pass unnoticed unless there is targeted and timely review of logs and system information.

Table 6.1 Factors inhibiting the analysis of cybercrime

Criminality in respect of systems and data can be targeted or random, depending on the skills, motivations and objectives of the criminals. The cloud context presents further complications when identifying the objectives of attackers unless there are clear reasons to believe that the attackers were specifically targeting the systems and data of one particular client of the service provider rather than attacking the service provider as a whole and looking to see what data might be vulnerable.

In addition, criminal activities might be undertaken inside or outside the organisation holding the data, meaning that different protective measures need to be taken in response to the two different risk scenarios. Protecting systems from external attacks will typically be based upon perimeter controls such as firewalls and strongly restricted opportunities for remote access, while protection from internal attacks will depend upon clearly defined access rights and the monitoring of staff activities. This is of course made more complicated by the need for some technical users to have wide ranging access rights in order to perform systems and data administration, but well-designed and effective security structures should take this into account and address these risks appropriately.

6.2.1.11 Summary

The eleven generic risks described above show clearly the range of issues of which management needs to be aware when considering or scoping changes in the operating environment or in practices. Before any deeper consideration of specific and technical details is performed, these provide a broad structure for management to use when identifying areas of concern and determining whether a proposed project is in principle feasible.

6.3 Data security

Of the eleven generic risks discussed above, the one that appears first on the list and typically attracts the most attention when discussing data and cloud technologies is that of security. This is partly as a result of its visibility and attractiveness as a subject, and partly because it is a subject that conceptually, at least in respect of its basic principles, is easy to grasp. The commonly employed security measures are also easy to envision and imagine in a proposed internal control structure.

In this section I will set out the overall principles of the management of system and data security, identifying the key constraints and limiting factors confronting security managers, and discussing why apparently straightforward concepts can be difficult to design and implement in practice.

6.3.1 Overall principles

Data security has always been a fundamental consideration for business and IT management. When transferring data outside the traditional secure perimeter, the topic remains crucial but becomes even more difficult to manage because of the lack of visibility over usage and access and the increased complexity of evaluating all the different risk scenarios and defining appropriate actions in response.

Data security needs to be considered at a number of levels and from a number of perspectives, including data in transit and at rest, data being processed, data consistency, data storage and data lifecycle management. Clearly not every aspect will be of equal and high importance for every organisation, but management need to evaluate their risk in relation to each facet in order to be sure that no significant risks have been forgotten.

Data security is very much a domain in which the principle of the weakest link applies and the challenge for organisations is to bring all aspects of security up to the same, appropriate, level. When considering the specific challenges related to outsourced or cloud-hosted data, it is important to understand that additional levels of control in place over aspects of security that the service provider can reasonably be expected to perform better than the client, given resources, experience and expertise, will not necessarily mitigate or compensate for additional risks that have been created by the transfer of services outside the client's traditional security perimeter. As an example, consider encryption. It is standard practice nowadays for sensitive data to be encrypted in storage and also, in many circumstances, in transit. But any use of those data needs to be in decrypted form and therefore clients of cloud services can be certain that at some point their data will be unencrypted, somewhere in the cloud. This is particularly prevalent where PaaS or SaaS facilities are being used.

Addressing these risks is difficult. The extreme position is that the only certain means of avoiding data leakage arising from such circumstances is simply to avoid putting sensitive or regulated data onto the public cloud at all, or at the very least only used the cloud for the storage of already-encrypted data. This is undoubtedly true, but is an extreme option that rather ignores the whole motivation for using the cloud at all. The real challenge is in developing means for using new technologies and technological opportunities without being opened up to excessive additional risks.

This is a delicate balancing act, especially in an evolving environment in which demands are increasingly moving towards greater remote access to data and transactional facilities. In the banking industry, for example, technologies have advanced so rapidly that the younger generation of clients is already showing a tendency to have forsaken PC-based Internet banking – itself a technology not yet two decades old – in favour of smartphone-based access. It is simply not an option for mainstream financial institutions to refuse to follow these demands, as the industry is competitive and customer loyalty cannot be assumed. Thus banks are being obliged as a customer service and customer retention strategy to address the issues of smartphone security and wireless network security in order to provide the services being requested.

6.3.2 Identity and Access Management (IAM)

The most fundamental challenge facing data owners and managers concerns controlling access to systems and data. This consists of knowing who should have access to data, who actually does have access, what kind of access they have, what they are doing with their access rights, and whether the access credentials that have been created are being used by the appropriate people.

In the cloud environment this situation is made additionally complicated by the need to manage both internal and external users. Typically there can be three groups of users requiring access to data: the data owners within the organisation; the technical staff working for the cloud provider; and legitimate standard users of the data, who might be connected to the owning organisation or be completely external. Keeping track of all these users is a critical and complicated task, especially in contexts where staff turnover is high (which can be the case of businesses with seasonal staffing patterns or in businesses undergoing significant structural or operational change) and IT and data management are expected to respond speedily to access creation or change requests.

6.3.2.1 Access policies

It is an accepted part of good IT management processes to have a clearly defined security policy that should include details of access requirements and practices. Good practice requires that a number of subprocesses should be documented:

- The process for determining which access rights should be associated with each type of user or role;
- The process for requesting and approving the creation of new user accounts;
- The process for modifying user account parameters and permissions;
- The process for deleting user accounts and removing permissions;
- The process for identifying the potentially anomalous use of user accounts;
- The process for monitoring whether the above processes are indeed being followed.

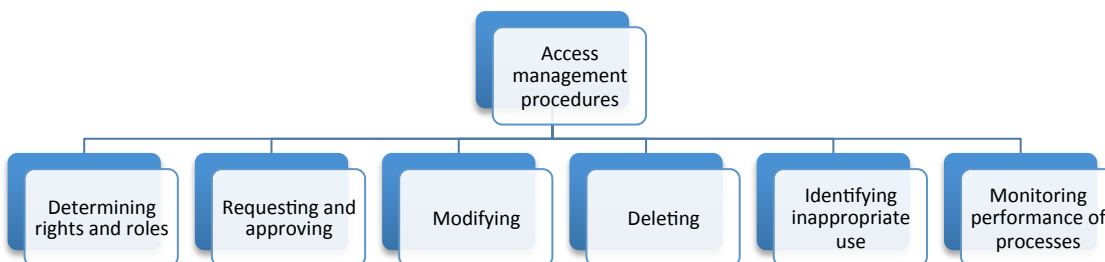


Figure 6.1 Access management sub-processes

Each of these subprocesses is on the face of it simple and straightforward, but each forms a necessary stage of the access management process. The process should start with the definition of which access rights are appropriate for each user or type of user. Very often even this initial step is neglected or badly performed, users tending to acquire access rights by accumulation without due consideration of what they really require for their daily roles. There can be pushback from users to any restrictions that are imposed, especially if such a scoping exercise is performed after the event with a view to imposing order on existing systems; this can be countered by demonstrating that it is in the users' own interests to have restricted rights, because it can reduce the risk of catastrophic error, because it enforces the adherence to internal controls policies and operating procedures, and because it ensures that accountability is maintained.

In my experience such access rights review and renewal is a worthwhile and valuable exercise for organisations that recognise that access to their systems has drifted out of control. In common with many IT-related projects, however, it is very much a management process rather than a technical one, and its success or failure will depend on the quality of leadership, communication and follow-up that management can apply.

Access requests should be formalized, documented, subject to a series of authorisations that the organisation deems appropriate (commonly a request will be originated by a department head and approved by someone more senior, but if the access rights requested include access to particularly sensitive or confidential data then additional approvals may be desirable from, for example, managers in Internal Controls or Security functions), and subject to follow-up to ensure that the account created and the rights granted do indeed correspond to the request.

Procedures for modifying accounts are frequently forgotten but are necessary in order to avoid situations of the accumulation of rights and the inadvertent transfer of excessive rights to new users. A situation frequently observed in large organisations concerns a staff member transferring, permanently or on secondment, to a different department or division. Very often the move will require the granting of additional access rights to allow the staff member to continue to work (for example, in a bank the staff member might be moving from a credits department to private banking in order to gain experience and will thus need access to the relevant functionality and transactions for the new department); granting these new rights is an imperative and typically gets actioned rapidly; reviewing the user's existing rights and simultaneously removing or suspending those that are no longer necessary is less common. Very often the justification for this is that in the context of a brief transfer or secondment it creates an intolerable administrative burden for IT staff to process such rights reductions. Such laxity presents control weaknesses in two areas, however. The first is that the user could find him- or herself in a position in which transactions can be passed or modifications made that bypass the organisation's internal controls and predefined workflows, either inadvertently or deliberately. The principle that this contravenes, that of the *segregation of duties*, is a fundamental element in developing internal control systems. The second is that very frequently new user accounts are created using an existing account as a template or by reusing the account of a staff member who has recently left. As a result of this, inappropriate and excessive access rights can be transferred and inherited without this being transparent to user managers, data managers, or even the users themselves.

The deletion of unnecessary user accounts is a key internal control activity as it removes the opportunity for activities to be performed within the information systems but outside the usual framework of accountability. The key concern in this respect relates to disgruntled ex-employees who have the technical

skills to connect to the organisation's systems from outside and use their still valid account to interfere with processing or modify or copy data. There is also a risk, though, from the use of the account by staff members still within the organisation. Should they discover the existence of a dormant and now unattributed account and either obtain the password or have the password reset to one of their choosing, the way is open to them to use that account for any purposes they choose.

Obviously there are situations where the immediate deletion of an account might not be the best practice, for example when files are explicitly tied to that account, but good practice requires that the account is at least suspended or has its password changed to one known only by appropriate staff until the deletion process can be completed according to defined procedures.

An aspect of access management that is often neglected, especially in environments where no obvious problems are being identified, concerns the regular and systematic review of the use being made of user accounts in order to identify any anomalous behaviour. This can be targeted at identifying the unexpected use of certain sensitive or powerful transactions, at logins from unexpected locations, or simply the use of accounts that are reserved for specific functions or that are known to be effectively dormant because of staff absence, for example. (Of course, the overall security policy should include a position statement and defined procedures on what action to take in the event of an account being likely to be dormant for an extended period as a result of staff absence for medical reasons or sabbatical or secondment. Good practice would suggest that such accounts be locked, following appropriate request by relevant management).

Identification of anomalies and subsequent investigation can help detect inappropriate behaviours and prevent lasting impacts on the integrity and accuracy of data. For example, if analysis of system logs showed the traces of logins late at night at weekends used to pass journal entries in the financial systems or edit data, this would require investigation because prima facie it would be inconsistent with any properly managed organisations policies and procedures. Similarly, a review of IP addresses that revealed a series of remote login attempts from a block of IP addresses allocated to a country in which the organisation has no consistent presence could be indicative of hacking attempts. Careful investigation is required because such traces could be the result of a staff member attempting to connect while travelling on company business, but security should be alert to the potential implications of such traces.

The final group of activities falls into the category of monitoring controls, control activities designed and performed in order to ensure that detailed control activities are being performed consistently and effectively and that exceptions and anomalies are being handled appropriately. In respect of access right management, typical monitoring controls will consist of a periodic – monthly, perhaps, or quarterly, depending of the volume of changes – review of user accounts that have been created, in order to ensure that the authorization procedures have been followed and that the accounts created and the related access rights do match those requested by the user managers. A similar review can be performed by obtaining a current staff list from Human Resources and matching this to a user list from the key systems, in order to ensure that no redundant accounts exist. Very often this is performed as a baseline review, with a more frequent interim review taking the form of obtaining a list of recent departures from the organisation and confirming from system user lists that all necessary accounts have been disabled or deleted.

Such monitoring controls are essential in order to obtain a reliable level of comfort that the internal controls relating to access management are functioning effectively. Each of the individual detailed controls

can be tested individually by managers or auditors, but the overview level is necessary to confirm – as far as is practical - that there have been no gaps in the overall control framework.

6.3.2.2 *Absence or confusion of access policies*

The above discussion presents an outline of the ideal situation; an organisation that is aware of the risks related to inappropriate access and has the means and the expertise to implement corresponding controls at each level. In reality, however, this is rarely the case. Different departments or groups within an organisation will have different policies, or none at all, and consistency and coordination will be absent. This will have impacts in respect of security, regulatory compliance, and reputational risks. Implementing an overall access policy and related measures is always a worthwhile endeavour, but can often be a lengthy and expensive process.

6.3.2.3 *Single-Sign-On – desirable or a single point of failure?*

The concerns about user access rights discussed above take on even greater significance when considered in the context of an environment in which Single-Sign-On (SSO) has been implemented. This greatly facilitates the lives of users, for whom the need to remember a variety of passwords and conform to each system's individual password strength requirements and rotation periods can be a burden that leads to weaknesses in controls through the writing down or reusing of simple passwords, but does present a single point of failure in respect of access to all of an organisation's systems. Simply stated, if user access procedures fail temporarily in respect of a single, not financially significant application, the impact is unlikely to be critical for the organisation. If the same procedures fail in respect of the user account belonging to a senior manager in an SSO environment, the potential for significant impact is very much greater.

6.3.2.4 *Key concepts in access management*

It is important to distinguish the key concepts in access management so that their significance can be recognized and their relative importance to each environment be appropriately assessing, thereby allowing management to design and implement the internal control mechanisms that best correspond to their control objectives and risk profile.

Authentication is the process of verifying the identity of a user or system. It should not be confused with *identification*, which is how a user or service first identifies itself to a service or platform, typically by means of presenting a user identifier (ID). Authentication is more robust and involves the verification of additional credentials provided by the requestor, such as a password or valid network information.

The next stage is *authorization*, which is the process of determining the privileges the user or system is entitled to once its identity has been established. Generally speaking authorization is used to determine whether the user or service has the necessary privileges to perform certain operations, and thus can be seen as a key element within the process of enforcing policies.

As discussed above, *auditing* is a key part of the identity and access management process. It consists of reviewing and examining the policies and procedures designed and implemented by management, and of examining the evidence of the operation of the internal controls, evidence that can take the form of authentication and authorization records, system log files, activity reports and manually-generated and stored documentation such as access request forms. When appropriate policies and procedures are in

place and sufficient evidence is retained, it is possible to verify compliance both with the established security policies and, by extension, with external requirements. Auditing is also a key tool for detecting failures in security services or structures and providing the basis for recommendations for improvement.

6.3.3 Summary

In this section I have outlined the major aspects of the management of security that management will always need to address. There is an almost instinctive tendency to view information security as a technical issue to be addressed through technical means; I have sought to demonstrate that above all it is a management issue and that security cannot be provided (within the constraints that security can never be definitively acquired) without clear, communicated and enforced management procedures. Technical approaches and tools have their part to play in the overall security framework, as I will discuss in respect of cryptography in Chapter 9, but without being firmly embedded in a properly structured process they cannot provide any more than a partial solution. Information security is a moving target that requires constant and focused effort to strike.

6.4 Privacy Concerns

Conceptually privacy is intrinsically linked to security. There is increasing understanding that adequate privacy cannot be ensured without appropriate security, although many recent media stories suggest that even seemingly adequate security is not on its own sufficient to provide privacy to the subjects of data.

6.4.1 Key concerns and considerations

Both data owners and the subjects of data can legitimately be concerned about aspects of cloud computing, from the perspectives of both security and privacy. A generalized lack of confidence in the capacity, ability, or willingness of service providers to ensure the privacy of data, particularly in the light of increasingly common revelations of the extent to which certain state agencies request and obtain access to data held by application and service providers, has emerged from a number of specific considerations. These include:

Access: knowing what personal information and data are held and who can have access to it, be this for processing and management purposes, through partnership agreements, through commercial contracts for, say, marketing purposes, or for law enforcement and security-focused monitoring purposes.

Audit and monitoring: knowing how organisations and individuals can monitor service providers and how stakeholders can obtain the comfort they need that their privacy requirements are being met.

Compliance: knowing which compliance requirements are applicable for any given service provider (which can be very opaque for the average end-user) and whether these requirements are being met. This issue is particularly complicated in the cloud environment because of the very nature of the cloud: to determine which requirements are applicable, it is necessary to know how many copies of data are stored, in production systems, test systems, backup systems and archives, and be able to catalogue where each copy is physically located. If these copies are stored across multiple jurisdictions, determining which jurisdiction and which requirements are applicable will require appropriate legal and procedural expertise.

Destruction: knowing how a service provider performs and manages data destruction, with particular consideration given to ensuring that all copies are indeed destroyed (from multiple locations, backups and archives) within a predetermined and appropriate timeframe.

Privacy breaches: knowing how a service provider is aware of security breaches and manages the process of notifying those affected.

Retention: knowing how long personal information is retained and whether this is governed by robust and enforced policies.

Storage: knowing where the data in the cloud are actually stored and whether they are kept separate from other datastores.

6.4.2 Key principles

There exist a number of key principles relating to privacy risk management and compliance in the context of cloud computing. Organisations intending to use such services should consider their risk profile and attitude in respect of these principles in order to ensure that they are not exposing themselves to undue risks (in respect of security, privacy, exposure or liability) as a result of transferring data and information outside their immediate security perimeter.

The *Collection Limitation Principle* states that the collection and accumulation of personal data should be restricted to the minimum possible for its purpose. Data collection should be lawful and fair. This is an area where there is confusion between the concepts of security and privacy. As discussed above, there are frameworks and standards available that explain security and enable organisations to implement appropriate measures, but there is no such consensus in respect of privacy. As a result, service providers may in good faith be proposing what they consider to be appropriate structures and safeguards, but these may not meet either a client organisation's requirements or particular legal or regulatory needs.

The *Use Limitation Principle* specifies the circumstances in which personal data can be disclosed or used. Although straightforward in theory, based on the consent of the subjects of the data and on the authority of law, this becomes more difficult to follow in practice as datasets are accumulated and new opportunities are created to link data and view datasets in different ways. The original conditions can easily be forgotten or ignored.

The *Security Principle* specifies that reasonable measures should be in place to safeguard data from inappropriate or unauthorized access, loss, destruction, use, modification or disclosure.

The *Retention and Destruction Principle* concerns the periods for which data should be retained, which broadly speaking is determined by the duration of the task or process for which it was gathered or by legal or regulatory requirements. The expectation is that data should be destroyed securely once these periods have expired. Once again, this is a principle that is simple to state but less simple to implement, for organisations have traditionally been poor at defining their own strategies for data retention and even at recognizing the purposes for which data had been collected. Very often a pragmatic approach is adopted, based upon legislation and regulations such as the Health Insurance Portability and Accountability Act of

1996 (HIPAA), the Sarbanes-Oxley Act (SOX), and other compliance requirements such as those maintained by the US Food and Drug Administration in respect of research and development and test data.⁸

The *Transfer Principle* states that data should not be transferred to countries that do not provide the same level of privacy protection as the organisation that collected the information. This is particularly relevant to the cloud computing environment, where infrastructure is routinely shared between organisations and processing sites. As services expand and become integrated in what is a dynamic environment, the opportunities for data transfer increase, and with them the chances for the principle to be contravened.

Adhering to this principle is particularly difficult in the cloud environment because of the nature of the cloud and cloud services: once outside the owners' own perimeter, the data could be stored anywhere within the service provider's network. Without carefully-worded agreements in place to define where and how data can be stored, this does not immediately coincide with legislation in place in some territories and regions (such as the EU Data Protection Directive - Directive 95/46/EC [79]) that obliges companies to know at all times the physical location of personal data in their possession.

Making this aspect even more complicated is the failure of reciprocal agreements and legal measures to provide full coverage of an organisation's obligations. As an example of this, the US Department of Commerce, in consultation with the European Union, created the US-EU Safe Harbor Framework [80]. This programme, approved by the EU in 2000, provides "a streamlined and cost-effective means for U.S. organisations to satisfy the Directive's "adequacy" requirement"⁹. Compliance with this programme may not, however, satisfy a multinational's EU legal obligations, because in cloud computing data could be stored on servers outside of both Europe and the United States, thereby falling outside of both jurisdictions and making the programme ineffective. In addition, this option may not be appropriate and available for certain industries and organisations not regulated by the US Federal Trade Commission, such as the financial services industry.

It is of course possible to avoid such difficulties by contracting cloud services with a provider that only operates in one territory or region, such as Safeswisscloud within Switzerland [81]. This does, however, present both financial and philosophical issues. Service providers proposing guarantees of services within one territory or region, especially in a part of the world where living and operating costs are high, are likely to have pricing structures that reflect these costs. At the same time, such solutions appear impractical and limiting because they exclude the flexibility and efficiency that cloud computing is designed to provide.

The *Accountability Principle* states that organisations, and designated individuals within organisations, are responsible for information under their control. This responsibility remains regardless of where data are processed or stored, and thus there is a requirement to ensure that policies and procedures are accepted and followed by all other parties who access, maintain or process the data.

⁸ The author can mention the case of a company active in the field of medical technologies and subject to the requirements of the FDA in all of its activities. Not only was it necessary to retain test data for its products, but because of the need to have these data available for inspection at short notice, and because of the rapid changes in technologies in that sector, it was necessary (where migration and emulation were not possible) to preserve the original processing and systems environments in full working order. The company's data centres contained a veritable museum of information systems of great interest to the technically-minded visitor.

⁹ A similar agreement is in place between the USA and Switzerland. See <http://export.gov/safeharbor/swiss/index.asp>

6.4.3 Summary

Privacy is an increasingly important topic nowadays and increasing numbers of organisations are being requested to demonstrate how they can ensure the privacy of the individuals or other organisations that are the subjects of the data they hold or process. Should an organisation choose to exploit the unstructured data it holds or has accumulated, process big data, and/or move onto cloud services, it should be aware of privacy concerns and be able to reply appropriately. In the context of cloud services or other outsourcing, it should also consider how it will be able to obtain responses and confirmations from the third party and share these with other interested parties.

6.5 Data lifecycle

The whole question of management of the *Data Lifecycle* requires careful consideration, as alongside security management it is the cornerstone of effective data management in all circumstances. The ideal for management is to ensure that only relevant, useful, consistent and valid data are retained and stored, thereby minimizing costs and ensuring that security policies are being applied to appropriate sets of data.

The data lifecycle can be broken down into a number of considerations, each consisting of a number of elements that it is necessary to address. These considerations cover the generation of the data, how the data are used, transferred, processed and modified, stored, archived, and destroyed. As presented here, the assumption is that data are being stored in the cloud, but the principles apply to all models of data storage.

6.5.1 Generation

- Source: are there procedures in place to record and evaluate the sources of data in order to evaluate data quality and consistency and trace any issues?
- Ownership: who has clear and acknowledged ownership of the data, and how is this maintained – or, if necessary, formally and knowingly transferred – if cloud services are used?
- Classification: what are the structure and methodology behind data classification? Has it been determined that any classes or groups of data should be excluded from external storage?
- Governance: has the organisation implemented any structures for the governance of data throughout the whole data lifecycle, whether this is internal or external?

6.5.2 Manner of use

- Location and use: is the information used only for internal purposes, or is it shared or distributed outside the organisation?
- Appropriateness: does the use of the information correspond to the purpose for which it was collected or created? Is this consistent with any commitments made to data subjects or regulators?
- Discovery/subpoena: is the management of the information in the cloud consistent with allowing compliance with legal requirements in the event of legal proceedings?

6.5.3 Transfer

- Transit across public networks: if information is transferred to or from a cloud using public networks, is it protected appropriately in respect of risk level and legal requirements?
- Encryption requirements: is information encrypted where necessary?
- Access controls: Are suitable access controls in place when information is being transferred to and from the cloud?

6.5.4 Transformation

- Derivation: can data owners be sure that original protections and limitations on use can be maintained when data are transformed or further processed in the cloud?
- Aggregation: are data owners aware of the circumstances in which data in the cloud might be aggregated, one consequence of which being that information might no longer be related to an identifiable individual and thus subject to less stringent security concerns?
- Integrity: is the integrity of information appropriately maintained and ensured when it is in the cloud?

6.5.5 Storage

- Access controls: have suitable access control measures been implemented to ensure that access is restricted to genuine users in the course of legitimate activities?
- Structured versus unstructured data: are data stored in such a way as to permit ongoing and future access and management?
- The CIA Triad: can management demonstrate how data integrity, availability, and confidentiality are maintained in the cloud?
- Encryption: can management demonstrate that stored data are encrypted properly in accordance with applicable legal requirements?

6.5.6 Archival

- Legal and compliance: does the service provider support specific requirements in respect of duration and location and can this be demonstrated?
- Off-site considerations: can the service provider offer long-term off-site storage that supports archival requirements?
- Longevity of media: will the media used for information storage be accessible in the medium- and long-terms? How will accessibility be monitored and what are the plans for renewal and transfer to newer media?
- Retention: how long are the data retention periods proposed by the service provider and do they meet the organisation's own policies and requirements?

6.5.7 Destruction

- Secure: does the service provider have procedures to securely destroy information obtained by customers or third parties in order to maintain its integrity and confidentiality?
- Complete: when information is destroyed, is it completely removed or can it be recovered?

6.5.8 Summary

The data lifecycle is reasonably straightforward to describe but, as this brief discussion demonstrates, there are a number of detailed processes to follow and factors to consider when adopting a rigorous and structured approach to data management. Fundamental to the whole process is a clear and detailed understanding of what data are held, for what they are being used, and for how long they should be retained. The concepts of governance, classification and compliance apply throughout the lifecycle and management need to have the tools and resources in place to address these issues.

6.6 Conclusion

This chapter has presented the first set of constraints faced by management in respect of the adoption of new technologies or practices. There exists a wide range of risks that need to be considered and addressed, while the single topic of information security has been shown to be vast and requiring significant and structured management input and control. The over-riding learning outcome for management is that what appears at the outset to be an extremely technical and technology-dependent project – be it to manage and exploit varied data sources, or to migrate onto the cloud, or both – is at its heart a traditional operational business problem that requires a great deal of management input and commitment and no small degree of analytical and communications skills.

Chapter 7 An Analysis of Taxonomies and Models

7.1 Introduction

“Taxonomy is described sometimes as a science and sometimes as an art, but really it’s a battleground.”¹⁰

It will be inescapable for any well-managed organisation to embark on any significant project without being aware of the taxonomies, standards and models that apply. These are both significant and useful for two reasons. They provide guidance on good practice to ensure that risks are addressed and that key issues are not omitted when defining scope and developing procedures. They also provide frameworks against which performance can be evaluated and compliance assessed. These are not only external constraints, although often compliance will indeed be an imposed concern for organisations. It is in the interest of every organisation to monitor its performance and identify any issues that occur, and such monitoring is facilitated by the use of standard documentation and models.

In this chapter I introduce and discuss a number of significant taxonomies, standards and models with which organisations should be familiar when considering significant information systems projects. Not every element will be relevant to every organisation in every context, but it is good practice to be familiar enough with the key frameworks in order to know which ones to select and which ones to discard.

7.2 Security event taxonomy

When designing and implementing a security framework, an important aspect concerns how security events will be classified and recorded. All organisations know – or should know – of the importance of implementing measures to prevent and detect security events. In my experience, however, not all are aware of the advantages to be gained from ensuring that they are recorded and classified in a standardised and robust way.

The reasoning behind this is straightforward and identical to recognised good practice in respect of helpdesk management, in which support calls are logged and classified. This helps with the identification and implementation of solutions, subsequent trend analysis aimed at identifying points of weakness and addressing these, and in monitoring the performance of systems and any areas where users might require more proactive support and training.

Security events can usefully be treated in the same way and yield similar benefits. If organisations know precisely what kinds of events are occurring, what is causing them, what is being targeted, if that is the

¹⁰ Bill Bryson, *A Short History of Nearly Everything*, p. 437

case, and whether the threats are coming from inside or outside the organisation, they will be able to take appropriate countermeasures and apply resources in the most efficient way.

In 2013 the European Telecommunications Standards Institute (ETSI) published a security event classification model and taxonomy [82] that provides a wide-ranging and solid framework for the type of analysis discussed above.

ETSI has published a series of group specifications based around applying an event based vision of security management to the three main types of security measures: security prevention measures, event detection measures, and event reaction measures. They define the main goals of security management as being:

- To reduce the residual risk incurred by information systems operationally and constantly; and
- To assess the actual application and effectiveness of security policies (whether part of an ISMS or not) for the purpose of continual improvement.

This definition is very risk and process based and very much aligned with the approaches adopted by COSO and Cobit that will be discussed later in this chapter.

7.2.1 Security Information and Event Management (SIEM)

The difficulty in managing security is well summarized by a diagram showing the scope and limitations of existing policies and procedures, with the differences between what is known and unknown in respect of risks and vulnerabilities.

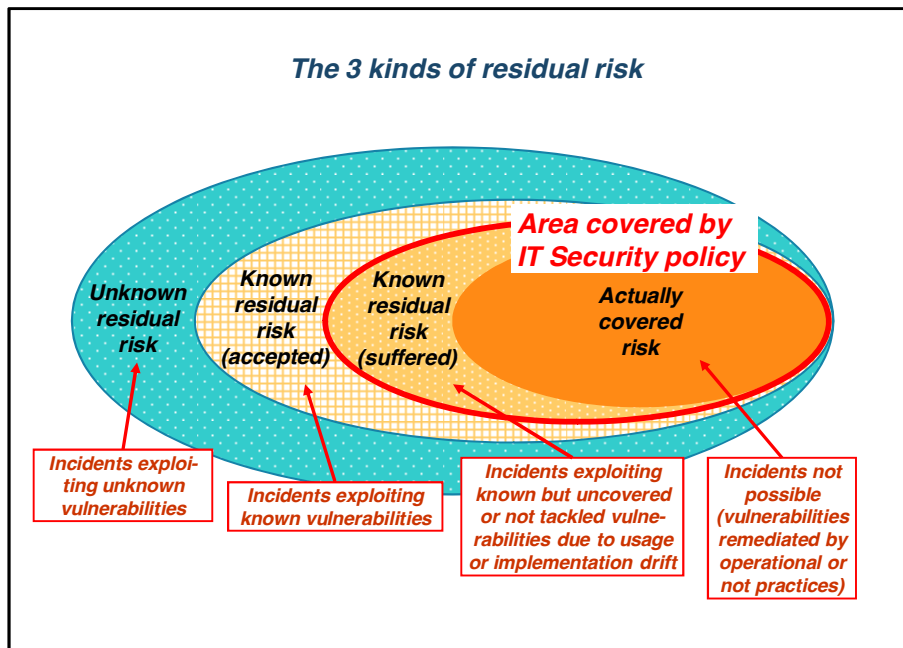


Figure 7.1 The three kinds of residual risk, from ETSI

ETSI explain that the objective of their framework is to “build a full taxonomy to thoroughly describe all IT security events (and when appropriate and necessary non-IT security events) and, based on it, to present an original representation that leverages the current international best practices and enables diversified and complex uses. The choice of a detailed taxonomy, which describes security events through a set of attributes (different for incidents and vulnerabilities), ensures that all possible situations can be taken into account with the required flexibility (especially thanks to the provided open dictionary), while the representation chosen for the taxonomy, highlighting the main categories generally accepted by industry consensus, makes the event classification model easier to understand and embrace for stakeholders.”

What is particularly interesting about the ETSI model, and why it is of relevance in this chapter, is that it draws upon current international best practices and thus is closely aligned to the other generally-accepted standards and models that will be discussed. This aspect of consistency is essential if quality results are hoped for.

7.2.2 Background to the taxonomy

The background discussion of the taxonomy refers, among other standards and models, to the ISO/IEC 27000 family of standards.

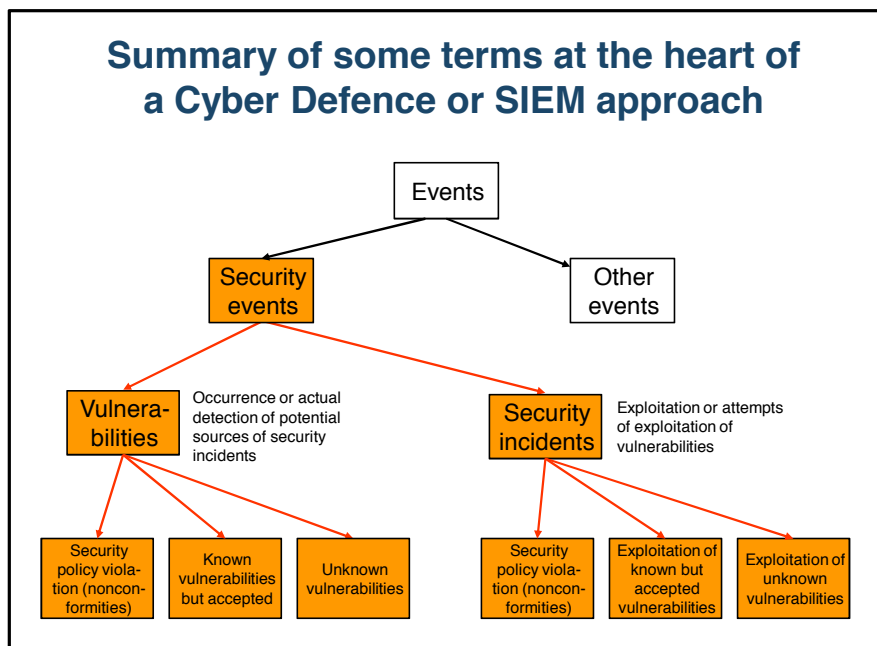


Figure 7.2 Definition of cyber-defence terms, from ETSI

There is a particular link to ISO 27004 – “a structure linking an information need to the relevant objects of measurement and their attributes ... The information needed generally relates to high level data used by IT security governance for decision support and improvements, whereas the objects of measurement (especially those related to technical controls) are often low level data produced by prevention and detection systems. Establishing a link between these two different worlds is easier when using an

intermediate level of abstraction such as an event classification model. Building indicators consists therefore simply in counting events during intervals of constant time”

This is a pragmatic approach designed to facilitate the visibility of linkage between the two areas, high-level and low-level.

The stated requirements for the model also position it clearly as a tool that is applicable across a wide range of organisations, capable of handling a wide range of events and causes of events, and being useful to a wide range of users. These high-level requirements can be summarized as:

- Being positioned at an appropriate level of abstraction;
- Being equally applicable to vulnerabilities and incidents;
- Including both a taxonomy (that ensures completeness and rigour) and a related representation model (that ensures easy understanding and use by all stakeholders and also enables the link with indicators);
- Handling complex security incidents represented as a combination of smaller elementary security incidents.

7.2.3 Philosophy and structure

The modular structure is of interest as a pragmatic and practical approach to developing a model that will be applicable across a range of circumstances and infrastructures. In order to be widely acceptable and applicable, models need to be sufficiently general to be relevant to a wide cross-section of the population but not so non-specific that no value is added.

The description of the model sets out the underlying philosophy, derived from a risk-based modelisation of the problem. “It is key in events and categories design and development to approach the model from the perspective of dealing with residual or not-covered risks hanging over the Information System. This implies that the model has to deal with the 2 components that constitute risk exposure (namely incidents and vulnerabilities, covering both detection and recovery).”

In order to do this, the model needs to sit between the model of the causes and motivations behind (malicious) acts and the model of the risks closely associated with protecting the key elements of confidentiality, integrity and availability in various industry sectors and business lines.

The authors argue that the relationship between the third and the second position makes it possible to establish very accurate links between the event classification model and the risk profile model, describing the links between the incidents categorized in the former and the results documented in the latter, and progressively get closer to an “accurate and exhaustive knowledge of IT systems disasters (notably Confidentiality and Integrity, which are not sufficiently and accurately measured) and to an ability to quantify them for internal purposes.”

This is interesting because it pragmatically and constructively views information security as an ongoing and iterative process. The model as presented is designed to allow management and users to improve their understanding of risks and the impacts of their manifestation, in order to increase their knowledge of incidents and how to handle them.

According to the developers of this model, “the challenge for companies or organisations is thus to implement their own genuine business-oriented IT security observatory. Such an advanced security approach, which is focused on strategy rather than on costs, brings to light IT security aspects that impact the success or image of the corresponding business line (also influenced by the geography, the local market and the competitors).” This is a very wide-ranging view of information security that considers the benefits to the organisation as a whole, in contrast to the traditional and still widespread view of security as a cost that is imposed upon organisations without yielding any clear or quantifiable benefits.

The authors explain the necessity for the model to rest on a detailed taxonomy. “To make sure that the model is exhaustive, and given that it is almost impossible to describe all possible situations, it is necessary to figure out how to define as accurately as possible security events with all the flexibility required. This leads to the definition of a detailed taxonomy (using distinct attributes for incidents and vulnerabilities), and to a complete and open dictionary of acceptable values for these attributes. With such an approach, the model needs however a way to highlight the main categories and sub-categories, to make it as understandable and usable as possible by all stakeholders.” This explanation illustrates the recognition at the design phase of the limits of a taxonomy to be applied to a disparate and variable population, and emphasizes the need for clarity and flexibility, as well as the critical aspect of usability. A common failing of the implementation of standards and models, in my experience and clearly in that of the developers of this model, is that for reasons of complexity or lack of clarity or lack of communication, the model is considered to be too difficult to use and thus is either ignored or used badly. Both outcomes can cause unwanted consequences: if the model is not used, there are no data that management can analyse to identify problems and design improvements; if the model is used badly, the data generated may be inaccurate and misleading and may skew management analysis sufficiently to induce the taking of inappropriate decisions.

7.2.4 Description of the taxonomy

7.2.4.1 *Classification of incidents*

The model structure and taxonomy used to describe incidents is as follows, with eight areas required to fully describe a change in a system:

- 1) Who and/or why (subject)
- 2) What (verb1)
- 3) How (verb 2)
- 4) Status of incident (attempt underway or success or failure)
- 5) Which vulnerability is being exploited
- 6) On what kind of asset (complement)
- 7) With what CIA consequence
- 8) With what kind of business impact

Each area may include up to three fields. Three areas are mandatory (who and/or why, what, status, asset) and the others being optional (how, vulnerability exploited, CIA consequence, impact). This is interesting as

although the mandatory fields ensure that the bare minimum of information is recorded, they do not necessarily allow sophisticated analysis of long-term trends.

7.2.4.2 Classification of vulnerabilities

The model structure and taxonomy used to describe vulnerabilities is as follows, with five areas required to fully describe a state:

- 1) What
- 2) On what kind of assets
- 3) Who (only for behavioral vulnerabilities)
- 4) For what purpose (only for behavioral vulnerabilities)
- 5) To what kind of possible exploitation

Each area may include up to three fields. Three areas are mandatory (what, asset, who) and the others are optional (purpose, exploitation). Clearly the use of the non-mandatory fields requires a sophisticated understanding of the nature and reasons for each attack. This is not necessarily available to every organisation, partly because of a lack of experience and education, and partly because in the case of sophisticated attacks it can be extremely difficult to identify the source, let alone the objectives and motivations [83].

7.2.4.3 Top-level choices in respect of incident areas

The top-level choices provided by this dictionary for incident areas can be summarized as follows:

Feature	Description
Who and/or why	Accident, unwitting or unintentional act (error), unawareness or carelessness or irresponsibility, malicious act (external or internal)
What	Unauthorized access to a system and/or to information, unauthorized action on the information system and/or against the organisation, installation of unauthorized software programs on a system (without owner's consent), information system remote disturbance, personal or organisation disturbance to induce stress, physical intrusion or illicit action, illicit activity carried out on the public Internet network (harming an organisation), various human errors (administration, handling, programming, general use), breakdown or malfunction of equipment, environmental events (information system component unavailability due to a natural disaster)
How	For each abovementioned "what" possibility, list of the methods and tools used (especially for attacks)
Status of incident	Security event attempt underway, successful security compromise, failed security compromise

Feature	Description
Vulnerability exploited	Based on the six possible types of vulnerabilities (behavioral, software, configuration, general security, conception, material)
Asset type	Databases and applications (perimeter, internal, public cloud, outsourcing), systems (perimeter, internal, public cloud, outsourcing), networks and telecommunications (low level devices, high level communication, middleware, wireless devices, security), offline storage devices (paper, electronic devices, magnetic devices, optical devices), end-user devices (local application software – user or organisation-owned, multipurpose-se workstations – user or organisation-owned), people (employee, business partner, on-premises or off-premises service provider), facilities and environment (real estate, physical security devices or systems, various utilities, office furniture)
CIA consequences	Loss of confidentiality (Personal Identifiable Information (PII), trade secrets, sensitive data, intellectual property, (military)-classified material, network and systems, security), loss of integrity (many diversified cases), loss of availability (performance decrease, full breakdown or malfunction, deletion, physical destruction, physical loss including theft)
Impact type	Direct impact (disruption to business operations, loss of productivity, fraud, incident recovery costs, life or health consequences), indirect impact (loss of competitive advantage, reputation damage, loss of market share, legal and regulatory costs)

Table 7.1 Top-level choices in respect of incident areas

7.2.4.4 Top-level choices in respect of vulnerabilities areas

The top-level choices provided by this dictionary for vulnerabilities areas can be summarized as follows:

Feature	Description
What	Behavioural vulnerabilities (further broken down into ten different sub-groups), software vulnerabilities, configuration vulnerabilities (see CCE), general security (organisational and technical) vulnerabilities, conception vulnerabilities (software, general design, environment), physical vulnerability (hardware, network and systems, personnel and site, environment)
On what kind of assets	Also covered in respect of incidents
Who (only for behavioral vulnerabilities)	Also covered in respect of incidents
For what purpose (only for behavioral vulnerabilities)	Also covered in respect of incidents

To what kind of possible exploitation	Also covered in respect of incidents
---------------------------------------	--------------------------------------

Table 7.2 Top-level choices in respect of vulnerabilities

7.2.5 Complex security incidents versus basic security incidents

7.2.5.1 *Distinction and significance*

One of the difficulties in describing security events (mainly security incidents) is the possible complexity of the events themselves. The framework document uses the well-known example of so-called APTs (Advanced Persistent Threats), to explain how a complex security incident can be broken down in more elementary security incidents.

To record such complex incidents in an internal database, it may be more pragmatic and practical to record each associated basic incident separately, in order to gather the richest information possible on such attacks. In particular, the breakdown of information can allow analysts to identify more precisely the attack vectors used (for example, which kinds of media and masking techniques were employed for malware installation. In the cases attacks that were not fully successful, it will be useful to know at what point, and for what reasons, they failed or were stopped.

Other examples given in the document of security incidents considered to be less complex are website defacement (intrusion followed by defacement), the misappropriation of resources by external attackers (intrusion followed by misappropriation), and backdoors on externally accessible servers acting as a foothold for future intrusions (intrusion followed by malware installation). Each of these consists of at least two sequential stages of malicious activity, and it is in the interest of those responsible for securing and monitoring systems to understand the precise sequence of events and the nature of each activity in order to be able to recognise the signs of attacks and implement appropriate defences.

In this model, security events are structured into seven major categories. The first three categories describe security incidents that constitute tangible threats to the information systems. The four remaining categories describe vulnerabilities and involve four of the six major types of vulnerabilities that can be exploited in order to implement an attack and create an IT security incident.

7.2.5.2 *Categories of security incidents constituting tangible threats*

The first three categories describing security incidents are as follows:

- The first category, called "*Intrusions and external attacks*", includes all types of incidents of a malicious nature coming from the outside.
- The second category, called "*Malfunctions*", includes all types of accidental (failures or breakdowns or natural disasters) or unintentional (human error) incidents.
- The third category, called "*Deviant internal behaviours*", includes all types of incidents of a malicious or intentional nature (negligence, recklessness and irresponsibility) originating within the company's or organisation's security perimeter.

7.2.5.3 Categories of vulnerabilities

The second group of four categories describing vulnerabilities are as follows:

- The first category, called "*Behavioural vulnerabilities*", includes events that are similar to some that are classified in the third category for incidents, without immediate CIA consequences, but where the deviant human origin of vulnerabilities needs to be pointed out.
- The second category, called "*Software vulnerabilities*", includes events stemming from upstream in-house or standard software development.
- The third category, called "*Configuration vulnerabilities*", includes events indicating a deviation from the accepted state-of-the-art in security policies or best practices, or a weakness that is exploited in attacks.
- The fourth category, called "*General security (technical or organisational) vulnerabilities*" includes vulnerabilities that can be defined as having an overall and significant effect on the security level of the information systems. Such vulnerabilities are positioned on a level equivalent with one of the control points set out in the ISO Standard 27002.

It is emphasized in the model that should that all security events, incidents or vulnerabilities, become examples of non-conformity when they violate the organisation's security policy and rules. The authors do, however, make an exception in respect of software vulnerabilities because these are of a somewhat different nature; the means of securing systems against such issues being different and the scope for autonomous security activities being limited.

While this model does provide a wide range of categories and classifications for incidents and vulnerabilities, it should once again be emphasized that there is a fundamental reliance upon users to apply these categories appropriately. In practice, in my experience, there can be a tendency to use one or more categories as a kind of catch-all to which non-standard or non-obvious incidents are assigned. As discussed above, this can lead to difficulties with later analysis and require effort and resources to rectify after the event.

7.2.6 Comparison with other event classification models

Several alternative or associated classifications are currently available within the framework of risk analysis and/or methods for defining security objectives definition (EBIOS, Mehari, CRAMM, Octave, etc.); ISO reference frameworks (for example, ISO 27005); MITRE reference frameworks (CAPEC); or recommendations within IT security communities (such as FIRST). These classifications generally cover all security incidents, and often vulnerabilities too, including hardware or software aspects, and whether they relate to breakdowns, malfunctions, human errors or malicious acts.

7.2.6.1 Risk analysis method classifications

The risk assessment methods and classifications are positioned mostly between position 2 and position 3 of the ETSI model. From a cyberdefence perspective, these classifications tend to exhibit limitations in comparison with the ETSI model: these limitations are primarily a result of their stated objectives. Typically the reader can identify:

- The absence of industry-wide statistics at the level of detail of the major event categories. In addition, it is sometimes impossible to obtain consistent figures for some events since they apply only to specific industry sectors;
- A very short list of events, lacking sufficient details for widespread use, and sometimes far removed from the current reality of cybercrime;
- A focus primarily on incidents at the expense of vulnerabilities;
- A difficulty in establishing a link with the events described in SIEM tools used to analyse and correlate logs;
- A difficulty in associating a consistent set of reaction plans with its various categories.

In summary, these classifications are structured primarily by the "complement" and the "who" (from the point of view, therefore, of upstream risk analysis), whereas the ETSI model is primarily structured according to the "how" and the "who" (from the point of view, therefore, of a downstream operational cyberdefence and SIEM approach that is focused on detecting residual risk and reacting towards the "who" in question).

7.2.6.2 CAPEC

The CAPEC (Common Attack Pattern Enumeration and Classification) reference framework, which is clearly positioned on position 2, deals with the same kinds of security incidents and complements the NIST SP 800-126 (SCAP) standard. The first CAPEC specifications were published at the end of 2007 under contract to the Department of Homeland Security in the US, and were later included by the Mitre Corporation in the complete series of Information Security Data Standards under the label "Making Security Measurable." CAPEC is a catalogue of attack patterns, along with a comprehensive classification taxonomy. In providing a standard mechanism for identifying, collecting, refining, and sharing attack patterns among the software community, CAPEC allows a more complete and thorough review of the security levels of information systems. Significant differences can, however, be identified with the ETSI approach:

- Security events are limited to malicious activity; the model does not take into account carelessness, accidents and vulnerabilities;
- The model does not include attributes related to the probability of the different types of attacks. The associated notion of priority is thus lacking, due to the absence of link to indicators and up-to-date statistical figures;
- It is impossible to extract from CAPEC a clear hierarchy and structure (for example, internal vs external attacks), and this absence of a tree structure, together with an unstructured list of incidents, can lead to difficulties in using it across a whole organisation;
- The lack of top-down support tends to lead to a focus on the technical description of security incidents, which can be difficult to understand by non-experts.

The contribution of CAPEC to the security community has been to enable understanding of the huge diversity of attack methods, thereby leading to more structured approaches to developing secure software and systems, and improved support for product certification, for example.

7.2.6.3 *FIRST*

The FIRST classifications, first introduced in 2004, are used for communication between CERTs and CSIRTs around the world. They are therefore mainly focused on external security incidents. In respect of the main criteria described above, they exhibit obvious limitations that restrict their use in a cyberdefence and SIEM approach. These include:

- A heterogeneous classification structure (for example, in the "Unlawful activity" and "Compromised information" categories);
- An inconsistently positioned model that can vary, depending on categories, between positions 2 and 3. It can be difficult to build a link with a structured and homogeneous reaction scheme;
- Indicator-related up-to-date statistical figures are in practice impossible to collect;
- Poor or missing distinctions between external and internal incidents for some types of incidents.

7.2.7 Details of the ETSI model

The ETSI model defines a large number of clear and structured categories and sub-categories of incidents. Tables setting out these categories can be found in Appendix A; what is important to note here is the high-level categorization of potential incidents. These are grouped into:

- Intrusions and external attacks (Category IEX)
- Malfunctions (Category IMF)
- Deviant internal behaviours (Category IDB)
- Behavioural vulnerabilities (Category VBH)
- Software vulnerabilities (Category VSW)
- Configuration vulnerabilities (Category VCF)
- General security vulnerabilities (Categories VTC – Technical and VOR – Organisational)

This categorization is interesting and useful because it forces users to consider the probable causes of each incident and allows similar incidents to be grouped together in ways that already indicate the nature of the underlying cause. Clear distinction is made between external and internal attacks, which as discussed earlier require different approaches to prevent and counter, and between attacks and malfunctions. Malfunctions can provoke security incidents in a variety of ways and should be recognized as such; measures to prevent these or lessen their impact may include some security-focused measures such as restricting even more tightly access to sensitive data and enforcing the approval of any changes made to such data, but are also likely to include controls over the robustness and resilience of systems and over the maintenance of a stable operating environment.

The many vulnerabilities facing information systems are grouped into four main categories reflecting their causes. These four categories are fairly wide-ranging: behavioural vulnerabilities regroup a reasonably well-defined range of practices and activities on the part of users that could lead to data leakage or inappropriate access to systems and information, while software and configuration vulnerabilities are both essentially catch-all headings for an extremely diverse range of weaknesses relating to how software is deployed, configured and managed. The final category of vulnerabilities, general security, is likewise wide-ranging and fairly high-level in its definitions. If there is a criticism to be made of the model, it is that there is inconsistency here: some aspects are treated in great detail and incidents or vulnerabilities can be assigned precisely to a specific sub-category, while others are left more vague. As I have argued throughout, this kind of vagueness should be avoided if effective reporting and efficient analysis is required.

7.2.8 Practical uses of the model

A number of practical uses of the model are described. These include:

- Providing input for security event testing to assess the effectiveness of the detection measures. The model (especially the taxonomy) is a basic input to producing lists of security events, testing and subsequently assessing the effectiveness of the detection measures in place. Testing and performance evaluation are key to improving the credibility and RoI (Return on Investment) of security measures by improving event detection rates. The ability to provide a solid ground for precise testing scenarios for a typical set of security events is therefore of the utmost importance to measure the performance of systems and tools;
- Acting as a gateway with the security event classifications used in risk analysis methods (more focused on "asset" and "impact" than on "what"). The aim is to extract general results from the risk analysis undertaken for specific individual systems, while refining the residual risk contents that are updated by means of the model event types and associated lists. A progressive enhancement of the classifications currently proposed within the risk analysis methods could eventually lead to their replacement by this model;
- Providing common objectives to jointly evaluate the risks of potential IT security incidents on an organisation's information systems and their consequences on the organisation itself;
- When implementing a Continuous Auditing scheme (corresponding to ISO 27002 or COBIT), tracking vulnerabilities and the application of security-related practices by relying on a comprehensive generic structure that deals with all kinds of security events (incidents, vulnerabilities and nonconformities), while using a common language and taxonomy. The contribution of the model is either to help with the definition of a series of control points where no control of this type had previously existed, or to refine and optimize existing controls. In all cases, according to the authors, the core value of the model is the overall consistency with other concomitant usages (such as an organised reaction scheme or benchmarking through indicators), and the possibility of defining a precise implementation of these two uses by means of an organisational interface between the monitored system and the possible central SOC in charge of this monitoring. Such a strictly formalized interface can contribute to progress and to the following clarifications:

- Accurate definition of residual risk taken into account (vulnerability x threat x impact), while identifying priority 1 action levers;
 - Prioritized list of events (incidents, vulnerabilities, nonconformities) on the IT system that is monitored within the framework of an internal service contract, while relying on the classification model;
 - Precise evaluation of event criticality (event severity x impacted target sensitivity);
 - Thorough identification of the people involved in the security chain (security managers or correspondents on various levels, network or system administrators, operational owners or managers of the monitored IT system components, etc.), and their exact role as part of handling such events;
 - List of standard reaction plan types that have to be launched;
 - Organisation of escalation and of crisis management structures in case of non-standard events and ones not categorized as critical and/or ones that escape the SOC's control;
 - Definition of the expected SLA-type service levels (time for notification, qualification, recommendation of solution, event resolution, etc).
- Detected security events that can be tied up to standard event types associated with representative public reference statistical figures (at industry sector level). The model provides an event classification at the right level (as determined in a convergent manner within a series of advanced SIEM projects), which makes it possible to produce an objective benchmarking of organisation's security level based on relevant operational indicators and the establishment of a corresponding state-of-the-art.
 - Possible reliance for insurance companies on this public reference and on the associated metrics to define new insurance offerings for cyber-risks, and to provide new potential "insurance / security investments" tradeoffs. Based on this state-of-the-art and on the associated metrics, another additional use is related to the definition by insurance companies of new insurance policies and associated premiums for cyber-risks. The aim here is to design and verify the relevance of new "insurance / security investments" trade-offs while covering, for example, increasing risks related to cybercrime; typical examples are DDoS attacks (occurring very randomly and irregularly within organisations, but with few prevention means and consequences that can be significant in terms of image) or internal fraud committed by employees or external service providers (somewhat more predictable occurrence frequency than the previous event, but also with prevention means and technical and procedural controls that cannot eliminate all risks, and financial consequences that can be considerable). Such events lend themselves well to the insurance notion, and thus to the pooling of risks within a profession.
 - In-depth structured analysis of practices and motivations of external malicious activity as well as of abnormal or deviant internal user behaviours, including the possibility to enrich this analysis with information coming from "Counterintelligence" activity.

- Preparation of more readable reports and dashboards, with indicators directly related to certain types of very technical events that can only be managed by specialists. This additional usage involves the definition of operational indicators for incidents and vulnerabilities / non-conformities extracted from the ones which can be associated with a state-of-the-art. This involves relying on the classification model and its major event types (in principle, extensively disseminated and known within company or organisation), preparing simpler and more readable reports and dashboards with indicators that are often innovative and also directly related to some types of very technical events that can only be grasped by specialists; for this purpose, it is also possible to rely on an existing reference framework.
- Linkage established between technical events detected by the available tools and a series of pre-defined corresponding reaction plans (with possible reliance on a reference framework in this domain). This relates to the possibility of associating, with many types of categorized events, standard reaction plans that are intended to deal with these events and to come back to a normal situation. To that end, argue the authors, it is better to rely on an existing reference framework, which typically includes some thirty standard plans generally organised into five strictly formalized steps. Experience has shown they would manage to cover, within the best international SIEM projects, nearly 95 % of the events occurring in the real world, and that more than 90 % of the launched plans were efficient for solving the problems resulting from events.

7.2.9 Summary

The ETSI model therefore sets out to provide a through, detailed, robust and flexible framework for recording and classifying security events, and subsequently allow sufficient high-quality analysis for trends to be identified and appropriate action plans developed and implemented.

As a very recent publication, the ETSI model has not yet had the opportunity to demonstrate its worth. Realistically it should be expected that even early adopters will require a couple of years to implement it, review results and draw conclusions, and so publications discussing its operational effectiveness should not be expected in the immediate future. It should also be anticipated, based on the lifecycles of other standards and models, that user feedback will lead to iterative improvements and periodic updates designed to address issues encountered and comments made. It will be interesting to return to the subject by the time the third version of the model, for example, has been published, in order to see what progress has been made and which elements have become most fundamental.

7.3 Security management and internal control standards

The above discussion of the range and scope of possible security incidents demonstrate that the owners of systems and data are confronted by a significant challenge in respect of keeping their assets secure, particularly when faced with the changing context of specific data management or a move onto the cloud.

There are a number of steps that need to be undertaken to ensure that significant risks are addressed in a structured and complete way. In order to achieve this, many organisations will choose to follow the methods set out in one or more commonly accepted standards.

The first step in a cloud project, for example, is to perform an inventory of systems and data and to define exactly where the control perimeter lies. This will involve mapping and understanding all the components and all the layers making up the infrastructure being used, including the network layers, servers, databases, middleware, applications, storage technologies and web services. For each, and overall, it is necessary to have a clear understanding of the scope and split of IT management, monitoring and maintenance responsibilities, for although some operational responsibilities will by necessity be transferred to the cloud provider, some aspects, depending on the nature of the services provided and the terms of the contract and service level agreement (SLA), will remain with the organisation.

This is specifically relevant to the question of security, because it may well be the case that the service provider is unable to provide security services that comply precisely with the internal security policies and thus one or more additional levels of security may need to be introduced or reinforced in-house. At the same time, even if the service provider does agree to provide a satisfactory level of security, procedures will need to be introduced to monitor the quality of the service provided.

7.3.1 Security management standards

It is rare to encounter an effective information systems management environment in which the structures in place are not drawn from recognized security management standards. Although legislation rarely, if ever, requires adherence to any specific standard (the Sarbanes-Oxley legislation, for example, demands only a robust and structured approach to internal controls without specifying further), it is far simpler for organisations to select a framework that suits their needs and purposes and then follow it to the degree of faithfulness that they deem appropriate.

A number of standards exist in the field of information management and security, with differing objectives and domains of applicability. Not every standard will provide coverage for all internal control requirements even if full compliance with the standard is achieved, and so management need to consider their requirements very carefully before committing to an implementation project.

The three most commonly used standards are ITIL, the ISO 27000 family, and COBIT, while COSO is important as a general organisation-wide internal controls framework. I have chosen these standards in particular because of their popularity, their utility, and because they all emphasize, with differing levels of precision and flexibility, the importance of an internal controls approach to both project management and corporate management. They thus correspond to the underlying approach to good practice and effective management of information systems that forms the core of this thesis.

7.3.1.1 ITIL

The *Information Technology Infrastructure Library (ITIL)* [84], originating in the 1980s within the UK civil service (specifically the Central Computer and Telecommunications Agency), was created as a set of recommendations intended to provide standardization of practices across government agencies and private sector contractors. It began as a collection of books covering specific practices within IT service management before being restructured in 2001 to regroup the content into nine logical sets that corresponded to different aspects of IT management, applications and services.

ITIL is designed to be applied across virtually every type of IT environment, including cloud operating environments. The philosophy behind the guidance is to ensure that effective information security

measures are taken at strategic, tactical, and operational levels. Within ITIL, information security is considered to be an iterative process that must be controlled, planned, implemented, evaluated, and maintained.

The most recent update, published in 2011, contains five volumes:

1. ITIL Service Strategy
2. ITIL Service Design
3. ITIL Service Transition
4. ITIL Service Operation
5. ITIL Continual Service Improvement

Within those volumes, the areas most relevant to security management and the use of the cloud are to be found in second, third and fourth.

ITIL Service Design is itself broken down into seven subdomains:

- i. Design Coordination
- ii. Service Catalogue
- iii. Service Level Management
- iv. Availability Management
- v. Capacity Management
- vi. Information Security Management System
- vii. Supplier Management

Of particular interest here to organisations contracting for cloud services are the domains relating to service level management, availability, and information security management.

As mentioned above, service level management is a critical aspect of any contract with a third party for the provision of services. Service level management ensures that appropriate arrangements are in place with suppliers and that the levels of services are monitored and reviewed. This monitoring should cover a range of measures of the quality of service, including availability, capacity and incident handling, and should also be tied into service continuity planning so that the organisation's continuity requirements are properly addressed.

Availability management might naively be considered more of an operational preoccupation for the service provider, but the data owners also have a significant role to play. They need to be able to anticipate changes in service requirements in order to transmit this information to the service provider and ensure that sudden and drastic changes do not adversely affect the provision of service. More profoundly, the concept of availability is intrinsically bound up with the definition of information security, as it forms one of the elements of the CIA triad. In retaining overall responsibility for the security of information, even when

stored or processed outside the traditional security perimeter, management needs to ensure the consistent and appropriate availability of data.

Information security management describes how structured and consistent information security fits into the management structures. The ITIL recommendations in respect of this are based on the codes of practice set out in the ISO/IEC 27001 and 27002 standards (see below) and are designed to protect information security with the further objective of protecting information assets against risks and hence maintaining their value.

ITIL Service Transition is also made up of seven subdomains:

- i. Transition planning and support
- ii. Change management
- iii. Service asset and configuration management
- iv. Release and deployment management
- v. Service validation and testing
- vi. Change evaluation
- vii. Knowledge management

The key subdomain here is that of change management, to the extent that either the organisation itself retains responsibility for requesting or approving changes, or needs to obtain comfort over the quality of the changes made by the service provider in order to maintain confidence in the security of systems and data.

Change management is a sector of activity that presents numerous risks in respect of the security of data. Clearly defined processes and strict segregation of duties are required to ensure that only authorized staff have access to the different environments (typically, sandbox, development, test, quality assurance and production, in larger organisations at least), that access to production data is restricted, that only properly authorized and tested changes are made to applications, platforms or data, and that all changes made are logged and subject to subsequent review.

ITIL Service Operation is itself made up of five subdomains:

- i. Event management
- ii. Incident management
- iii. Request fulfilment
- iv. Problem management
- v. Identity management

The key elements within this domain are event management and identity management. Clearly the purely operational aspect of event management – an event being defined as an anomaly or exception, or even as

the completion of a process in the course of normal operations – will become the responsibility of the service provider and handled as part of daily operational activities. The data owners should take care, however, to define exactly which kinds of events they would like to be notified about, either in real-time or as a part of periodic reporting, and ensure that such notifications are actually provided, along with details of the subsequent actions undertaken and any trends or root causes identified.

Identity management as discussed above is the responsibility of both service providers from an operational perspective and the data owners from a strategic perspective. Typically the organisation's management will define access rights for users and manage the process of creating accounts and granting the appropriate rights, but there may be circumstances in which it is the service provider that performs the account creation, modification or deletion process. The management of technical and system user accounts within the service provider's infrastructure will of course be entirely the responsibility of the service provider and the data owners will typically have very little visibility over such activities. They will be within their rights, however, to be informed about the processes are designed, performed and managed.

In common with the ISO 27000 standards, ITIL breaks information security down into a series of straightforward and logical steps:

- Policies
- Definition of the overall objectives an organisation is attempting to achieve
- Processes
- Requirements for achieving the objectives
- Procedures
- Responsibilities and duties for achieving objectives
- Work instructions
- Instructions for taking specific actions

Summary

ITIL is thus a robust and well-documented framework for a number of aspects of internal control. In practice it is very often implemented in and around helpdesk and service desk functions, a domain in which its qualities are recognised, but it applies equally well to the domain of information security, decomposing what can be a complicated area for management into a series of simple and sequential steps.

7.3.1.2 ISO 27000

The ISO/IEC 27000 family of standards addresses information security matters, aiming to provide a common framework and vocabulary for the consistent application and evaluation of security matters. The introduction to ISO/IEC 27000:2012 states that “Through the use of the ISMS family of standards, organisations can develop and implement a framework for managing the security of their information assets including financial information, intellectual property, and employee details, or information entrusted

to them by customers or third parties. These standards can also be used to prepare for an independent assessment of their ISMS applied to the protection of information” [85].

The family of standards has been enlarged over recent years to encompass many aspects of information security management, with the list of publications set out in ISO/IEC 27000:2012 including seventeen documents, some of which were still in the approval phase. The currently valid documents are as follows:

Reference Number	Subject
ISO/IEC 27000:2009	Information security management systems — Overview and vocabulary
ISO/IEC 27001:2005	Information security management systems — Requirements
ISO/IEC 27002:2005	Code of practice for information security management
ISO/IEC 27003:2010	Information security management system implementation guidance
ISO/IEC 27004:2009	Information security management — Measurement
ISO/IEC 27005:2011	Information security risk management
ISO/IEC 27006:2011	Requirements for bodies providing audit and certification of information security management systems
ISO/IEC 27007:2011	Guidelines for information security management systems auditing
ISO/IEC TR 27008:2011	Guidelines for auditors on information security management systems controls
ISO/IEC 27010:2012	Information security management guidelines for inter-sector and inter-organisational communications
ITU-T X.1051 ISO/IEC 27011:2008	Information security management guidelines for telecommunications organisations based on ISO/IEC 27002
ISO/IEC FDIS 27013	Guidance on the integrated implementation of ISO/IEC 27001 and ISO/IEC 20000-1
ITU-T X.1054 ISO/IEC FDIS 27014	Governance of information security
ISO/IEC TR 27015	Information security management guidelines for financial services
ISO/IEC WD 27016	Information security management – Organisational economics

Table 7.3 The ISO 27000 family of standards

Of these, the most significant from a practical perspective are ISO/IEC 27000, which provides an overview and a standardized vocabulary; ISO/IEC 27001, which sets out the requirements for information security management systems; and ISO/IEC 27002, which presents a code of practice for information security

management and provides details of the requirements for a management system to be considered complete and effective.

It is important to note that ISO/IEC 27001 formally defines the mandatory requirements for an Information Security Management System (ISMS) and is also a certification standard. Organisations can be audited against and certified compliant with that standard. ISO/IEC 27002 by contrast is used to indicate and identify suitable information security controls within the ISMS, but is not a certification standard. As a guideline, it presents a list of control objectives and possible good practices that organisations can select from according to their requirements, motivations and resources.

7.3.1.3 ISO 27001

ISO/IEC 27001, of which an updated release version was published in September 2013, contains eleven domains, designed to cover the important aspects of the management and organisation of information security, address the key detailed concerns related to the subject, and present how compliance can be demonstrated. This last aspect is particularly significant given the nature of the standard as a basis for certification.

The standard was first published as BS 7799 Part 2, a standard developed by the Department of Trade and Industry within the UK Government and published by the BSI Group in 1999. After discussions with international standards bodies, it was adopted by the ISO in 2005. BS 7799 Part 3, covering risk analysis and management, was published by BSI in 2005 and also aligns with ISO/IEC 27001

The eleven domains are as follows:

- Security policy - management direction
- Organisation of information security – governance of information security
- Asset management - inventory and classification of information assets
- Human resources security - security aspects for employees joining, moving within, and leaving an organisation
- Physical and environmental security - protection of the computer facilities
- Communications and operations management - management of technical security controls in systems and networks
- Access control - restriction of access rights to networks, systems, applications, functions and data
- Information systems acquisition, development and maintenance - building security into applications
- Information security incident management - anticipating and responding appropriately to information security breaches
- Business continuity management - protecting, maintaining and recovering business-critical processes and systems

- Compliance - ensuring conformity with information security policies, standards, laws and regulations

Even in an outsourced or cloud environment, all of these domains are relevant to an organisation's management because of their retained overall responsibility for security and compliance, even if daily operational responsibility is delegated to the service provider.

A particular area of interest is the asset and data classification requirement, under which assets are identified and classified so that an appropriate level of protection can be assigned and implemented. There is no specification of the nature of the classification to be performed, however many organisations choose to follow a four-level system similar to the following:

- Restricted: highly sensitive information the disclosure of which could cause significant damage to the position or reputation of the organisation;
- Confidential: sensitive information the disclosure of which could adversely affect the organisation, its partners or staff;
- Internal: information that is appropriate and useful for internal distribution within an organisation but that should not be made public because its content could be useful to outsiders;
- Public: information that could be or should be made public, such as marketing materials, public statements, or performance information.

A key feature of the standard is its adoption of and reliance on the Plan-Do-Check-Act (PDCA) process model. This iterative four-step management method, also known as the Deming or Shewhart cycle, is intended to promote continuous improvement. In summary, processes are designed that should deliver results in accordance with expectations, then the processes are implemented and executed. The results are obtained and checked against expectations, and then any corrective actions deemed necessary are requested, based on root cause analysis and the requirement to change. This whole process is followed iteratively to ensure ongoing and continuous improvement. Although originally conceived as a tool for manufacturing and processing quality control, the technique also lends itself to systems management processes because of its focus on iteratively improving the quality of information and its method of inducing improvements as a result of systematic problem-solving exercises.

7.3.1.4 ISO 27002

ISO/IEC 27002 is also the result of the efforts of the BSI in developing a national standard for the UK. BS 7799 Part 1 was published in 1995 and revised in 1998, and adopted by the ISO as ISO/IEC 17799 in 2000. The standard was revised in 2005 and finally incorporated into the ISO 27000 family as ISO/IEC 27002 in 2007.

The 2013 version has been renamed "Code of practice for information security controls" in order to emphasise the focus of controls rather than management, which remains the focus of ISO/IEC 27001. This latest version of the standard consists of twenty sections, most of which are broken down into subsections as follows:

Section 0: Introduction: sets out the background, details the contents of the standard, and points to ISO/IEC 27000 for the overall structure and glossary for ISO27k.

Section 1: Scope: describes the applicability of the standard and makes recommendations for those responsible for selecting, implementing and managing information security.

Section 2: Normative references: explains that ISO/IEC 27000 is the only standard considered absolutely indispensable for the use of ISO/IEC 27002; ISO/IEC 27001 will often be useful when considering the overall management of information security but there is no requirement to use the two together. Other standards are mentioned and a bibliography is provided.

Section 3: Terms and definitions: explains that all relevant specialist terms and definitions are now defined in ISO/IEC 27000 and that most of these apply across the entire ISO27k family of standards.

Section 4: Structure of this standard: describes how fourteen of the twenty sections of the standard contain control objectives and activities. Each “security control clause” follows the same standard structure: subsections containing a control objective, each of which is followed by one or more suggested control activities, each of which in turn is supported by implementation guidance and, if necessary, explanatory notes. There are 35 fairly generic control objectives defined with the shared objective of preserving the confidentiality, integrity and availability of information, and more than one hundred proposed control activities, which are noted not to be comprehensive nor universally applicable.

Section 5: Information security policies: management direction for information security, including defining a set of policies set out their management of information security.

Section 6: Organisation of information security: roles and responsibilities for information security should be defined and allocated to individuals taking care to avoid conflicts of interest and prevent inappropriate activities; information security should be an integral part of the management of all types of project; security policies and controls should exist for mobile devices and teleworking.

Section 7: Human resource security: security responsibilities should be taken into account when recruiting permanent employees, contractors and temporary staff, and included in contracts; management should ensure that employees and contractors are aware of, and motivated to comply with, their information security obligations; formal disciplinary processes should exist to handle information security breaches; the security aspects of a person’s exit from the organisation or significant changes of roles should be appropriately managed.

Section 8: Asset management: all information assets should be inventoried and accountable owners should be identified; acceptable use policies should be defined; asset management should be active in respect of leavers; information should be classified, labelled and handled by its owners according to the level of protection needed; information storage media should be managed, controlled, transported and disposed of in such a way that the information content is not compromised.

Section 9: Access control: requirements to control access to information assets should be clearly documented in an access control policy and related procedures; network access and connections should be appropriately managed; the whole lifecycle of user access rights should be managed, including privileged user accounts, with password management and regular reviews and updates of access rights; users should

be aware of their own responsibilities in respect of effective access controls; access to information should be restricted in accordance with the access control policy.

Section 10: Cryptography: there should be a policy on the use of encryption, including cryptographic authentication and integrity controls.

Section 11: Physical and environmental security: premises should be protected by defined physical perimeters and barriers, with physical entry controls and working procedures; specialist advice should be sought regarding protection against fires, floods, earthquakes, bombs and similar threats; ICT equipment and supporting utilities (such as power and air conditioning) and cabling should be secured and maintained; nothing should be taken offsite without authorization; information must be destroyed prior to storage media being disposed of or reused; unattended equipment should be secured and there should be a clear desk and clear screen policy.

Section 12: Operations management: IT operating responsibilities and procedures should be documented and changes to IT facilities and systems should be controlled; capacity and performance need to be managed; different systems environments (such as development, test and production) should be separated and secured; protection from malware should be implemented and kept up-to-date; a backup policy should exist and corresponding backups made and retained; sensitive system events and activities should be logged and reviewed and the logs retained and secured; software installation procedures should be defined and the necessary rights restricted; bugs and vulnerabilities should be patched; IT audits should be planned and carried out with care to minimize both adverse effects on production systems and any inappropriate data access.

Section 13: Communications security: networks and network services should be appropriately secured; policies, procedures and agreements should be in place in respect of information transfer to and from third parties, including electronic messaging.

Section 14: System acquisition, development and maintenance: the security control requirements of new systems should be analysed and specified; policies should be enacted governing secure software/systems development; changes to systems (both applications and operating systems) should be controlled; the development environment should be secured and outsourced development subject to the internal control framework; acceptance criteria for new developments and software changes should be defined to include security aspects; test data should be carefully selected/generated and controlled.

Section 15: Supplier relationships: policies, procedures and awareness programmes should exist in order to protect information that is accessible to IT outsourcers and other external suppliers; service delivery by external suppliers should be monitored and reviewed against contracts; service changes should be controlled.

Section 16: Information security incident management: policies and procedures should exist and responsibilities be defined in respect of managing information security events, incidents and weaknesses and collecting forensic evidence.

Section 17: Information security aspects of business continuity management: information security continuity should be planned, implemented and reviewed as a core part of overall business continuity planning; IT facilities should possess sufficient redundancy to satisfy availability requirements.

Section 18: Compliance: obligations to external authorities and other third parties in relation to information security, including intellectual property, records, privacy, personally identifiable information and cryptography, should be identified and documented; information security measures should be regularly subject to independent audit and the results of these audits reported to management; compliance with policies and procedures should also be systematically reviewed by management.

This standard is thus wide-ranging and contains a large number of control objectives and activities designed to assist management to address the risks that their organisation faces and thus to implement and operate effective controls. The standard does remain, however, broad guidance rather than a recipe book for immediate and unthinking implementation. Management does need to perform a careful and considered risk assessment process in order to ensure that all relevant and significant risks have been properly considered, and then take care in designing the subsequent control activities to ensure that the most effective and efficient responses are being chosen, both in terms of addressing the specific risks and using resources in the most appropriate way.

7.3.1.5 *ISO 27004*

ISO/IEC 27004:2009 sets out to help organisations measure and report on their information security management systems with a view to continuous improvement. This is a very important element in security management because, as in all fields, effective decision-taking is only possible on the basis of reliable information. One specific difficulty in relation to information security management concerns, however, the difficulty in identifying and recording security events, as very often breaches will not be visible to users or administrators and will only become evident if their impact is noticeable.

The main sections of the standard are as follows:

- Information security measurement overview;
- Management responsibilities;
- Measures and measurement development;
- Measurement operation;
- Data analysis and measurement results reporting;
- Information Security Measurement Program evaluation and improvement.

7.3.1.6 *Summary*

The ISO 27000 family of standards is very wide-ranging and very widely-adopted across industries and across international borders. Many organisations seek attestation of their compliance with one or more of the standards as a public recognition of the quality of their management systems. An interesting observation from the field, however, concerns how far many organisations progress in working through and implementing policies and procedures in conformity with successive standards. Very often the first ambition of management is to assess how closely they are aligned to the requirements of ISO 27001, and some form of gap analysis is then performed. The results of this should enable them to position themselves both in respect of ISO 27001 (the existence and operation of information security management systems) and ISO 27002 (the details of information security management). Professional experience shows that often

the findings of the initial gap analysis, combined with a realistic estimate of the amount of work and resources required to achieve both initial and sustainable compliance with the standards, are sufficiently off-putting to make management reconsider their early enthusiasm for such initiatives.

7.3.1.7 COBIT

COBIT – Control Objectives for Information Technology – is described by its publishers, the Information Systems Audit and Control Association (ISACA) as a “globally accepted framework, providing an end-to-end business view of the governance of enterprise IT that reflects the central role of information and technology in creating value for enterprises” [3]. As a business framework for the governance and management of corporate IT, it “provides globally accepted principles, practices, analytical tools and models to help increase the trust in, and value from, information systems” and is widely used by large corporation as part of the structures in place to ensure compliance with national and international internal control requirements. Version 5 of the framework, which was finally published in 2012, integrates other standards and frameworks such as ISACA’s own Val IT and Risk IT as well as ITIL and some of the ISO standards. As such, it reflects current thinking in respect of control issues, scope and best practices.

COBIT is based on five main principles for the governance and management of enterprise IT:

- Principle 1: Meeting Stakeholder Needs
- Principle 2: Covering the Enterprise End-to- End
- Principle 3: Applying a Single, Integrated Framework
- Principle 4: Enabling a Holistic Approach
- Principle 5: Separating Governance From Management

It is important to note that the framework employs a wide definition of stakeholders, including internal and external entities. It also draws a clear distinction between *governance* – assuring that investments in IT generate business value and mitigate the risks that are associated with IT projects - and *management*, which is aimed at ensuring that the information technology resources of an organisation are managed in accordance with its needs and priorities.

The framework then describes seven categories of enablers – the activities, structures, skills and processes that allow effective internal controls to be designed, implemented and operated and allow relevant risks to be addressed. These are:

- Principles, policies and frameworks: the vehicle to translate the desired behaviour into practical guidance for day-to-day management;
- Processes: an organised set of practices and activities to achieve certain objectives and produce a set of outputs in support of achieving overall IT-related goals;
- Organisational structures: the key decision-making entities in an enterprise;
- Culture, ethics and behaviour of individuals and of the enterprise: a critical success factor in governance and management activities;

- Information: required for keeping the organisation running and well governed, and at the operational level very often the key product of the enterprise itself;
- Services, infrastructure and applications: the infrastructure, technology and applications that provide the enterprise with information technology processing and services;
- People, skills and competencies: required for successful completion of all activities, and for making correct decisions and taking corrective actions.

The most fundamental feature of COBIT is that internal controls, including those over security, are seen as part of a holistic and multi-layered approach to corporate governance. Guidance is given in the details of the framework as to the nature of the risks that management could expect to encounter as part of the use of information systems, and control objectives and sample control activities are suggested. This section of COBIT is not, however, a recipe book from which various ideas can be selected and implemented in a haphazard and unstructured way. Management need to define a top-down approach that considers culture, risks, objectives and strategy before effective internal controls can be designed.

This approach is shared with the COSO approach to internal controls, which although not specifically aimed at IT controls issues has been widely adopted for compliance purposes and presents a more generally business-focused framework to accompany COBIT.

Summary

COBIT is a highly regarded and widely followed model for implementing and monitoring risk based internal controls. It is popular with users because of the clear, sequential, drill down structure and the catalogues of proposed control activities, and is particularly popular with auditors, both internal and external, because it corresponds so closely to the demands of the audit methodologies employed by the main institutes and large audit firms. Implementing internal controls in accordance with the tenets of COBIT can be a long and expensive process, but the results are viewed positively by many organisations in respect of the increased control they have over their operations and their ability to satisfy audit demands without undue additional effort.

7.3.1.8 COSO

COSO (the Committee of Sponsoring Organisations of the Treadway Commission) is “a joint initiative of ... five private sector organisations ... and is dedicated to providing thought leadership through the development of frameworks and guidance on enterprise risk management, internal control and fraud deterrence” [86].

The five private sector organisations are:

- American Accounting Association
- American Institute of Certified Public Accountants
- Financial Executives International
- The Association of Accountants and Financial Professionals in Business

- The Institute of Internal Auditors

COSO was initially organised in 1985 to sponsor the National Commission on Fraudulent Financial Reporting, studying the causal factors that can lead to fraudulent financial reporting. Since then, the mission of COSO, drawing on the experience and expertise of the members of the five founding bodies, has been extended “to provide thought leadership through the development of comprehensive frameworks and guidance on enterprise risk management, internal control and fraud deterrence designed to improve organisational performance and governance and to reduce the extent of fraud in organisations” [87].

One of the best known outputs from COSO’s deliberations is the internal controls framework, known as *Internal Control – Integrated Framework*, first published in 1992 and periodically revised afterwards, with the most recent version being published in 2013. Conceptually the framework has been little altered through the revision process, with the significant changes resulting from the need to link the framework to the requirements for compliance with Sarbanes-Oxley of many of the users of the framework, and from developments in reporting possibilities and requirements. The 1992 framework will be considered superseded at the end of December 2014, but the underlying concepts and principles are considered by COSO still to be fundamentally sound and valid.

The key graphical aid to understanding the COSO Framework takes the form of the COSO Cube.



Figure 7.3 The Internal Control - Integrated Framework, from COSO (1992)

This cube represents the multiple dimensions of internal controls (defined by COSO as “a process effected by an entity’s board of directors, management, and other personnel, designed to provide reasonable assurance regarding the achievement of objectives relating to operations, reporting, and compliance” [88]), showing how activities and processes affect different parts of the organisation.

The top face of the cube shows the three major domains of activity for an organisation, while the side face shows how the organisation might typically be structured into different entities, divisions, departments, units and functions.

The front face presents the five groups of activities – “foundational components” in COSO terminology - involved in effective internal controls, drilling down from the most overarching concepts to the details of daily activities.

In the updated framework, each component is now formally described as consisting of a minimum of two internal control principles that are essential for a fully-functioning and effective internal control environment. Taken in turn, these seventeen principles are as follows:

Control Environment – the context in which processes and internal controls operate. This includes the attitudes and communicated values of senior management – the “tone at the top” – and how commitment and conformity are demonstrated and enforced.

1. Demonstrates commitment to integrity and ethical values
2. Exercises oversight responsibility
3. Establishes structure, authority, and responsibility
4. Demonstrates commitment to competence
5. Enforces accountability

Risk Assessment – internal controls can only be effective and efficient if they are based on a realistic and up-to-date evaluation of all relevant risks, internal and external.

6. Specifies suitable objectives
7. Identifies and analyses risk
8. Assesses fraud risk
9. Identifies and analyses significant change

Control Activities – the range of policies, procedures, activities and mechanisms designed and implemented to address the risks identified as significant to the organisation, within the strategic control environment chosen by senior management.

10. Selects and develops control activities
11. Selects and develops general controls over technology
12. Deploys through policies and procedures

Information & Communication – effective communication must ensure that appropriate information flows within organisations and also to external parties such as customers, suppliers, regulators and shareholders.

13. Uses relevant information
14. Communicates internally
15. Communicates externally

Monitoring – internal control processes should be monitored over time to ensure that their performance is adequate and that opportunities for improvement are seized by exceptions and anomalies being reported and corrective actions designed and undertaken.

16. Conducts ongoing and/or separate evaluations
17. Evaluates and communicates deficiencies

Although meeting commonly-shared concerns about corporate governance and the need for documented control environments, COSO is not without its limitations.

The first of these concerns issues with the necessary human input into control activities. Internal controls can fail because of error or misjudgement, while they can be overridden by collusion amongst employees and failures of the framework of separation of duties, and as a result of coercion by top management. Simply stated, no system of internal controls can prevent fraud or manipulation by sufficiently well-organised, highly-placed and determined officers. If sufficient access rights to information systems exist and monitoring controls can be subverted, inappropriate activities can be carried out.

The second concern relates to the scale and complexity of the model when implemented in its entirety. An article on CFO Magazine in 2011 reported that "One of the biggest problems: limiting internal audits to one of the three key objectives of the framework. In the COSO model, those objectives are applied to five key components (control environment, risk assessment, control activities, information and communication, and monitoring). Given the number of possible matrices, it's not surprising that the number of audits can get out of hand" [89]. In practice, many organisations that implemented the framework did so from an approach of trying to satisfy external auditors and thus documented and validated all of their internal

control activities as audit-relevant. This was time-consuming and very costly and did not necessarily assist management in ensuring that the control environment as documented and attested actually corresponded to their requirements. Subsequent iterations of the process, very often linked to ongoing Sarbanes-Oxley certifications, saw organisations taking a much more selective approach to control definition and reducing the administrative and financial burden significantly.

In 2004 COSO published an updated and extended model (developed by the audit firm PricewaterhouseCoopers, one of whose predecessor firms had largely developed the original COSO model) designed to help organisations to evaluate their *Enterprise Risk Management – Integrated Framework* (ERM). This newer model extends the number of categories (top face) to four, adding the Strategic element to ensure that the framework as implemented considers and incorporates the organisation’s high-level goals, and that these are aligned with and support its mission

The number of framework components (front face) is extended to eight.

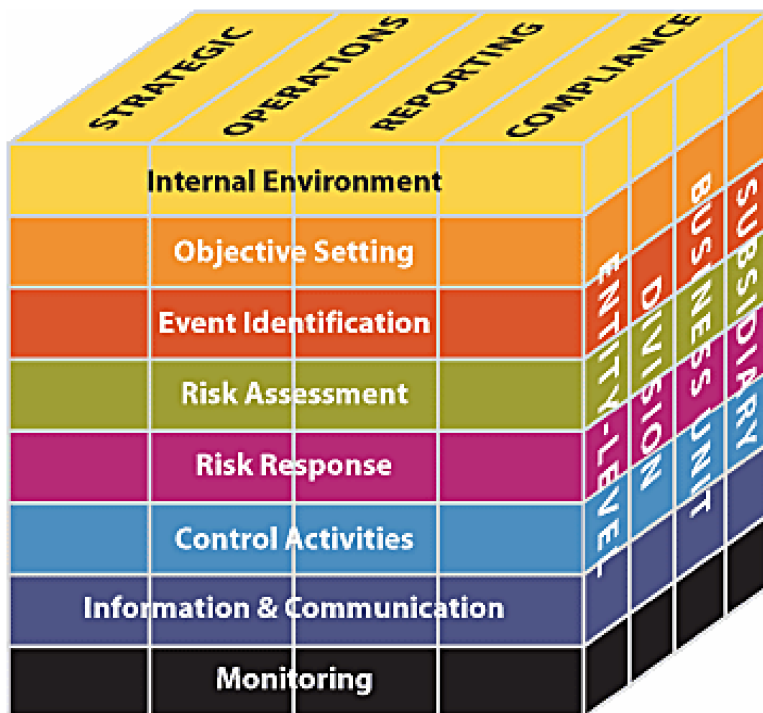


Figure 7.4 COSO Enterprise Risk Management - Integrated Framework (2004)

The previous five components of the Internal Control - Integrated Framework are retained but further developed, while the model is expanded to meet the growing demand for risk management:

Internal environment: The internal environment covers the tone, atmosphere and attitude of an organisation. It sets the basis for how risk is viewed and addressed by an entity's people, including, but not restricted to, risk management philosophy and risk appetite, integrity and ethical values, and the environment in which they operate.

Objective setting: Objectives must exist before management can identify potential events affecting their achievement. Enterprise risk management ensures both that management operate a process to set objectives and that the chosen objectives align with the entity's mission and with its risk appetite.

Event identification: Internal and external events affecting achievement of an entity's objectives must be identified, distinguishing between risks and opportunities. Opportunities for improvement and change need to be passed into management's iterative strategy or objective-setting processes.

Risk assessment: Risks should be analysed, considering likelihood and impact, as a basis for determining how they should be managed, both individually and as a whole.

Risk response: Management selects risk responses – avoiding, accepting, reducing, or sharing risk –in accordance with the organisation's risk tolerances and risk appetite.

Control activities: Policies and procedures are established and implemented to help ensure the risk responses are effectively carried out.

Information and communication: Relevant information is identified, captured, and communicated in forms and timeframes that enable people to perform their duties.

Monitoring: The entirety of enterprise risk management is monitored and modifications made as necessary. Monitoring is accomplished through ongoing management activities, separate evaluations, or both.

Summary

The ERM framework can be subject to the same criticisms as the Internal Control one, in that there remains a great deal of dependence for efficacy on human interventions. In addition, the role and significance of the monitoring component is critical but in practice it is poorly understood and difficult to implement. In 2009 COSO published additional *Guidance on Monitoring Internal Control Systems*, promoting “monitoring based on three main principles:

- Establishing a foundation for monitoring, including (a) a proper tone at the top; (b) an effective organisational structure that assigns monitoring roles to people with appropriate capabilities, objectivity and authority; and (c) a starting point or "baseline" of known effective internal control from which ongoing monitoring and separate evaluations can be implemented;
- Designing and executing monitoring procedures focused on persuasive information about the operation of key controls that address meaningful risks to organisational objectives; and
- Assessing and reporting results, which includes evaluating the severity of any identified deficiencies and reporting the monitoring results to the appropriate personnel and the board for timely action and follow-up if needed” [90].

The final key element in internal controls emphasized by the COSO frameworks is that of the importance of, and indeed necessity for, frequent and effective auditing. It is made clear that both internal and external auditing have roles to play: *internal audit* has the task of assessing the internal control systems implemented by the organisation and contributing to ongoing effectiveness, using its independent

reporting line directly to top management to bring significant issues to their attention and, if necessary, thereby bypass hesitations or reticence on the part of middle-managers or lower-level staff to accept findings and address issues that have been identified; while *external audit* will consider the effectiveness of the internal control environment, either explicitly in the context of a compliance audit or implicitly in the context of a controls-based audit of financial reporting, and will identify and report on weaknesses and areas for improvement directly to senior management and offer independent and focused recommendations for improvement.

7.4 Conclusion

The frameworks and standards discussed above provide a great deal of information to management, mainly in the form of guidelines rather than instructions towards achieving certification. There is a great deal of common ground between the publications, an aspect that should provide a level of comfort to management that they are not striking out into the unknown, nor are they taking the risk of adopting non-standard or poorly regarded measures.

Knowledge of the generally accepted good practices and control objectives in information systems management, and in particular in respect of security, equips management to address a number of challenges in respect of significant IT projects. In particular, they will be able to ask the correct questions of internal project managers, of consultants, and of potential service providers, in respect of organisational, structural, procedural and technical matters, in order to assess whether appropriate levels of service and security can be provided and whether required compliance standards can consistently be attained.

The standards will also allow management to evaluate, in the context of management information and comfort over controls, how any issues, anomalies or exceptions can be identified, classified and reported, and how management can receive comfort over the impact of these events and how the events were handled.

The frameworks and standards can also be compared in order to identify the key areas of focus for data and security management. These can be summarized as the following key concepts:

- Availability management
- Access control
- Vulnerability management
- Patch management
- Configuration management
- Incident response
- System use and access monitoring

Overall, the understanding and application of relevant models and standards is a positive and beneficial step for management. They can be confident that their approach is structured and complete, with no significant risks or concerns omitted, and that they are communicating on a common platform and with a shared language with partners, auditors and regulators.

Chapter 8 The Need for Auditing and Compliance

8.1 Introduction

“Quis custodiet ipsos custodes?”

“Who watches the watchmen?”¹¹

In order to determine how successfully policies and procedures are being followed and are contributing to controlling activities, and in order to be able to demonstrate compliance with internal and external requirements, most organisations will be confronted by the need for audit. In this chapter I discuss the theoretical backgrounds to audit and compliance, describe the major legislative constraints in place, discuss the most significant audit frameworks, and finally set out the specificities of the Swiss situation in respect of the above.

In common with previous chapters, the choices reflect my particular interest in audit and compliance from a structured and documented perspective, with a focus on the requirements of large organisations likely to be subjected to international requirements in respect of compliance with regulations. There is an underlying linkage of operational compliance with financial compliance, as well as an underlying assumption about the size and nature of the organisations in question, their internal cultures, and their resources in respect of these issues.

8.2 Auditing and compliance

These two processes are inextricably linked. Compliance is an inescapable facet of corporate governance and an end in itself in many industries and regulatory contexts, while auditing is a necessary part of the process of achieving and demonstrating compliance. Both internal and external audits have roles to play in the ongoing compliance process.

8.2.1 Internal and external audit

A definition of terms is required. *Internal audit* is work performed by a team internal to the organisation or brought in specifically from outside and mandated by management to, amongst other tasks, evaluate and report upon issues of adherence to policies and procedures, respect for and quality of internal controls, and the management of errors, exceptions and omissions in operations.

¹¹ Juvenal, Satire 6, 346–348

External audit is performed by independent bodies outside the organisation. Their scope can vary widely depending on the reasons for their mandates, but typically they will be seeking to evaluate compliance with a given standard or, very commonly, be reviewing internal controls over and around the information systems in order to provide comfort over the quality of the information used for financial reporting.

Financial auditing, an annual operational necessity for many organisations, has evolved and developed greatly as information systems have become more significant. Traditionally it has been possible to audit a company's financial position and performance on a transactional basis, looking at the substance of its transactions and activities and obtaining assurance over the integrity and accuracy of management's assertions through the examination of evidence relating to transactions. As organisations have become increasingly dependent on their information systems for all aspects of their business and reporting, however, from production planning and inventory management to processing and recording online sales and calculating depreciation and revaluations automatically, there is less hands-on human intervention in some processes and less traditional formal evidence to be consulted. At the same time, the size of some businesses and organisations has become so great that traditional methods of auditing individual business cycles are no longer possible. While a company's activities in respect of, say, fixed asset purchases can perhaps still be audited on a transactional basis, identifying the major acquisitions and consulting the invoices, the same company's sales cycle, which might perhaps involve the sale of a hundred million small widgets, cannot be approached in the same way.

The way that auditing practices have evolved in response to this is to develop means of obtaining comfort based on internal controls around key business processes. If the auditors can be sure, based on evaluation and targeted testing of the operation of controls, that internal controls are robust and effective, they can use that comfort to draw conclusions about the overall operation of the process.

It is useful to define clearly what is meant by audit comfort. Audit comfort is the support for audit conclusions that is obtained from the examination of evidence, the evaluation of internal controls, and the assessment of relevant information provided by third parties. It is the necessary basis for audit opinions, and limitations in the acquisition of sufficient comfort will lead to reservations and qualifications in any opinion given. It should be noted that audit comfort, particularly in respect of the design and operation of internal controls, will be drawn from the same sources as management's own confidence that processes are operating consistently and in a controlled manner. Management can therefore implement control activities knowing that the results of these activities will provide two layers of benefit: daily operational assurance; and evidence with which the auditors will be satisfied.

Given the importance of information systems in the operations of businesses and the generation of financial reports, it is therefore the case that in order to conclude on processes and related internal controls, it is necessary to audit the controls around the management and operation of the relevant information systems.

The fundamental principles behind audit and compliance for an organisation are to ensure that it has properly addressed the key elements necessary to demonstrate that it is respecting requirements. This consists of:

- The identification of all relevant requirements, internal, external, legal or contractual
- The implementation of strategies, policies, procedures and activities to address the requirements

- The performance of monitoring to ensure that the detailed measures are operating consistently and effectively and that anomalies and exceptions are being identified and handled appropriately.

Although not every organisation is required to have an internal audit function, and many do not have the resources to finance such a group, it can be beneficial to have a function and for it to liaise closely with the external auditors. This can help to avoid surprises when the external auditors visit, by having verified the effectiveness of control activities and the existence of appropriate evidence of their operation, and to ensure that staff and management are fully conversant with the concepts of risk and control, through briefing, training and inspection, before having to explain things to the external team.

A systematic approach to compliance is necessarily particularly important in environments where the impact of a control failure could be severe. From the perspective of a service provider, it is inadvisable to wait for an annual external audit to reveal problems in the design or operation of controls during the past year, because such control failures may have impacted multiple customers. Key controls must be identified and monitored so that any potential issues can be investigated and, if necessary, rectified, in a timely manner. Such control monitoring often forms part of the scope of an internal audit function, as they have the necessary skills and independence. It could, however, be performed by an internal controls team, operations group, or security team, depending on the range and nature of controls in question.

8.2.2 Governance, Risk and Compliance

Governance, Risk and Compliance (GRC) has become a very hot topic in the corporate world over the last decade, particularly as a consequence of the introduction of the Sarbanes-Oxley legislation (see below). Corporate management found itself obliged to take a structured and measured overall approach to control issues, ensuring that the three necessary pillars are integrated and monitored to avoid conflicts, wasteful overlaps and gaps.

It is useful to define each of the three key concepts.

8.2.2.1 Governance

Governance describes the management approach used by senior executives to direct and control the entire organisation, typically involving a combination of management information and hierarchical management control structures. The objective of governance activities is to ensure that the management information reaching the appropriate levels of the hierarchy is sufficiently complete, accurate and timely to enable decision making. This is reflected in the COSO Cubes presented and discussed in Chapter 7.

8.2.2.2 Risk management

Risk management is the process of identifying, evaluating and responding to risks relevant to the organisation's objectives. Risk management is a wide-ranging and common activity in all facets of corporate activity, but in the context of GRC it is particularly concerned with external legal and regulatory compliance risks.

8.2.2.3 Compliance

Compliance, as discussed above, is the objective and process of conforming with stated requirements. These might be internal or external, regulatory or contractual, mandatory or voluntary.

There is nothing new in the individual concepts behind GRC, but the value to organisations has come from the creation of structures and frameworks that can be applied across an organisation in a structured way and which emphasise that each of the three disciplines consists of four basic components: strategy, processes, technology and people. In addition, three basic rules can be identified: the organisation's risk appetite, its internal policies and external regulations. In applying these disciplines, components and rules to the organisation's objectives, an organisation can work towards its GRC objectives of ethically correct behaviour and improved efficiency and effectiveness [91].

Of specific interest in the field of information systems is the sub-branch of GRC that considers IT: ensuring that the IT organisation supports the current and expected future needs of the organisation and complies with relevant constraints. IT GRC can itself be split into a number of key capabilities, including, depending on the particular package or methodology that has been selected:

- IT governance and policy management;
- IT asset tracking;
- IT risk assessment and response;
- IT control implementation;
- IT regulatory compliance and reporting;
- IT incident and threat management;
- IT vendor risk and performance management;
- business continuity planning; and
- ongoing IT auditing [92]

The overall principles of GRC are reflected in existing standards. For example, section 2.2 ISO/IEC 38500: 2008, "Corporate governance of information technology" defines the responsibility of executive management:

"Directors should govern IT through three main tasks:

- a) Evaluate the current and future use of IT.
- b) Direct preparation and implementation of plans and policies to ensure that use of IT meets business objectives.
- c) Monitor conformance to policies, and performance against the plans" [93]

COBIT 5 (see Chapter 7) also includes the key principles of GRC. The Governance domain contains five governance processes within which Evaluate, Direct and Monitor (EDM) practices are defined:

- 01 Ensure governance framework setting and maintenance
- 02 Ensure benefits delivery

- 03 Ensure risk optimization
- 04 Ensure resource optimization
- 05 Ensure stakeholder transparency.

Risk management is covered by both EDM03 Ensure risk management, and by the Management process Align, Plan and Organise 12 (APO12) Manage risk

Compliance is covered by a process within the Management Monitor, Evaluate and Assess domain: *MEA03 Monitor, evaluate and assess compliance with external requirements.* [3]

8.2.3 Summary

Audit practices and methodologies are extremely mature and the standard principles and terminology are widely accepted. There is an overwhelming focus on tight management controls that are designed to respond to specific risks and vulnerabilities.

8.3 Legislative constraints

Although there are good operational and governance reasons why organisations would voluntarily mandate or submit to an audit process, in reality the major motivation for undergoing audit is to comply with legal requirements. Depending on the nature of an organisation's activities, its status, ownership and location, it will find itself subjected to different regimes with different reporting, auditing and compliance requirements. Below I set out a selection of the most significant pieces of legislation from a global perspective, based on my experience of the impact they have had on organisations.

8.3.1 Sarbanes-Oxley Act

As a comprehensive response to significant and eye-catching cases of financial reporting fraud in 2001 and 2002, the US Government determined that was necessary to impose upon publicly quoted companies (and some other organisations) more stringent requirements in respect of corporate financial reporting. Among the measures included in the Sarbanes-Oxley Act of 2002 (SOX), the most significant are:

- Public company CEOs and CFOs are required to certify the effectiveness of their internal controls over financial reporting (ICOFr) on a quarterly and annual basis, thereby making themselves personally responsible for the implementation and operation of effective internal controls within their organisation, for being appropriately informed of significant errors and exceptions, and handling such exceptions in a suitable manner.
- Management is required to perform an annual assessment of their ICOFR.
- External auditors are required to express an opinion on the effectiveness of management's ICOFR as at the company's fiscal year end.

This has had a significant effect on the way that external financial audits were performed and documented. Although the financial audits of large companies were already largely systems – controls – driven, because of their size and for the reasons explained above, for the first time outside of specific industries the external auditors were obliged to give an opinion on the quality and reliability of internal controls.

Previously internal controls could be reviewed and relied upon for audit comfort purposes if suitable or ignored if not, with the audit approach reverting to a more traditional substantive one and the failings of the internal controls perhaps being mentioned in a management letter, but there was no requirement to go further.

SOX was also responsible for the creation of the Public Company Accounting Oversight Board (PCAOB) which was charged with establishing audit standards and inspecting the work of audit firms to ensure that the auditors were in compliance with the legislation and standards. Until then, the profession had been self-regulated.

The PCAOB is overseen by the Securities and Exchange Commission (SEC), of which the mission is to “protect investors, maintain fair, orderly, and efficient markets, and facilitate capital formation” [94]. It is the responsibility of the SEC to:

- “interpret and enforce federal securities laws;
- issue new rules and amend existing rules;
- oversee the inspection of securities firms, brokers, investment advisers, and ratings agencies;
- oversee private regulatory organisations in the securities, accounting, and auditing fields; and
- coordinate U.S. securities regulation with federal, state, and foreign authorities” [94].

The PCAOB thus benefits from significant political and legal support in the performance of its activities.

8.3.2 Auditing Standards

Since December 2003 the PCAOB has published standards for the auditing of financial statements. The most significant of these from an IT controls perspective are *Auditing Standard No. 2: An Audit of Internal Control over Financial Reporting Performed in Conjunction with an Audit of Financial Statements (AS2)*, published on 9 March 2004 and the subject of a *Policy Statement Regarding Implementation* issued on 16 May 2005 and a *Report on the Initial Implementation* issued on 30 November 2005; and its successor standard *Auditing Standard No. 5: An Audit of Internal Control Over Financial Reporting That Is Integrated with An Audit of Financial Statements (AS5)* published on 24 May 2007.

AS2 emphasised the importance of Information Technology General Controls (ITGCs), explaining the need for focus on the effectiveness of ITGCs thus: “Some controls ... might have a pervasive effect on the achievement of many overall objectives of the control criteria. For example, information technology general controls over program development, program changes, computer operations, and access to programs and data help ensure that specific controls over the processing of transactions are operating effectively” [95].

The IT Governance Institute developed a set of IT Control Objectives for Sarbanes-Oxley that became the de facto industry standard for the ITGC needed to achieve the requirements of SOX. This standard included control objectives and recommended supporting control procedures in the following areas [96]:

- Program Development and Program Change

- Acquire or develop application system software
- Acquire technology infrastructure
- Develop and maintain policies and procedures
- Install and test application software and technology infrastructure
- Manage changes.
- Computer Operations and Access to Programs and Data
 - Define and manage service levels
 - Manage third-party services
 - Ensure system security
 - Manage the configuration
 - Manage problems and incidents
 - Manage data
 - Manage operations

These IT controls underpin all the processes and application controls that refer to or rely upon the organisation's information systems.

AS5 changed the emphasis to a more narrowly risk-based approach, enabling most reporting companies to reduce the scope of their compliance activities. In the first years of SOX many companies adopted a very prudent approach to compliance, documenting and testing most of the internal controls in operation. The publication of AS5 allowed them, under guidance from their corporate advisors and auditors, to revisit their control matrices and reclassify their individual control activities as critical, key, or non-key, and focus testing and verification efforts only on the most important ones.

8.3.3 The impact of SOX on outsourcing and cloud computing

The focus of SOX is on the effectiveness and reliability of a company's financial reporting process, which is driven by its information systems. Systems that have a material impact on financial recording and reporting process are thus in scope for SOX and both the detailed controls in place over those systems and the ITGCs in place over the wider IT environment should therefore be included in compliance efforts. This applies to both systems and platforms maintained in-house and those that are outsourced.

Whether or not any particular system falls into scope for SOX is ultimately a question for corporate management, to be resolved after developing a cartography of applications and platforms and seeing how these support the audit-significant business processes. In the case of cloud services, the question will revolve around the precise nature of the service provided and the impact that this has on audit-significant processes.

8.3.4 PCI DSS (Payment Card Industry Data Security Standard)

For organisations that handle payment card data, a fundamental concern is compliance with the Payment Card Industry Data Security Standard (PCI DSS). The standard sets out twelve requirements for any business that stores, processes, or transmits payment cardholder data [97], and the scope of responsibility is such that organisations are also required to ensure that their contracts with third-party service providers include PCI DSS compliance should their service providers be storing or processing cardholder data on their behalf.

In a cloud environment, the organisation and their service provider should clearly define and agree their relative responsibilities, making it clear whether any are shared or can be imputed to one party. In order to do this, it is fundamental to be able to map information flows and identify where storage and processing occur.

8.3.5 HIPAA (Health Insurance Portability and Accountability Act)

The US Health Insurance Portability and Accountability Act (HIPAA) of 1996 was designed to improve the portability and continuity of health insurance coverage, protecting health insurance coverage for workers and their families when they change or lose their jobs and requiring the establishment of national standards for electronic health care transactions and national identifiers for providers, health insurance plans, and employers. The security and privacy rules associated with HIPAA emphasize in particular the obligations of health organisations to ensure that individually identifiable health information is adequately protected when entrusted to business associates (e.g., third-party service providers). Organisations that have chosen to outsource data storage or processing should therefore be able to demonstrate compliance on the part of their service provider with the provisions of the Act.

8.3.6 Summary

These three major pieces of United States legislation, accompanied by the guidance issued by auditors and other bodies, have had a significant impact on the way that many organisations around the world approach both the management of systems and information and the way they view auditing. The obligation to define and document policies, procedures and controls and to have the quality of these measures demonstrated through independent audit brought about a significant change to the operating environment for many companies, previously unaffected by stringent compliance requirements, who fell into scope for such legislation. In Switzerland in particular, with a traditional approach to internal controls in many industries that was far less formal than the requirements of Sarbanes-Oxley demanded, the move towards compliance often required a great deal of effort and resources.

8.4 Audit frameworks and approaches

In Chapter 7 I discussed the major internal controls and audit frameworks. It is important to emphasise the differences encountered between auditing in-house and auditing a third party, because although the overall objectives are the same and similar techniques will be employed, there are a number of logistical, procedural and organisational elements that make such audit projects more difficult to carry out.

Typically in an outsourced environment there are a number of ways for the organisation purchasing the services to acquire at least a reasonable level of control and assurance. These can be described as obtaining

comfort from audit procedures performed by the organisation itself or by an audit entity specifically mandated to audit on its behalf; by obtaining a service auditor's report for the environment in question; or by implementing, performing and monitoring the success of a series of internal control activities.

8.4.1 Self-performed audit procedures

The first of these methods is the most direct and is identical in form and process to the use of internal audit functions to evaluate and report upon business operations that are carried out in-house, as well as the use of the findings of external audits, where internal controls are often evaluated in the context of a controls-based audit. The audit team works with the organisation's management to define scope and timeframe, documents the business procedures, identifies the key control activities, assesses these for both design effectiveness (whether as designed the control activities do address the risks to which they relate and the control objectives to which they correspond) and operating effectiveness (whether the control activities are actually working as designed, with necessary inputs being received, the control being performed, and appropriate conclusions being drawn and actions undertaken). Based on the results of these audit activities, management will be in a position to draw informed conclusions on whether or not internal controls are working as intended.

The right to perform such audits will need to be specified in the contract for service provision.

The difficulties of performing such reviews in an outsourced environment are many. Firstly, the organisation may not have an appropriately skilled, experienced or available audit department – it is frequently the case that internal audit functions tend to have greater expertise in, and requirements to focus on, financial management issues such as the use and disposal of assets, Value For Money, and purchase and inventory management. Detailed technical skills in the field of information systems are more likely to be found in larger and more IT-dependent organisations.

Secondly, it might be prohibitively expensive to mandate a third party to perform the necessary audit on their behalf.

Thirdly, it might be logistically difficult to identify all the areas in which processing or data storage or IT management tasks are performed at the service provider, especially if multiple sites in multiple locations are used. This will add to the complexity of scoping and scheduling such an audit. Audit scoping is in itself an issue that requires experience and expertise, but it is essential in respect of the quality and relevance of the results that will be obtained and the resources and timeframes necessary to perform the audits. Management should be able to determine which systems, applications and datasets are particularly relevant for their purposes and ensure that these are included in the scope, while not necessarily including all of their information assets. This is frequently performed on the basis of a cartography of systems and data, identifying the most critical and the most sensitive applications and data and determining where these are located and used.

Fourthly, but far from being the least significant issue, there is the impact of being audited on the service provider. In order to respond to audit questions, staff and documentation need to be made available, and then resources need to be provided to assist with the testing of individual controls, extracting system information and reports and explaining their contents. This can have multiple impacts on the service provider, taking up human and technical resources, distracting staff from their daily activities, and using office space. If a service provider were to allow all of its clients to pay audit visits independently and at

times that suited them, it is not difficult to imagine that this could cause significant disruption to the provider's activities.

8.4.2 Service audit approach

Many service providers thus prefer the second method of allowing their clients to obtain comfort over internal controls: the service audit report. In this situation, the service provider itself mandates an external auditor to review its internal controls and provide an opinion on the effectiveness and efficiency of these controls. This report is then made available to the service provider's customers who can incorporate its findings into their own evaluation of the control environment.

The advantages of such an approach for the service provider are clear. They meet their requirements to provide reliable information about their control environment, while avoiding the inefficiencies of having multiple audit teams visiting their premises and having to provide repeated briefings, explanations and copies of documentation. If the audit partner chosen has a size and structure that allows consistency of team membership and a minimum of rotation, there will also be the advantage of familiarity year to year with business practices and documentation standards as applied by the service provider, which will also increase the efficiency of the audit process.

The report issued by the auditor can take many forms, but for many years the de facto international standard to be followed was the Statement on Auditing Standards No. 70: Service Organisations, commonly abbreviated as SAS 70.

This standard specifies two kinds of reports, named Type I and Type II. A Type I service auditor's report includes the service auditor's opinion on the fairness of the presentation of the description of controls that had been placed in operation by the service organisation and on the suitability of the design of the controls to achieve the specified control objectives (thus corresponding to the concept of design effectiveness as discussed above). A Type II service auditor's report includes the information contained in a Type I service auditor's report and also includes the service auditor's opinion on whether the specific controls were operating effectively during the period under review. Because this opinion has to be supported by evidence of the operation of those controls, a Type II report therefore also includes a description of the service auditor's tests of operating effectiveness and the results of those tests.

SAS 70 was introduced in 1993 and effectively superseded in 2010 when the Auditing Standards Board of the American Institute of Certified Public Accountants (AICPA) restructured its guidance to service auditors, grouping it into Statements on Standards for Attestation Engagements (SSAE), and naming the new standard "Reporting on Controls at a Service Organisation." The related guidance for User Auditors (that is, those auditors making use of service auditors' reports in the evaluation of the business practices or financial statements of organisations making use of the facilities provided by service providers) would remain in AU section 324 (codified location of SAS 70) but would be renamed Audit Considerations Relating to an Entity Using a Service Organisation. The updated and restructured guidance for Service Auditors to the Statements on Standards for Attestation Engagements No. 16 (SSAE 16) was formally issued in June 2010 and became effective on 15 June 2011. SSAE 16 reports (also known as "SOC 1" reports) are produced in line with these standards, which retain the underlying principles and philosophy of the SAS 70 framework. One significant change is that management of the service organisation must now provide a written assertion regarding the effectiveness of controls, which is now included in the final service auditor's report.

Internationally, the International Standard on Assurance Engagements (ISAE) No. 3402, Assurance Reports on Controls at a Service Organisation, was issued in December 2009 by the International Auditing and Assurance Standards Board (IAASB), which is part of the International Federation of Accountants (IFAC). ISAE 3402 was developed to provide a first international assurance standard for allowing public accountants to issue a report for use by user organisations and their auditors (user auditors) on the controls at a service organisation that are likely to impact or be a part of the user organisation's system of internal control over financial reporting, and thus corresponds very closely to the old SAS 70 and the American SSAE 16. ISAE 3402 also became effective on 15 June 2011.

The importance of understanding the related guidance for user auditors (which also applies, by extension, to user management) is critical. When a service audit report is received, the reader need to follow a number of careful steps in order to be certain that the report is both useful and valid before beginning to draw any conclusions from it.

These steps include:

1. Confirming that the report applies to the totality of the period in question. This is of particular importance when using a service audit report in the context of obtaining third-party audit comfort for a specific accounting period, but also applies to more general use of the report. If management's concern is over the effectiveness of internal controls for the period from 1 January to 31 December 2012, say, and the audit report covers the period from 1 April 2012 to 31 March 2013, how valid is it for management's purposes and how much use can they make of it? In the simplest of cases, a prior period report will also be available that will provide the necessary coverage, but in other circumstances this may not be the case and management will have to turn to other methods to acquire the comfort that they need. These could include obtaining representations from the service manager that no changes had been made to the control environment during the period outside the coverage of the audit report and that no weaknesses in internal control effectiveness, either design or operational, had been identified during that period. Other methods could involve the performance and review of internal controls within the client organisation, a subject to which I will return below.
2. Confirming that all the systems and environments of operational significance to the organisation were included in the scope of the audit report. Management will need to have an understanding of the platforms and applications that are being used for their purposes at the service provider, including operating systems, databases, middleware, application systems and network technologies, and be able to confirm that these were all appropriately evaluated. Very often in the case of large service providers a common control environment will exist, under which they apply identical internal control procedures to all of their environments: if the auditors have been able to confirm that this is the case, then it is not inappropriate for them to test the internal controls in operation around a sample of the operating environments, and management can accept the validity of their conclusions without needing the confirmation that their particular instance of the database, for example, had been tested. Should the coverage of the audit not meet management's requirements, again management would need to evaluate the size and significance of the gap and consider means by which they could obtain the missing assurance.
3. Understanding the results of the work done and the significance of any exceptions or weaknesses noted by the auditors. Generally speaking, if the work has been performed to appropriate standards and documented sufficiently, and if the conclusions drawn by the auditors are solidly based on the evidence,

this step should be reasonably straightforward. Management should, however, guard against skipping to the conclusion and, if the report contains it, the service provider's management attestation, and blindly accepting the absence of negative conclusions as being sufficient for their purposes. Each weakness in the design or the operation of controls should be considered, both individually and cumulatively, to identify any possible causes for concern. This is because it is possible that weaknesses identified during the audit, considered to be insignificant within the overall framework of internal controls could potentially be significant in respect of the specific circumstances (use of systems and combination of technologies, for example) of one particular customer.

8.4.3 Other procedures

Should such an independent audit report not be available, or only provide partial coverage of the control environment or the period in question, and should it in addition be impossible or impractical to the organisation to perform their own audit procedures at the service provider's premises, a third way of obtaining comfort over operations is needed. This method can be summarized as consisting of two groups of activities: internal controls performed within the organisation over the activities of the service provider, or audit procedures designed to evaluate the completeness, accuracy and integrity of the service provider's processing and outputs.

Firstly, internal controls can be designed that allow local management to monitor and evaluate the activities of a third party. These can take the form, for example, of procedures to track the responses of the service provider to requests for changes: if there is a process by which the customer asks the service provider to change access rights to an application or a datastore to correspond to the arrival or departure of a member of staff, management can track these change requests and the responses of the service provider in order to ensure that the correct actions have been undertaken.

Very often the contract terms with the service provider will include regular meetings and reporting mechanisms through which the provider will present status up-dates, usually in the form of progress against KPIs and lists of open points, and management can ask questions and ensure that everything is under control. These meetings can form the central points of control activities for management, as a structured means of ensuring that they are monitoring the performance of the service provider in a regular and consistent way, observing long-term trends and identifying anomalies.

Secondly, audit procedures can be designed along similar principles to the above, based on the expected results from the service provider's activities. An example of this would be extracting at the period end a list of user access rights to the organisation's applications and comparing these to expectations, expectations based on management's understanding of the access requirements and on the instructions given during the year to create, modify or delete access rights. If the rights correspond, this can provide a layer of assurance to management that both the service provider's internal procedures and controls are operational and that access to their applications and data is being appropriately managed.

With an appropriately selected range of internal controls and audit procedures, management can therefore obtain a certain level of comfort over the existence and quality of the control environment in place at the service provider.

8.5 Specifically IT based audit frameworks

8.5.1 SysTrust

This provides for the audit of controls based on predefined principles and criteria in respect of security, notably the traditional triad of availability, confidentiality, and processing integrity. This framework is designed to allow auditors and users to attest to the reliability of any system.

8.5.2 WebTrust

This framework is similar conceptually to SysTrust above, but is specifically designed to online and e-commerce systems.

8.5.3 ISO 27001

As discussed above, this provides for the systematic and structured audit of an organisation's Information Security Management System (ISMS), on the basis that this ISMS has been documented and clearly structured.

8.6 Summary

Obtaining audit comfort over services and facilities over which one does not have direct control or clear visibility is possible and frameworks and standard practices exist for doing so. Such efforts require a high degree of experience and expertise, however, as well as clarity and commitment from all parties. In the absence of these qualities, alternative methods of acquiring a level of assurance must be designed and implemented.

8.7 The Swiss context

Although the principles of the Swiss approach to data management, security and compliance are not significantly different to those in place in other countries, its approach to data protection for individuals, and in particular in respect of certain types of data, makes it of particular interest.

The use of cloud computing services presents a number of difficulties in relation to regulations and compliance.

8.7.1 Legal requirements

In the domain of data protection, aspects such as the location, transmission, recovery of data, the ownership of and responsibility for data, and data processing itself are all subject to strict regulation, particularly when performed by third parties.

The laws in place apply to both the data owner and the service provider and this can cause both strategic and operational difficulties for both and hence present a barrier to the uptake and development of such services and agreements.

8.7.1.1 LPD

The most significant block of legislation applicable to this domain in Switzerland is the Law on the Protection of Data (Loi fédérale sur la protection des données {LPD, RS 235.1}) [98] that came into force on 19 June 1992, and the associated Ordonnance on the Law on the Protection of Data (Ordonnance sur la loi fédérale sur la protection des données {OLPD, RS 235.11}) [99] that came into force on 14 June 1993.

Additional definitions, instructions and measures are contained in other legal texts such as the Code des Obligations (CO, article 328b) [100], which sets out the limitations on the nature of data that employers can hold and process in respect of their employees.

The LPD is monitored and enforced by two bodies. The first of these is the Federal Officer for Data Protection and Transparency (le Préposé fédéral à la protection des données et à la transparence) who is appointed by the Conseil Fédéral and administratively a member of the Federal Department of Justice and Police. This official has the authority to perform investigations, offer advice, make recommendations and promote public information.

8.7.1.2 Federal Commission for Data Protection

The second is the Federal Commission for Data Protection (La Commission fédérale de la protection des données), which is a body for arbitration and the processing of appeals mandated to deliberate on the implementation or not of the recommendations made by the Federal Officer and to judge appeals against the decisions made by federal bodies in respect of data protection.

The fundamental objective of the LPD is to protect the integrity and the fundamental rights of individuals and entities whose data are being processed, whether this processing is being performed by individuals, private bodies or federal entities.

Although wide-ranging in its principles, the scope of application of the LPD is in practice subject to a number of restrictions.

It does not apply, for example, to:

- Personal data that an individual is using exclusively for personal purposes and is not communicated to any third parties
- The workings and deliberations of the federal parliament and parliamentary commissions
- Any ongoing proceedings in respect of civil and criminal cases and of international legal requests, and of public and administrative law, with the exception of the initial hearings of administrative procedures
- Public registers related to legal reports of private law
- Personal data processed by the International Committee of the Red Cross

The LPD does set out a number of responsibilities of the part of data holders to ensure that the rights of the subjects of the data are not infringed. These include:

- The requirement to declare the existence of the datasets, who can use them and who can be contacted for access; these declarations are maintained in a publicly accessible register
- The requirement to ensure that data are not used or abused such that the subjects have their integrity or reputation compromised
- The requirement to implement technical and organisational measures to secure the data appropriately against accidental or unauthorised destruction, accidental loss, technical errors, the falsification, theft or illicit use of the data, or the unauthorised modification, copying or access to the data
- The requirement to ensure that when data are passed to a third party for processing or storage, the third party only performs the specific processing that it has been contracted to do and that it can guarantee the security of the data
- The requirement to respect the right to access and to the rectification of data, although these rights are limited in certain restricted contexts;
- The requirement to respect the limitations in place over the transfer of data outside Swiss territory: these restrict the nature of the data that can be transferred across borders and oblige data owners to ensure that local legislation in the countries to which the data are been exported offers the same protection to the subjects of the data as Swiss law, or to provide contractual details ensuring that the use of the data is appropriately protected, that only data necessary for the fulfilment of a contract and relating to the contracting parties is involved, that the data have been made public and there has been no formal opposition to the processing of the data abroad, or that the data transfer is between subsidiaries of the same organisation and the same adequacy of security and privacy can be assured
- The procedures for certification of the processing of data, with numerous facets requiring documentation and validation.

8.7.1.3 *Relevance to data and the cloud*

These requirements have numerous impacts on the field of cloud computing. Cloud service providers based in Switzerland are under the obligation to guarantee the security of the data they are holding and processing, in accordance with art 7 LPD. In order to do this, they need to acquire a certificate granted according to the requirements of the Ordonnance on certifications in respect of the protection of data (“Ordonnance sur les certifications en matière de protection des données”, (OCPD) 235.13)) of 2007 [101].

This certificate removes from the service provider the obligation to declare the nature and contents of each of the files being processed to the Federal Officer, thus making data processing and storage activities administratively feasible.

In order to obtain the necessary certificate, the service producer is supposed to undergo a traditional audit procedure designed to evaluate the organisation itself, its data protection procedures, its information

systems, and its data processing services, similar in nature and scope to the type of review typically performed under SAS70. The technical difficulties involved in performing such an audit, however, along with the costs and timescales involved, the perceived inconsistency of the process and a vagueness in the current legislation, have led the Federal Officer to suspend the requirement for the evaluation of the systems and services in place until an appropriate solution can be found [102].

Swiss companies that regularly process data need, according to art 5 LPD, to confirm the accuracy of those data. They are required to take necessary steps to allow the people or entities concerned by the data to have modified or deleted any data that they feel are inaccurate or incomplete.

In practical terms, this means that any company operating a Client Relationship Management system in the cloud should be able to allow either the data owners or the subjects of the data to inspect, review and request changes to the data at all times.

A result of the most recent update of the LPD was that the requirement to declare individual files was replaced by the requirement for diligence. This means that organisations using cloud services hosted or partially operated abroad are required to confirm conformity with the general principles set out in the LPD (art 4) and ensure the adequacy of the relevant legislation in the host country. In essence, any organisation wishing to outsource its data abroad is required not only to guarantee the security of those data but also guarantee that the local legislation in force will ensure an adequate level of data security.

A list of countries considered to provide adequate legislation is available from the website of the Federal Officer [103].

If the safeguards in place are not adequate, however, various solutions do exist. Should the outsourcing be located in a country where data security legislation is insufficient to ensure the confidentiality of personal data, art 6 LPD does permit the outsourcing to go ahead if sufficient guarantees, in the form of robust contract terms in line with international standards, can be provided.

As discussed above, American companies are required to respect the data protection principles enshrined in the Swiss Safe Harbor Framework and to be registered with the US Department of Commerce. As a result, the organisation requiring the service should ensure that the potential service provider is indeed registered.

8.7.1.4 Explanatory guidance

In October 2011 the Preposé published an explanatory document in respect of cloud computing [104]. The objective of this document was to provide some standard definitions and a shared vocabulary for organisations considering using cloud services, to set out the major risks, and to discuss how these impacted the protection of data.

The document distinguishes between private, public, hybrid and community clouds, and also discusses the different offers possible: Infrastructure as a Service, Platform as a Service, and Software as a Service.

It then outlines the key risks related to cloud computing. For the authors, the most significant risks are:

- The loss of control over data once data owners no longer know precisely where data are stored

- The absence of separation and isolation of data if platforms and infrastructure are shared with other organisations and managed by a third party
- Non-respect of legal requirements, especially if data are copied many times and stored in numerous different jurisdictions
- Access to data by foreign authorities if the data are stored in or accessible from a country with legal procedures that allow such measures
- Becoming captive clients of a service provider if customised or non-standard technologies are used.

In addition three general risks in relation to the storage and processing of data are listed:

- The loss of data through theft, deletion, modification or system crash
- Network or system breakdowns leading to the unavailability of data, with consequences for the ongoing operations of the organisation
- The misuse of data by users, administrators or intruders.

Four key areas of concern are then spelled out for data owners, referring to the legislation in place and the legal responsibilities that have been created.

1. Article 10a of the LPD specifies that cloud computing is subject to the same legal restrictions as traditional outsourcing. Data processing can be transferred to a third party as long as this is allowed by law or contract and as long as two conditions are met: only those elements of processing that the service user has the right to perform are actually performed, and that there is no legal or contractual obligation in respect of maintaining secrecy that forbids such outsourcing. The party requesting the services should, in particular, ensure that the third party can guarantee the security of the data.
2. The third party should guarantee the security of the data according to article 7 of the LPD and articles 8ss and 20ss of the OLPD. These clauses insist that personal data should be protected by technical and organisational measures against any unauthorised processing and that the confidentiality, integrity and availability of the data are ensured.
3. As the use of cloud services often involves the communication of data across borders, personal data may find itself transferred to countries in which the level of data protection is lower than that in Switzerland. Data owners should be aware of article 6 of the LPD, which insists on confirming that other territories to which data might be transferred can guarantee an adequate level of data protection before any transfers of sensitive data can take place.
4. Data users also need to confirm the right of access of the subjects of data to that data (article 8 of the LPD) and the right to delete or correct data (article 5 of the LPD). The fact of having outsourced or contracted cloud services does not free the data owners from their obligations under these regulations.

The conclusion is that the choice of service provider needs to be the result of a complete analysis of the organisational, legal and technical risks that might apply. If the data owner has any doubts over the ability of the service provider to comply with the regulations in place, the outsourcing should not be undertaken.

8.7.2 Audit Standards

8.7.2.1 *Theoretical background*

From an audit perspective, the Swiss Audit Standard (Norme d'Audit Suisse 890 {NAS 890}) that came into effect on 17 December 2007 [105] introduced for the first time a systematic and documented set of requirements for the documentation and attestation of internal controls. The legal requirement for companies to be able to demonstrate the existence, if not the effectiveness, of such internal controls, and have this confirmed by external auditors, was introduced in the Code des Obligations in 2007 (art 728a CO).

According to articles 716a, 662a and 957 of the Code des Obligations, businesses that are signed up to the "registre de commerce" are required to adhere to prescribed frameworks for accounting and the presentation of accounts. In addition, companies that are subject to the requirements of the "contrôle ordinaire" (the requirement for an independent scrutiny of their accounts) set out in art 727 CO have the additional requirement to implement and maintain an internal control system, the existence of which should be verified by the external auditor (art 728a CO).

The objective of such a system is to guarantee, at least in principle, the keeping of appropriate books and records for accounting purposes and the creation of an adequate financial report.

NAS 890 defines the conditions for the verification of the internal control system as well as the minimal requirements for confirming its existence.

The scope, extent and quantities of the inquiries and tests to be performed by the auditor in order to verify the existence of the internal control system within an organisation must necessarily take account of the size, complexity and risk profile of that organisation.

In contrast to the detailed evaluations and tests required by auditing standards in respect of certain financial aspects of the annual accounts, the procedures required in respect of internal controls are rather light and focus on demonstrating its existence rather than its effective and efficient operation.

This is in clear contrast to the requirements of the Sarbanes-Oxley legislation as discussed above, under which not only the design effectiveness but also the operating effectiveness of the internal controls needs to be demonstrated, verified, and reported upon as a specific and separate element of the independent financial reporting process.

8.7.2.2 *Practical application*

In order to confirm the existence of an internal control system in respect of art 728 CO, the auditor will typically inspect documentation of processes and procedures, of risk assessments, of control design, and of the monitoring of the operation of control activities.

Not every business in Switzerland is subject to these requirements. The specific criteria for qualification are set out in the Code des Obligations and concern the size, turnover and staffing levels of the company over a

period of years, in order to ensure that only companies of a certain sustained size and level of activity are concerned by the requirements to introduce and maintain a formal approach to internal controls.

In order to conclude that an internal control system is indeed in place, the auditor needs to consider the following points:

- That the ICS is clearly defined and documented
- That the ICS corresponds to the real risks and the commercial activities of the company
- That the staff in positions of responsibility are familiar with the ICS
- That the ICS in place does correspond to the ICS designed by management
- That there is a culture of awareness of internal controls within the company

Typically the last point is the most difficult to ascertain because of the common gap between what is documented formally, often the result of working groups meeting at a high level within the organisation, and what permeates down throughout the organisation.

There is also frequently a conceptual gap or a lack of common vocabulary within organisations, between management and staff and also between auditors and employees. Without a clear understanding of what is meant by terms such as “control objective” and “control activity”, there is always a risk of confusion.

Each annual audit will end with the submission of two reports by the auditor. One will be addressed to the Annual General Meeting of the company, while the other is addressed to the Board of Directors. In both reports the auditor will document an opinion of the existence or otherwise of the internal control system, with scope for the expression of qualifications or reserves.

The existence of the ICS will typically be able to be confirmed if:

- The deliberations of the board of directors in respect of the definition of the principles of the ICS are documented
- This definition corresponds at least to the minimum requirements for an ICS as defined by the legislation in force and taking into account the size, complexity, and risk profile of the company
- The implementation of the ICS is documented and verifications have been able to confirm that the key controls defined by the Board have been introduced

If such documentation and support cannot be obtained and inspected, the auditor may be obliged to conclude that the existence of the ICS cannot be confirmed.

This also applies to partial satisfaction of the requirements: if, for example, an ICS has been designed and implemented but is not the result of a tailored risk identification and assessment process, and it is clear that significant risks exist that are not addressed by the ICS in its current state, the conclusion must be that the ICS is insufficient.

Within a group of companies, if the existence of an ICS for another group entity has been attested to by another auditor or auditors, and if the auditor of the parent company wishes to rely upon this confirmation,

standard audit practices require that the work of the other auditor should be evaluated and the conclusions examined.

Such an approach is not specifically required under NAS 890, particularly in the situation where the subsidiary is able to provide a SOX 404 or an ISA 3152 report, but good practice in the field of external audit recommends that all third-party attestations are reviewed and considered on their merits.

Very often there will be questions to be asked about the scope of the certification, the period it relates to, and the quantity of testing performed and evidence received.

8.7.2.3 *The impact upon cloud computing*

The requirement to demonstrate the existence of an ICS is particularly significant for those organisations that use external service providers, especially cloud services.

Although obtaining sufficient comfort over the management and use of a private cloud will often be reasonably straightforward, obtaining similar comfort over the use of a public cloud will confront the organisation with the same difficulties traditionally encountered when addressing control issues over which they have no direct influence.

In the case of the use of a private cloud, the company can simply follow the directives set out in NAS 890 in respect of companies and groups of companies.

In the case of the use of a public cloud or an externalised private cloud, the company subject to the requirements of the “contrôle ordinaire” will need to demonstrate the existence of its own ICS, of that of its cloud service provider, and of those operated by any subcontractors or other third-parties who might participate in data storage or processing activities.

It is critical in these situations too to ensure that there are no gaps in responsibility, coverage or scope.

The audit of third parties is very often a problematic area, both from a planning perspective and in terms of actually performing the necessary work and drawing the appropriate conclusions.

A number of questions commonly need to be answered:

- Will the service provider accept to be audited and is this requirement documented in the formal contract with the service demander?
- Who will bear the costs of such an audit?
- Do similar arrangements exist for auditing subcontractors or additional third parties?
- Who will bear the costs of these additional audits: the company receiving the service, the main service provider, or the subcontractor?
- To what degree should the service provider and particularly subcontractors be audited? Their entire ICS or simply the aspects most relevant to the services being provided?
- Can deadlines be respected while still providing the range and level of audit comfort required?

Of course, reliance can be placed on previous or existing third-party attestations as to the existence and operation of internal controls, subject to the confirmations and requirements for re-evaluation mentioned above.

8.7.3 Summary

Switzerland is very cautious and conservative in respect of security, privacy and data management and the regulations in place reflect this. The protection of personal data and of the fundamental rights of individuals and entities is at the core of data protection and privacy legislation.

From an audit perspective, the principles of reliable and consistent internal controls are well established but the requirements in place under the LCI are lighter and less constraining than those imposed by such regulations as Sarbanes-Oxley.

8.8 Conclusion

There is an ever-increasing need for audit within organisations, both to provide management with comfort over the quality of their operations and to provide regulators and other third parties with evidence and attestations on respect of quality and compliance. The regulatory environment has grown increasingly demanding over recent years and organisations have had to respond to this by adapting their internal procedures and determining how their compliance can be demonstrated.

New and more demanding regulations have not been published in a vacuum, however. Detailed guidance accompanied the introduction of Sarbanes-Oxley, for example, while professional bodies and standards organisations have also published regularly updated frameworks, guidance and work programmes designed to assist management, their advisers and their auditors to design and implement appropriate controls and verify their operation.

Chapter 9 Identification of Key Technological Elements

9.1 Introduction

“Better be despised for too anxious apprehensions, than ruined by too confident security.”¹²

In previous chapters I have discussed the evident need for effective security measures throughout organisations that are handling sensitive or commercially valuable data. Security is fundamentally a management question, one of organisation, structures, methods and training, in which purely technological aspects are the details that fill out the picture.

Although specific security measures cannot provide all the answers to security questions and cannot individually address each information security objective, it is nonetheless essential to consider the techniques and technologies available when designing security strategies and frameworks, in order to understand which concerns can be addressed by which technologies and which concerns will potentially require additional or more focused manual measures, for example.

In Chapter 6 I discussed concepts such as single sign-on and management practices for restricting and controlling user access to systems and data. In this chapter I will discuss the most frequently-mentioned technology for data security, outlining its principles and use and identifying its limitations.

9.2 The use of cryptography

In any context in which data security is critical, consideration will be given to the use of encryption technologies to prevent data visibility and leakage.

A detailed ENISA publication [106] describes protection measures applied to safeguard sensitive and/or personal data, which has been acquired legitimately by a data controller. As such, “it discusses how information technology users, who have a basic knowledge of information security, can employ cryptographic techniques to protect personal data. Finally, it addresses the need for a minimum level of requirements for cryptography across European Union (EU) Member States (MSs) in their effort to protect personal and/or sensitive data.”

The key recommendations and findings are that:

- The cryptographic measures are only one piece of a puzzle when we refer to privacy and data protection. Cryptographic measures provide a layer of protection and may reduce the impact of breaches.

¹² Edmund Burke, “Reflections on the French Revolution”, § 14.

- The relevant stakeholders (Data Protection Authorities, member states authorities, service providers) should recommend and implement security measures for protecting personal data and should rely on state-of-the-art solutions and configurations for this purpose.
- Specialized personnel are needed for the correct implementation of cryptographic protective measures. Such measures need also to be updated.

9.2.1 Encryption as an element of security

ENISA view information security as a puzzle, the solution for which requires measures and techniques drawn from different domains.

An effective security policy can only be designed after the completion of a detailed risk assessment, in order to ensure that genuine risks related to clearly identified assets and vulnerabilities are being addressed.

9.2.1.1 *Security frameworks*

Logical security by itself forms only one part of the security framework. Encryption is a powerful tool but only addresses one small aspect of data security. Logical access controls need to be supplemented by measures addressing the following aspects:

- Physical security: preventing unauthorised physical access to equipment, machines, storage media and infrastructure elements;
- Organisational security: preventing inappropriate access to systems and data by separating responsibilities and ensuring that roles and responsibilities are clearly defined;
- Personnel security: mitigating the risk related to hiring and employing staff by vetting candidates, implementing education and awareness-raising schemes, ensuring that staff transfer and departure policies and procedures are in place and effective.

9.2.1.2 *Elements of logical security*

Logical security itself is composed of many layers of activities and measures. These include:

- Network security: data flow permanently across and between networks and are vulnerable to interception or modification in transit. This is a particular risk because the fundamental components of the Internet, the core protocols, were not designed with security in mind;
- Operating system security: every device and platform relies upon an operating system to act as an interface between users and application programmes and the hardware and to manage the use and performance of the hardware;
- Database security: many systems depend on the presence of data stored in specialised databases. Management need to ensure that direct access to the data in the databases, using specialist tools or operating system functions rather than passing through application programmes, is limited to the minimum necessary and is appropriately monitored, logged and reviewed;

- Application security: the most common way to access, read or modify data is through the application system to which it relates. It is essential that user access rights are appropriately managed to ensure that no inappropriate modifications or deletions of data are performed;
- Logging and monitoring: effective control environments contain both preventive and detective controls. In practice it is realistically impossible to prevent every accidental or deliberate negative event from taking place, and so it is necessary to implement detective measures to identify possible errors or anomalies after the event and to allow corrective actions to be performed as required. The most fundamental aspects of this are logging and monitoring: logging to ensure that all system events are recorded and available for review; and monitoring, reviewing the logs and system information to identify potential issues. Logs are critical for system management but they do not constitute a control measure by themselves; and
- Cryptographic techniques: rendering information unreadable to those without the means to decipher it is a fundamental means of ensuring privacy and confidentiality.

9.2.1.3 *The role of cryptography*

ENISA describe cryptography as “an essential technical means to provide privacy and privacy-related services.” They cite Gürses, who identifies and defines three privacy paradigms, namely confidentiality, control and practice [107], and show how cryptography has a role to play in these.

“In the first paradigm, privacy is ensured by keeping personal data confidential, i.e. protecting it so that unauthorised people can’t access or modify it. This definition of privacy is clearly and strongly linked to classical cryptographic schemes: electronic data secure can be kept secure by using strong cryptography (and strong keys).

The second paradigm, which is more common in legal texts, adds the ability to control what happens with personal data. This is also called the right to informational self-determination. In this definition, several advanced cryptographic schemes can play an important role, e.g. techniques to reduce the amount of personal data that is released to the strictly required minimum.

The third paradigm defines privacy as transparency on the ways in which information is collected, aggregated and used. Cryptography plays here a much less visible role, but it is still present underneath as the method to enforce the policies formulated with respect to the treatment of personal data.”

9.2.1.4 *The applicability and value of cryptography*

In common with all security measures, cryptography is not a technology to be applied blindly to all aspects of an organisation’s activities. It should be applied where appropriate and in response to clearly-defined and appropriate security requirements.

Security requirements may differ at different stages of the data lifecycle. Five stages are identified, as set out below.

Stage of the data lifecycle	Information security requirements	Comments
Collection and storage	Confidentiality Integrity	It is important to note that although individual data elements may not be sensitive or controversial, when accumulated they may become so. Data should be protected from unauthorised reading access and unauthorised modification.
Transmission	Confidentiality Integrity	Data are often vulnerable during transmission, both over networks where interception is possible, and by transfer media over which standard access controls are often weaker than over storage media.
Retrieval	Availability Integrity	Data must be available when requested and provided with no loss of integrity.
Processing	Confidentiality Integrity	The risk of unauthorised access is higher than for data at rest because the processing environment is typically less well secured than the storage one.
Deletion	N/A	If secure deletion is required, care must be taken to ensure that the data cannot be reconstructed from any traces left on storage media, and that the deletion process applies to every storage and transfer medium that has been used.

Table 9.1 Security requirements through the data lifecycle

9.2.2 Basic cryptographic techniques

There are four basic techniques that can be applied to protecting data.

9.2.2.1 Encryption

This is the process that renders data unintelligible by transforming it into ciphertext, a converted version of the original text that is unreadable without being transformed back. Efficient encryption processes ensure that the reverse conversion is possible only when in possession of the required key. Encryption is useful for both data at rest and data being transmitted, as any unauthorised interceptor of the ciphertext will also need the decryption key to be able to read and use the information contained in the source file.

9.2.2.2 Data authentication

This process makes it possible to guarantee the integrity of data, or at least to identify cases where data have been modified illicitly. This is done by adding a tag (a cryptographic checksum) or a digital signature to the encrypted data. The process is simple: a tag or checksum is calculated by applying an algorithm to the original file or message in the form of a cryptographic authentication key. The file and the associated tag or signature can then be transmitted together. The recipient can verify the validity of the tag using a

verification key that is paired with the authentication key, and any anomaly can be investigated before relying on the integrity of the transmitted file.

Both of the above techniques can be carried out in full knowledge that third-parties may gain access to the ciphertext, source data or tag. In many cases the techniques will be combined for more effective and wide-ranging security, in the form of authenticated encryption. In this process the original message will be encrypted and a tag calculated based on the ciphertext. It is this ciphertext that is transmitted, along with the tag. In this way, as long as the cryptographic keys are properly protected, it is possible to ensure both the confidentiality and the integrity of the source message.

9.2.2.3 *Hashing*

This is a process that calculates a short and unique representation of data in an irreversible way. It is a process that has long been used in internal control procedures in a variety of contexts, such as in verifying the completeness of batch input into information systems by performing a calculation upon a numeric field and comparing the initial calculation to an identical one generated by the system at the completion of the input process. Hashing does not use a cryptographic key as input and is not an authentication mechanism by itself. Its use comes from providing a value, a fingerprint of a message or set of data, that can be used as the basis of a digital tag. Speed and efficiency are key to the effective use of cryptographic technologies in daily life, and it is far quicker to calculate the digital signature of a hash value – and to recalculate it on reception – than to perform the same process on a larger message or set of data.

9.2.2.4 *Digital signing*

This is a process that creates a short string that depends on the content of the message and a secret – a key – known only to the signer. It is important to note that the validity of a signature can be verified without knowledge of the secret itself. Digital signing provides both data authentication and identification.

9.2.3 Symmetric and asymmetric encryption

9.2.3.1 *Symmetric encryption*

Symmetric encryption is based on all parties using the same keys. Both encryption and decryption operations use the same keys, or keys that can easily be derived from each other. Both keys need to be kept secret for the system to be effective.

9.2.3.2 *Asymmetric encryption*

Asymmetric encryption is as its name suggests based on pairs of keys that cannot be derived from each other. The keys used for decryption are separate from those used for encryption, and the keys used for data authentication are separate from those used for validation. Because of the asymmetric nature of the system, it is not necessary – or even desirable – to keep all of the keys secret. Only the keys needed for decryption and for data authentication need to remain secret, and these are known as private keys. Their counterparts used for encryption and validation can be published and are known as public keys. The reason behind this usage and need for availability of keys is practical. If someone wants to send someone a message that only the intended recipient will be able to read, the message can be encrypted with the recipient's public key. Only the intended recipient has access to the corresponding private key and will be able to decrypt and read the message. In the same way, a person wishing to send a message of guaranteed

integrity can create the tag with their own private data authentication key. When the recipient receives the message, they can verify the authenticity of the tag by using the sender's public key, and be confident that the sender was indeed the originator of the message.

9.2.3.3 Key management

Effective key management is the basis of success cryptographic mechanisms, and indeed "cryptographic mechanisms reduce the problem of data security to the problem of key management. This is known as Kerckhoffs' principle: 'the security of a cryptosystem should not rely on the secrecy of any of its workings, except for the value of the secret key'."

ENISA report that "By necessity, symmetric cryptography is based on secret keys. In order to keep the application secure, there needs to be a mechanism to protect the confidentiality and the authenticity of the keys."

By way of contrast, asymmetric cryptography brings with it asymmetric security requirements for the keys. In brief:

"In order to control the confidentiality of data, we need to control only the access to the decryption key. In principle, the encryption key may be known by other entities. In order to guarantee the authenticity of data, we need to control only the access to the authentication key. In principle, the verification key may be known by other entities."

The objectives of key management correspond very closely to the fundamental objectives of data security management:

1. Protecting the confidentiality and authenticity of secret and private keys, as well as protecting secret and private keys against unauthorized use.
2. Protecting the authenticity of public keys.
3. Ensuring the availability of secret and public keys.

The fundamentals of key management also correspond in many aspects to the fundamentals of lifecycle management of other forms of data. ENISA identify seven major concerns.

- Key generation: keys need to be unpredictable and, in the case of asymmetric primitives, need to demonstrate a sufficient amount of entropy in order to ensure confidentiality. In practice generating suitable private keys can be a challenging task.
- Key registration/certification: keys need to be indelibly associated with their owner. For public keys, this is done through public-key certificates issued by registration or certification authorities, guaranteeing that a certain key belongs to a certain user, and can be used for specific purposes during its period of validity. Certificates are public documents authenticated in their turn by digital signatures.
- Key distribution and installation: as has always been the case for physical code-books and encryption machines, digital keys need to be distributed to their users. This generates problems of confidentiality and authenticity of the kind that cryptographic keys themselves are designed to

resolve. “For systems based on symmetric cryptography, both the sender and the receiver need to obtain a copy of the key, hence the key needs to be transported securely (protection of confidentiality and authenticity) at least once. All the copies of the key need to be installed and stored securely. For systems based on asymmetric cryptography, the private key is often generated in place, such that no transport is needed. (In secure hardware, the functionality to export the private key of an asymmetric key pair is usually deliberately not implemented.) Otherwise, it needs to be protected like a symmetric key. The public key still needs to be transported, but only the authenticity (and hence the integrity) needs to be protected, which is achieved by the use of certificates.” Key distribution centres do also exist to simplify the process.

- Key use: throughout the lifetime, keys should be protected against unauthorised use both by attackers and by the owner of the key. This includes exporting the key or using it in an insecure environment.
- Key storage: secure hardware can ensure that keys are never exported, but as is true for all critical data there is also a need to maintain backup copies. These might prove necessary to replace a lost key, to be able to access data after an employee leaves, to access old data requiring expired but archived data, or to provide suitably motivated and justified access to data to law enforcement agencies.
- Key revocation/validation: keys expire and need to be replaced, or sometimes need to be withdrawn before their planned end of life because of leakage or advances in cryptanalysis. Procedures need to be in place to ensure that keys can be properly revoked and that employees no longer use or rely on keys that have been withdrawn early. Validation in this context means confirming that an operation was performed with a key that was valid at the time of the operation.
- Key destruction: keys need to be removed from hardware after they have expired. In common with other forms of sensitive data, management needs to ensure that such deletion is both timely and permanent, so that the key cannot be recovered from the hardware using forensic techniques.

9.2.4 Types of attacks on standing data, and possible counter-measures

ENISA identify five possible types of attack on standing data and discuss the nature of the counter-measures that could be used to ensure protection.

9.2.4.1 *Hashed-data-only leak*

The first is a *hashed-data-only leak*, in which only the hashed values of data are exposed, not the data themselves. In this scenario, the immediate impact is low because hashing is by definition a non-invertible function and thus attackers cannot derive the original data directly from the hashed values. Attackers can, however, compare the hashed values to the hashed values of data that they are hashing themselves to confirm the accuracy of guesses or processing. This is particularly relevant to a situation in which a file containing hashed passwords (for example, the `/etc/passwd` file in many Unix systems) is leaked.

In this scenario, if the hash function used is a secure cryptographic keyed hash function and that key has remained secret, leakage of the hashed values does not cause any immediate leakage of data. The leakage will only allow attackers to confirm some guessed or brute-forced value for individual data items.

9.2.4.2 *Logical leak*

The second case is a *logical leak*. In this situation, encrypted data have been exposed, through an everyday occurrence such as a portable media device being lost or stolen or transmissions across a network being intercepted. This kind of situation is analogous to the loss of data in printed form, in that it is possible to determine exactly which sets of data were exposed. Because there are no complicating factors to consider, with only the communications channel being attacked, the data will remain secure as long as current authenticated encryption has been used, the source data have not been exposed, and the decryption key remains secret.

9.2.4.3 *Hardware leak*

The third case is a *hardware leak*. In this scenario, attackers have gained access to or control over data-storing hardware such as powered-down laptops, external hard disks or USB sticks. In such a situation, the data on the hardware are secure as long as they are protected by authenticated encryption and if all unencrypted copies of the data have been deleted securely. When considering a stolen or lost laptop, it is also essential that swap and hibernation files have always been encrypted, because these can contain sensitive data. It is also necessary to ensure that the cryptographic keys in use are also properly protected, in particular ensuring that the top-level cryptographic key is never stored on the device itself but is required to be input manually or read from a separate token.

9.2.4.4 *Break in on a live device*

The fourth case is a *break in on a live device*. In this scenario, attackers have remote, logical access to a device that is using sensitive information, perhaps by hijacking a computer or through a virus that is transmitting information. This is difficult to defend against unless data remains subject to authenticated encryption throughout the attack and cryptographic keys have likewise been encrypted. Because of the way computers operate, with data files opened in RAM during sessions and thus residing in RAM in unencrypted form, both data and keys present on the machine during the attack should be assumed to have been leaked during such an attack. Protection against such attacks consists of only unencrypting a minimum of data at a time, as necessary, and in avoiding illicit access to cryptographic keys by keeping as many of them as possible inside secure hardware.

9.2.4.5 *Full user impersonation*

The fifth case is *full user impersonation*, during which a user loses passwords or access tokens and attackers can exploit this knowledge to impersonate the user. Such an attack will logically give the attackers access to the user's working environment as far as the stolen credentials allow this, and so the response has to be based upon reducing to a minimum the amount of data that can be accessed. Typically this will be done by keeping as many keys as possible within secure hardware and by using an appropriate key architecture so that keys that are required to reside in RAM during usage each only encrypt small amounts of data. In addition, access controls should be designed so that encryption keys are not present on the access tokens.

9.2.5 Countermeasures

A number of available countermeasures can be seen to be applicable across these attack scenarios. In brief, these consist of:

- Secure deletion of files: ensuring that data remanence is avoided by systematically overwriting the sectors of a disk or other storage medium used for storing data until recovery becomes difficult. If a physical solution is required, disks can be degaussed to return the media to pristine state, or, if the media are no longer required, they can be physically destroyed.
- Authenticated encryption: while encryption ensures that data are unintelligible except to users who possess the relevant decryption key, data authentication ensures that users who possess the validation key can detect if unauthorised modifications have been applied to the data during transit or storage. The two techniques are almost always employed together in modern systems because there are few contexts in which it would make sense to have data confidential but have no assurance over its integrity, and because encryption systems without authentication are more vulnerable to attacks than those with it.
- Secure hardware: this refers to dedicated processors – in the forms of smartcards, hardware security modules or trusted platform modules – that are used for handling sensitive or confidential data. Such hardware can be programmed to provide only a restricted set of services and by limiting the inputs that it accepts and the outputs that it delivers, its security can be increased in relation to more general purpose hardware. An example of secure hardware could be a system on which all cryptographic keys are stored and used, ensuring that they never pass onto less secure elements of the system.
- Key architecture: good practice insists that the amount of data encrypted with any one key should be limited. This limits the exposure should a key be revealed or limits the extent of security degradation suffered should significant amounts of data be encrypted or authenticated with the same key. A key architecture will consist of two or more levels of keys used at the lowest level to encrypt-authenticate or decrypt-validate previously defined subsets of data, and at higher levels to encrypt-authenticate or to derive the lower level keys. Typically keys at higher levels will be kept on a restricted set of devices and kept apart from keys used for direct source data processing purposes. The architecture will also provide for the controlled updating, expiring and deleting of keys.
- Access control system: passwords will be used in combination with tokens containing cryptographic keys.

9.2.6 Data minimisation

An important trend in data protection is that of data minimisation. The underlying principle is that the best way to secure data is never to have collected them at all: the practice is therefore to avoid storing sensitive data, therefore alleviating the need for stringent and widespread security. This is done by performing as much local processing as possible and transmitting and storing only the results of the processing; by using attribute-based credentials that provide the minimum amount of information necessary to confirm an identity or a characteristic; or by using “private information retrieval”, in which users can retrieve elements from a database without the database owner knowing precisely which element has been retrieved.

9.2.7 Data stripping and separation

A novel approach to improving the security of data held on the cloud was presented by Leistikow and Tavangarian [108] at the AINA 2013 conference in Barcelona. They presented a novel secure partitioning method based on data separation by using facial recognition and stripping algorithms, essentially making the data unusable to any unauthorized user gaining access to it. Although implemented in an initial phase in respect of image files, in an attempt to improve security for individuals using web services such as Dropbox, Picasa, or Flickr for managing, saving and organising digital pictures, the principles could be applicable to other types of files and provide an additional level of security, one that is not dependent on the data management processes of the Cloud provider.

9.3 Directives

A number of directives are in place in respect of data protection and these refer to and impose the use of cryptography in order to ensure that data are stored in a form that avoids the identification of the subjects of the data. Two of the most important of these directives from a European perspective and the EU Data Protection Directive and the e-Privacy Directive, which are discussed below.

Such directives are important to all organisations that are handling, processing and storing personal data because they set out the precise requirements with which the organisations, and any partners with which they work or service providers to whom they outsource all or part of their processing or storage, must comply. From a legal, regulatory and compliance perspective, a proper understanding of these directives is essential.

9.3.1 EU Data Protection Directive

One of the key drivers behind the protection of personal data is the Data Protection Directive 95/46/EC [109] on the “protection of individuals with regard to the processing of personal data and on the free movement of such data.” Article 17 of this directive addresses aspects of the security of processing, mentioning that “1. Member States shall provide that the controller must implement appropriate technical and organisational measures to protect personal data against accidental or unlawful destruction or accidental loss, alteration, unauthorized disclosure or access, in particular where the processing involves the transmission of data over a network, and against all other unlawful forms of processing. [...] [S]uch measures shall ensure a level of security appropriate to the risks represented by the processing and the nature of the data to be protected. [...] 2. The Member States shall provide that the controller must, where processing is carried out on his behalf, choose a processor providing sufficient guarantees in respect of the technical security measures and organisational measures governing the processing to be carried out, and must ensure compliance with those measures.” Such measures in respect of data protection could include encryption.

Article 6 concerns data quality and specifies that personal data must be kept “in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the data were collected or for which they are further processed.” Safeguards to ensure compliance with this article could involve data minimisation and anonymisation techniques, both of which can be implemented using encryption technologies.

It is thus clear that cryptography has a key role to play in ensuring data security and in ensuring compliance with regulatory requirements.

In January 2012 the EC proposed a new regulation [110] to replace the existing Data Protection Directive. This regulation builds on the existing requirements, for example obliging, in Section 2, Article 30, data controllers and processors to implement appropriate measures for the security of processing; this is an extension of previous requirements and sees the obligation extended to processors, irrespective of their contract with the controller. Requirements will subsequently be defined for various situations, including to: “(a) prevent any unauthorised access to personal data; (b) prevent any unauthorised disclosure, reading, copying, modification, erasure or removal of personal data; (c) ensure the verification of the lawfulness of processing operations.”

9.3.2 e-Privacy Directive

The data breach notification provision in Article 4(3) of the e-Privacy Directive 2002/58/EC [111] is built upon with the introduction of an obligation to notify personal data breaches. Significantly, some exceptions are provided for, for example if “appropriate technical protective measures” are in place in order to “render the data unintelligible to any person who is not authorized to access it.”

Article 4 also specifies security requirements in respect of privacy and electronic communications. “1. The provider of a publicly available electronic communications service must take appropriate technical and organisational measures to safeguard security of its services, if necessary in conjunction with the provider of the public communications network with respect to network security. Having regard to the state of the art and the cost of their implementation, these measures shall ensure a level of security appropriate to the risk presented.”

Directive 2002/58/EC has been amended by directive 2009/136/EC [112]. Among the amendments to the above-mentioned Article 4, it is specified that data owners and operators should: “ensure that personal data can be accessed only by authorised personnel for legally authorised purposes”, “protect personal data stored or transmitted against accidental or unlawful destruction, accidental loss or alteration, and unauthorised or unlawful storage, processing, access or disclosure”, and “ensure the implementation of a security policy with respect to the processing of personal data.”

Again, exceptions are provided for. “Notification of a personal data breach to a subscriber or individual concerned shall not be required if the provider has demonstrated to the satisfaction of the competent authority that it has implemented appropriate technological protection measures, and that those measures were applied to the data concerned by the security breach. Such technological protection measures shall render the data unintelligible to any person who is not authorised to access it.”

It is important in this context to understand what is formally meant by “unintelligible.” This is defined in an EU regulation published in 2013 [113].

“Data shall be considered unintelligible if:

1. It has been securely encrypted with a standardised algorithm, the key used to decrypt the data has not been compromised in any security breach, and the key used to decrypt the data has been generated so that it cannot be ascertained by available technological means by any person who is not authorised to access the key; or

2. It has been replaced by its hashed value calculated with a standardised cryptographic keyed hash function, the key used to hash the data has not been compromised in any security breach, and the key used to hash the data has been generated in a way that it cannot be ascertained by available technological means by any person who is not authorised to access the key.”

9.3.3 Summary

When applied correctly and properly administered, cryptography represents a powerful tool for data owners and administrators to protect the privacy and confidentiality of their data and to ensure compliance with standards and directives that apply.

It should be remembered, however, that cryptography is only ever a partial solution to security problems. If we return to the fundamental principles of information security and consider the three pillars, we can see the cryptography can help ensure *Confidentiality*, if properly implemented, and can also ensure some elements of *Integrity* when considering the completeness and accuracy of data at rest or in transmission. It cannot, however, provide any assurance over the *Availability* of data, and indeed if not properly controlled can present problems in that respect.

9.4 Conclusion

This chapter and the three preceding it have demonstrated how much theoretical work exists in respect of internal controls, risk management, security management and management comfort. Frameworks and technical methods exist for identifying risks and designing solutions, as well as for monitoring the ongoing security of data. The impact of compliance requirements is also clearly documented, as well as the means of ensuring that such requirements are met and that management can obtain the comfort they require over the adequate security of their systems and data.

It is essential to view data security both as a process and as a complex combination of elements. Information security is never achieved definitively [114], as technologies, risks, vulnerabilities as weaknesses change, and it is the task of management to identify the correct combination of measures that address current risks and vulnerabilities, as well as implement procedures for the regular and systematic review of the operating environment in order to identify changes and, if necessary, make appropriate changes to the risk assessment and control environment.

Security must be approached in practice from a consideration of the objectives, overall and detailed, that management have in respect of their systems and data. These objectives may be driven by external compliance requirements, by commitments made to partners or customers, by principles of good internal control, or by a combination of all of these. Only when objectives are clearly defined can realistic, coherent and systematic measures be designed and implemented; addressing a subset of risks in an inconsistent, non-prioritised and haphazard manner will lead to inconsistencies and weaknesses in the measures chosen, as well as potentially provoking an unrealistic and misleading sense of confidence in the overall quality of the measures chosen.

Technologies and techniques such as cryptography, user access controls and physical access controls are individual elements that properly combined can provide the level of comfort that management require. It is important to remember the limitations of each technology or method, however. No single approach will provide the entire range of comfort that is sought; rather, each is a link in a chain. It is also essential to

remember the management and operational requirements of each approach: cryptographic solutions need to be monitored, managed and updated on a constant basis, in the same way as user access rights need to be subject both to appropriate control procedures over the granting, modification and deletion of rights and to robust and effective ongoing monitoring of existing rights.

Not every technique will be applicable in every situation. Precise requirements will change, depending on the compliance environment, the sensitivity of the systems and data, the budgets available, and the perceived risks. Not every security framework will involve encryption, although nowadays it is extremely widespread. Instead, an approach may be adopted that involves very restricted access to data at rest and in transit and also involves an appropriate level of regular monitoring. Management should therefore avoid making assumptions about the measures to implement until the initial objective-setting and risk assessment stages of the security process have been completed.

A key element in the ongoing monitoring and continuous improvement processes consists of the identification, classification and treatment of incidents. A number of frameworks exist for this: management need to choose the framework that is the most appropriate for the organisation, taking into account risks, benefits and resources. It is essential that the approach chosen will be sustainable and practical for the organisation, because management comfort is based on both the design effectiveness and operating effectiveness of control procedures. Implementing a control activity or framework and then not ensuring that it is properly monitored and reviewed offer at best no more value than having no controls in place, while such situations can be shown to be more harmful than having no relevant internal controls because of the impacts on decision-taking and behaviour of a misplaced and inappropriate confidence in the quality of the control environment.

Chapter 10 A high-level proposal to address the issues

10.1 Introduction

“Simplify, simplify”¹³

It is clear that there are a number of issues to be addressed when designing a framework solution for addressing the security concerns related to the management of unstructured data within the cloud environment. These issues include: the familiarity of members of staff and users at all levels of the hierarchy with the concepts of structured and unstructured data and the various options for managing and storing data; understanding of the risks – legal, operational and technical – related to data management; knowledge of what data are held within the organisation and how vital and valuable these might be; and understanding of the responsibilities and duties of every person who uses or handles the data.

In the previous chapters I have identified the current state of analysis and attitudes in respect of the use of these technologies. In Chapters 3 and 4 the published literature has been reviewed; in Chapter 5 the results of market surveys were presented and analysed; and in Chapters 6 to 9 the key constraints for management and users were identified and discussed.

In proposing a pragmatic solution for the issues identified, it is necessary to drill down to the core elements of the problems and address those elements fully and sequentially. A useful and relevant analogy is that of improving the general level of information security within an organisation: the only systematic and practical method of doing this is to address the key issues from a top-down perspective and implement the correct structures for coherent and methodical security practices. Attempting to effect improvements by modifying details in isolation – for example, by changing the minimum length of passwords – will have no useful impact.

The approach I have adopted, then, is to simplify the issues as much as possible and to follow a top-down approach based on clearly-identified risks and consistent and documented responses.

The framework solution is intended for the use of management, to assist them to address their own risks and ensure that appropriate responses are designed and implemented. The language and structure of the framework is consistent with the approach adopted by internal controls specialists and auditors and the outputs, if documented and validated, correspond to the requirements of auditors in respect of demonstrating compliance with good practice and the design and operating effectiveness of internal controls, but this is a practical consequence of the design of the framework rather than its key objective. The key objective is to allow management to oversee and perform a successful project to their own

¹³ H. D. Thoreau, “Walden, or, Life in the Woods”, chapter 2

satisfaction. As such, the solution seeks to address both governance and compliance requirements in respect of data handling. As discussed in Chapter 8, this solution is primarily aimed at management and is designed to address their requirements for appropriate internal controls, allowing them to carry out projects and implement ongoing control procedures in accordance with their specific requirements for demonstrable governance; the outputs from the processes, if properly documented, approved and retained, will also help them demonstrate compliance with applicable requirements to their satisfaction and that of their auditors, if necessary.

The choice of vocabulary is a result of some discussion. During development, the solution was referred to as a template, a model and a framework, almost interchangeably. After reflection, the term “framework” was chosen as the most appropriate term to be used consistently because it best describes both the nature and the intentions of the solution. It is a framework that encompasses the whole range of project activities and within which other, more specific, frameworks and models can be incorporated should that be desired. An example of how this might be possible is given in Section 10.16 below.

10.2 Understanding the concepts and related vocabulary

At the most senior levels of the hierarchy, the board members and members of management committees need to be equipped to take the most significant strategic decisions. They do not necessarily need to be able to decide on choices of technology, as such matters tend to be tactical and operational decisions that can effectively and appropriately be delegated to managers with responsibility for technical and control issues, with the ultimate decisions obviously requiring high-level approval and sponsoring. The language for the education of these senior levels, and for their discussions and decision-making processes, will by necessity be the language of corporate control, of risks and exposures, potential losses and returns on investment.

The middle-management layers, on both the technical and business sides, need to understand both the big picture in respect of overall risks and strategic directions and concerns and the technical and operational details of both the current operating environment and any new solutions that are to be chosen.

Users need to have a clear understanding of the scope and purpose of the data description and inventory project, in order to ensure their buy-in and enthusiasm, and also be completely familiar with the terminology being used and with the systems and data repositories that fall into the scope of the discussions.

10.3 Scoping: types of data, uses of data, and storage options

Management and users need to have a very clear understanding of the type and nature of the data they possess.

10.3.1 Structured data

Typically a large organisation will possess one or more substantial data stores, usually containing the core financial information and also encompassing inventory, purchasing, manufacturing, customer and supplier management, and human resource information. In cases where the organisation is using a centralized ERP system it is possible that all this information will be sitting in one large central database; in other circumstances the information will be stored in the dedicated databases used for each application or

system. Of course, it can also be the case that a central ERP system can be supported by additional external applications, for a variety of reasons: the equivalent module for the ERP may not exist or be sufficiently mature to meet customers' requirements; the implementation of the ERP, which can often be spread over several years, may not yet be sufficiently robust or mature to allow the implementation of additional modules; or the organisation may prefer to retain legacy systems for reasons of efficiency or practicality.

Regardless of the decisions that have led to the current situation, the organisation needs to identify all of the formal and structured data sources that are in use and be aware of the nature of the information within them. This means understanding what areas of activity it relates to, what format it is in, which periods it covers, and how complete it is.

10.3.2 Unstructured data

Typically it will require more effort on the part of management and users to identify the sources and storage locations of unstructured data. Whereas structured data linked to specific applications and platforms will reside in documented and identifiable databases, unstructured data, as discussed in Chapter 3, can be found throughout an organisation, in a variety of formats and stored on a variety of platforms and media. Management must therefore draw up and implement detailed plans for the systematic review of storage media in order to identify files and databases that might meet to criteria to be considered useful sources of data. In particular, care should be taken to review departmental networks and local storage devices to ensure that nothing is missed. Frequently this process will be automated, using dedicated machines and specialized software such as spiders to crawl through the organisation's networks and catalogue files.

10.3.3 Storage media

It should not be forgotten that there are numerous possible locations for data storage, and all should be considered when searching and cataloging the data held within an organisation.

10.3.4 Central databases

Located on central servers and networked to allow shared access, these are the main sources of structured data.

10.3.5 File servers

These are used to permit the central storage and sharing of files of all types.

10.3.6 Mail servers

Mail servers will contain copies of emails and their attachments, which are a frequent source of unstructured information.

10.3.7 Local hard drives

Many users have the habit or requirement to store documents on the local drives of their workstations or laptop computers.

10.3.8 Backup media

In overall security management it is essential to consider the contents and location of backup media; this also applies when performing a data identification and cataloging exercise as there may be previous versions of existing data on those media that could be valuable to catalogue and recover, as well as files that had subsequently been deleted from production systems.

10.3.9 Smart devices

Over recent years smart devices such as smartphones and tablet computers have become essential tools for individuals and for employees and useful information can be stored on these in the form of notes, documents, and mail attachments.

10.3.10 Portable storage

The domain of portable storage was revolutionized at the beginning of the 21st century by the widespread adoption (and improvement) of the USB interface and by the rapid growth in the storage capacities of USB storage devices, both solid state and rotating disk. In a context in which users are encouraged to backup and maintain multiple copies of their private and professional data, portable media need to be included in the cataloging process.

A great deal of use is also made of card media such as SDHC cards, commonly used in digital cameras, smartphones and portable music players to provide mobile storage but equally capable of being used as removal external storage media for PCs.

10.3.11 Private cloud or remote storage services

Many individuals already use cloud services such as Dropbox or Skydrive to ensure permanent access to, and synchronization of, their most frequently used files. Even if corporate security policies expressly forbid the use of such services for storing professional documents, management should be aware that such behavior is possible and that useful files may be stored in such locations.

10.3.12 Uses of the data

In order to be able to assess the usefulness and validity of data effectively, management and users need to know exactly what those data are used for. In certain cases this will be obvious and not require any analysis, as in the case of core accounting data. In other cases, however, this will potentially be more difficult, particularly in situations where data have been extracted from a larger system. Usually this occurs for the purposes of analysis or reporting, but the more remote the extracted data becomes from the source, both in time and after manipulation, the less obvious the uses will be unless they are clearly documented.

10.3.13 Quality of the data

Management and users need to know whether the data catalogued during the search process are worth retaining and preparing for transfer. There are many reasons why this might not be the case: data extracted from databases or applications may be incomplete or have been modified so that they are no longer consistent with the source data; the data may be old and no longer valid; the data may refer to aspects of

business or operations that are no longer of interest to the organisation. In all cases the Integrity criterion is paramount, however, as are the related audit objectives of completeness, accuracy and validity.

10.3.14 Simplifying the scoping

In order to simplify the scoping and discovery process as much as possible, management should approach the process in a structured manner, taking care to define parameter in advance to avoid wasting time. Consideration should be given to determining which kinds of data should be excluded from the search on the grounds of relevance or age, while procedures should be defined in respect of when decisions can be taken, and by whom, to halt particular lines of searching. As mentioned in 10.3.2 above, consideration should also be given to using automated tools, developed in-house or acquired commercially, to automate the research as far as possible. As a matter of good practice, of course, the decisions over which data to exclude and how these will be identified should be documented, along with the justification and approval of these decisions.

10.4 Understanding the risks

10.4.1 Legal risks

As discussed in Chapter 4, there are numerous legal concerns in relation to transferring data outside the direct perimeter of control of an organisation and entrusting it to a third-party. The data owners retain overall responsibility for the security and use of the data while losing certain aspects of direct control over them.

10.4.2 Operational risks

The transfer of data outside the organisation's perimeter can affect operational aspects of system management because of the absence of direct control over the availability of data, the speed of access to data, and the response to support or change requests. When all is held and supported in-house, there are usually more direct lines for resolving problems based on closer relationships, whilst the first-line support itself is often efficient because of an underlying in-house understanding of business practices, processes, and objectives.

10.4.3 Technical risks

A number of technical risks are presented by a move onto third-party hosting. Primarily, there is the dependence on the service provided by that third-party for the availability and integrity of data. If connectivity to the service provider is lost, then unless effective and appropriate contingency plans are in place and can be activated, operations may be interrupted. Access to systems and data can also become more complicated, remote logging being required rather than direct access through consoles.

10.5 Roles and responsibilities

As discussed in Chapter 4 when considering service level agreements and the split of roles and responsibilities, these issues need to be clarified and agreed upon very early in the planning phase in order to ensure that there are no gaps and no disputes. Two general principles apply: the data owners retain overall and final responsibility for the security and confidentiality of data; while service providers will, for

legitimate and risk management-related reasons, only assume the roles, responsibilities and risks that their contracts and service level agreements require them to do.

When the decision has been taken in principle to move data onto the cloud, careful consideration then needs to be given to each of the phases of the transfer project and the significant issues that arise within each.

10.6 Planning the transfer

10.6.1 Identifying the data to be moved

A pragmatic approach includes a phase of detailed consideration of which data are useful to the organisation, which data within that subset should be moved to the service provider, and which data can potentially be deleted. Not all useful data will necessarily be included in the migration project: some may be deemed too sensitive to leave the existing security perimeter; other data may be consulted so infrequently, or have so little link to other data, that there is little perceived return on the cost of preparing it for transfer and having it hosted.

10.6.2 Cleaning the data and ensuring consistency

Any data cataloguing and migrating project will necessarily include a phase, often lengthy and costly in resources, during which the data are cleaned and checked for consistency. This is particularly important in the context of a migration onto the cloud for two reasons: firstly, if cloud services are to be billed on the basis of the quantity of data stored, it is only common sense to ensure that the data transferred does not contain irrelevant and out of date items, nor duplicates; secondly, there will be an automatic assumption on the part of users that after such a migration project has been performed the data being accessed will be, automatically, consistent and coherent and accurate. It is thus an essential step of the migration process to ensure that the data are as clean and accurate as possible, in order to avoid analytical or decision taking issues further down the line.

This phase can present its own challenges in respect of security. A case can be made within many organisations to perform such cleaning and processing on its own dedicated infrastructure, in order to avoid intensive use of processing power on the existing infrastructure that might slow down or interfere with standard daily operational tasks, and also in order to avoid local area networks being overloaded with additional traffic. The argument could then be made to use facilities provided by third-party service providers, servers and networks configured to the client's designs and rapidly available. Doing so, however, would then introduce additional complications in respect of security, both in respect of the transfer of the raw data to the new platform and in respect of access to that data once it was hosted by the service provider. Conceptually this is exactly the same security problem as is presented by the full move to outsourced platforms, but encountered earlier in the process and perhaps before detailed security planning had had a chance to be performed.

10.6.3 Formatting and portability

Very often the data catalogued and identified for migration will have been taken from different systems running on different platforms, and because of the differences in the way different technologies store and handle data, there will almost inevitably be differences in the way files are structured and in the way their

contents are ordered. Thanks to the widespread adoption of open data sources and standard formats, the mutual legibility of files is no longer the barrier to the transferability of files and data as it was in the early days of micro-computing, but the staff responsible for the migration project should remain aware of the potential difficulties to be encountered when accumulating data from multiple platforms, especially when using older and less standards-compliant hardware and operating systems.

10.6.4 Determining what should happen to local copies of data being transferred

Clearly the data transfer process will involve the creation of data files on local systems that will then be transferred onto the service provider's platforms. Management should already have decided what will happen to these local copies once the transfer has been completed and verified. Essentially there are three options: to keep the local copies, which may be the most prudent approach but which will logically incur local storage and maintenance costs which the move to outsourced might have been intended to remove from the operating budget; to copy the local copies onto backup media so they would always be available again locally if needed, and then delete them from production systems; or simply to delete them as soon as possible. This last option may sound drastic, but if local storage space is limited or if existing platforms are being decommissioned, it could be argued to be the best decision.

10.6.5 Determining the completeness, integrity and accuracy of the data transfer process

A fundamental step in any data migration process, be it from one accounting system to another within an organisation or a transfer of data outside the organisation, is to verify the completeness, integrity and accuracy of the transfer process. Typically this will be performed using a variety of analytical techniques, the precise nature of which will depend on the types of files and data being transferred. File counts and file size comparisons will always form the backbone of the approach, but additional comfort can also be gained from hash total calculations performed upon numeric fields pre- and post-transfer, while other fields may lend themselves to reporting on statistics and content.

This will often be a multi-stage process as the same techniques will apply to the consolidation and cleaning of data, extracting data from their original files and formats and generating repacked transfer files and needing to compare the results of these processes with reports run on the source files. When the final transfer of data files onto the service provider's platforms is complete, a complete audit trail should be available to attest to the quality of the process and the upholding of the three key objectives of completeness, integrity and accuracy.

10.7 Handling other data appropriately

10.7.1 Reorganisation

Decisions will need to be taken as to how data that are not to be transferred are to be treated during this process. Such data will typically consist of data that are deliberately being retained in-house for reasons of confidentiality or accessibility, or data that are not deemed to be of sufficient value to be worthwhile classifying, cleaning and transferring within the framework of the current project. For each dataset, it will be necessary to decide whether cleaning and reorganising is appropriate, taking into account the

economies of scale of performing such a task as part of the overall migration preparation project rather than doing so as a stand-alone project at some other time.

10.7.2 Deletion

For data that are considered valueless for the organisation and therefore appropriate for deletion, a procedure needs to be defined and implemented for carrying out this deletion. There are two major aspects that need to be considered when defining this process:

10.7.2.1 *Completeness*

When data are to be deleted, consideration should be given to the desirability and practicality of deleting every copy from every storage medium. This includes servers, workstations, external disks, USB keys, portable devices and backup media. Clearly the deletion of live copies from servers should be reasonably straightforward, but identifying where all other copies reside may have been a time-consuming part of the whole data scoping exercise and actually retrieving those files and deleting them is likely to prove even more time consuming. Depending on the rotation policy in place in respect of backup media, and of course on the backup policy itself in respect of scope and completeness and whether incremental or complete backups are taken, copies of unwanted data will gradually be removed from circulation, although the data will have a tendency to linger on the media that are allocated to long-term storage, such as annual year-end backups intended to be retained for ten years for certain compliance purposes.

10.7.2.2 *Security*

If the data to be deleted are considered to be sensitive, for example containing personal details of former employees or clients, it may prove necessary to ensure that the deletion is permanent and irrevocable. In order to achieve this, it will not be sufficient to simply delete the files from the various storage media upon which they are stored. It will also be necessary to overwrite the media securely so that traces of the files in question cannot be retrieved, or in the case of media that can be replaced to physically destroy them sufficiently for there to be no means of reading their contents again. This can be done mechanically, physically destroying the disk and its platters or shattering the memory chips on a non-volatile device, or through electromagnetic means.

10.8 Addressing legal considerations

10.8.1 Confirming that no requirements or obligations are being contravened

All organisations holding personal or sensitive data must, as discussed in chapter 8, ensure that they will remain in compliance with appropriate requirements and obligations both during the data migration phase and after the data have been transferred and have started to be used. Usually such organisations will have been aware of such obligations while all data were stored in-house and therefore will be familiar with relevant requirements: where particular attention needs to be paid, however, is in ensuring that the facilities provided by the service provider and the ongoing guarantees of privacy and confidentiality are sufficient to meet the requirements of users and of the entities or administrations that define the legal obligations, and then throughout the migration process. In general it is during times of change that control procedures come under strain and where failures and exceptions can be expected. The privacy and security of data are thus under greater pressure during change processes.

10.8.2 Informing the subjects of data

Management needs to be aware of any requirements to inform the subjects of data of any changes in where it is stored or how it is accessed.

10.8.3 Informing regulators and authorities

Similarly, management should confirm at an early stage in their project planning whether it will be necessary or not to inform regulators or authorities of the proposed move to outsourced facilities.

10.9 Contingency planning for problems with the transfer

A fundamental part of good practice for all information systems change projects, be they software changes, hardware changes, modifications to parameters or data migrations, to give just a few examples, to have a structured, complete and verified contingency plan for the data transfer project. Structurally this should not differ in any significant way from traditional disaster recovery and business continuity planning, which is understandable given that depending on the size and nature of the data transfer, the organisation could potentially find itself in a situation where a failure in the transfer could leave it unable to function.

We can imagine the situation in which an organisation has decided to retain its core accounting systems in-house but has decided to outsource the hosting of its Client Relationship Management (CRM) system. If problems are encountered with this transfer, and data are corrupt or incomplete after the process, the ability of the organisation to continue with its core daily operations may be compromised. Central accounting functions would be unaffected, but any process that relied upon CRM data could be hindered or, in an extreme case, rendered impossible.

It is therefore unarguable that a well-managed organisation will have contingency plans in place to ensure that any problems encountered during the transfer process can be identified and addressed in a controlled manner.

In common with traditional contingency plans, the specific plan for the data transfer should be broken down into a number of phases. These should include the following steps:

10.9.1 Risk analysis

Management should identify the realistic and relevant risks that they face in undertaking such a project, and classify them according to likelihood.

10.9.2 Business impact analysis

Management should analyse the potential impact of the occurrence of each of the risk scenarios in order to anticipate its severity and be able to determine the level of response necessary.

The output of this analysis will vary markedly from one organisation to the next because of the inevitable difference between any two projects in scope, significance of the data, and importance of the migration to ongoing operations. For that reason management should attempt to perform both the risk analysis and the business impact analysis from a blank sheet rather than try to modify the analyses performed in other organisations.

It is also essential that these analyses draw upon the experience and expertise of users from across the organisation. IT users will have a clear picture of the technical aspects of the migration, but it is the data owners and key users who will be in the best position to explain which data are used when and how, and be able to identify the most critical elements and functionalities of the data for ongoing business activities.

10.9.3 Design of monitoring activities

Within a well-managed project it is essential to be able to identify at an appropriately early stage any errors, exceptions or anomalies that may be occurring. In the same way as ordinary control procedures over business processes contain internal control activities that allow exceptions to be highlighted and addressed by management, the migration project should contain control activities that will identify issues.

In practice these will often take a similar form to the control activities described in Section 10.6.5 above.

10.9.4 Design of contingency activities

An essential part of the contingency planning process is to have decided in advance what the possible contingency activities will be. It is a truism of general contingency planning that the time to think is before the incident strikes; once the crisis is underway, it is the time to act.

Management should therefore decide upon and document their possible courses of action in the event of each kind of issue. The responses could range from essentially do nothing at all, continuing with the transfer and simply taking note of the problems encountered, to halting and reversing the process. Between these extremes could lie options such as reversing and reperforming parts of the transfer or extending the timeframe allocated for completion of the process.

10.9.5 Decision taking

Management needs to have defined the key decisions that will need to be taken, the sequence in which they will be taken, the expected timings, and who has responsibility for each decision. How this is organised will depend on the organisation's culture and practices; what is important is that there is complete clarity throughout the transfer team over the correct person to turn to for each decision at each stage of the process.

The most critical decisions will be any that radically affect the direction and completion of the transfer process. The key among these will be the final "go live" decision, taken after verification procedures have been completed and confirmations have been received by the senior members of the project team.

Once data have been successfully transferred, the focus of risk management and controls activities will change to be concentrated on monitoring and responding.

10.10 Awareness of problems

Management will need to implement procedures for being informed of issues and incidents related to the transferred data.

10.10.1 Definition of incidents

The first element among these procedures should be a clear definition of possible events and what would constitute an incident requiring being recorded. Not every quirk or anomaly is significant enough to require management action, but it is essential that those that do are clearly identified and highlighted.

10.10.2 Being informed of incidents or detecting them

Management need to determine how they wish incidents to be identified and reported. Typically within a framework of internal controls there will be scope for both detailed controls that identify errors and omissions and monitoring controls that consider the performance of controls and indicators over a longer period. Management will typically wish to employ a mixture of such controls, both manual and automated, in order to provide the widest possible coverage of business processes. Typically they will also wish to provide mechanisms by which users can report, in a structured and coherent way, any inconsistencies or potential anomalies that they perceive in the course of business activities, perhaps in the context of generating a report that presents unexpected information or in encountering a data item that looks unusual.

10.10.3 Incidents affecting the confidentiality of data

Management need to be aware of any incidents in which the confidentiality of data is potentially or actually compromised. These could take the form of incidents of data leakage, such as data being downloaded inappropriately or a legitimate data download being copied or its support being lost or stolen. They could also be the result of the abuse of access rights by legitimate users, either within the organisation or at the service provider, of the creation or use of unauthorized or inappropriate access rights, or of simple illicit access attempts from outside the organisation.

10.10.4 Incidents affecting the integrity of data

The integrity of data can be affected by the modification of data structures and by other database maintenance activities that could lead to data elements being lost or misplaced.

10.10.5 Incidents affecting the accuracy of data

The accuracy of data can be affected in several ways. It can happen in the course of normal daily operations, with a user modifying, inadvertently or deliberately, data elements. It can also happen as a result of administration actions at a database or platform level as technical users rebuild or consolidate data or perform database maintenance activities. It can also be affected by complete or partial restore of data.

10.10.6 Respect of the reporting lines

Every organisation is free to define and implement its own reporting lines for incident handling. Good practice exists and is promoted by frameworks such as ITIL, but there are no absolute rules that need to be followed. What is essential for management, however, is the assurance that the reporting lines as defined are being respected by users, so that incidents are being appropriately identified, appropriately reported, appropriately escalated, if necessary, and appropriately addressed. Organisations are free to choose whether they appoint a specific incident manager attached to the helpdesk function, or whether they wish

to schedule regular meetings with operational management to discuss recent and outstanding incidents. Senior management simply need to understand their roles and more essentially their responsibilities in respect of incident management, and organise themselves accordingly.

10.11 Problem tracking over time

A key principle recommended by frameworks such as ITIL concerns the systematic and regular tracking of problems. This is critical to be able to perform trend analysis and identify issues that recur or which follow a discernable pattern.

10.11.1 Recording incidents

Procedures and a supporting platform should be implemented through which incidents can systematically be recorded. This can be largely user-driven, with individual users having access to a helpdesk or incident reporting system; for reasons of consistency and efficiency many organisations prefer to restrict write access to such platforms to specialist, usually helpdesk, staff, who receive calls or mails from users and log the details in the system. Given that the quality and consistency of the information in this system provide all of its value, it is often beneficial to restrict update access to a small number of staff who have been trained in the importance of consistency and accuracy.

10.11.2 Follow-up and closure

In the same way as a log file is of no value to management unless it is reviewed critically and any issues addressed, a log of incidents is of no value unless management can be certain that each incident has been followed up and closed, or is at least in the process of being followed up. To ensure that this is being done, management will often implement a schedule of regular – perhaps monthly – reviews of system statistics and open items. Management will be relying upon data about the data in order to ensure that items are being addressed and closed in an effective manner.

10.12 Monitoring the quality of service

As discussed in Chapter 4, it will be essential for management to be able to monitor and evaluate the quality of the service that they are receiving from their service provider. Service provision can easily turn into a metaphorical financial black hole for customers, who receive regular bills while having little visibility over the reality of the service for which they are paying. In order to ensure that the services being provided meet their requirements, correspond to the contract terms, and represent value for money, management need to obtain regular, up-to-date and accurate information about the quality of service.

10.12.1 Regular reporting

Management should insist on regular reports from the service provider covering a number of areas of interest, including system uptime and data availability, access records, maintenance activities, processing and storage statistics, issues encountered and potential security issues. Ideally such reports should be produced at a frequency between weekly and quarterly, depending on the quantity of information and the perception of risk on the part of management, and should contain historical data for comparison purposes and notice of perceived future issues or opportunities.

10.12.2 Local dashboard

A very useful tool for management to visualize performance and highlight potential issues is to implement a dashboard in which key indicators are selected and their performance over time compared to budgets and expectations. With carefully selected indicators and accurate and valid data, such a tool can provide management with a reasonable level of comfort over the quality of service provision and also indicate areas of priority for further investigation and, if necessary, rectification.

10.12.3 Access to statistics

The quality of decision taking is limited by the quality of the information available as a basis for those decisions. Management should therefore ensure that they identify the statistics upon which they will need to rely for their monitoring and decision taking, and also determine how those statistics will be obtained. Any or a combination of three methods could be employed: receiving performance data directly from the service provider; extracting data from the service providers' systems; or generating the data locally based on the analysis of the use of systems.

10.12.4 Management meetings

A key element in monitoring service provision is the scheduling of regular client service meetings with representatives of the service provider at which the regular reports can be discussed, issues raised, trends highlighted and discussed, and any performance issues clarified and explained. Not only do such meetings allow the effective communication of concerns and questions between management and the service provider, but they also allow the development of more effective and efficient working relationships between the two parties.

10.13 Monitoring of technical and compliance developments

In opting for third-party solutions, organisations are committing themselves to monitoring changes in two different rapidly evolving fields.

10.13.1 Technical developments

Even though it will be the responsibility of the service provider to maintain the environments and infrastructure, the data owners retain the responsibility to be aware of any changes that might be significant in helping to protect their data or in providing better solutions for their specific requirements.

A great deal of the success of responding to changes will depend on the relationship between the client and the service provider. Management will need to ensure that they are comfortable with the philosophy of the service provider, in that the service provider can be relied upon to make suggestions and recommendations that are specifically in the client's best interest rather than trying to leverage the existing relationship and ongoing contracts to sell additional services that may not be entirely necessary.

10.13.1.1 Improvements in cryptography

It is almost a given that the initial set-up and operation of the outsourced services will involve cryptography in some form, be it over data at rest or over data in motion, or both. Cryptography is a dynamic field and both the service provider and the client will need to be aware of changes and improvements in the

technologies in use and in the environment in which they operate. Changes might mean that more efficient methods of encryption and decryption become available, which would allow more rapid access to data. Changes might mean that stronger algorithms and encryption techniques become available at comparable cost and with no apparent impact for end-users, making an upgrade a reasonable and viable option. Changes might also occur with researchers discovering flaws in the encryption system in use, making an upgrade or reimplementing necessary.

Clearly corporate management cannot be expected to have a detailed and up-to-date understanding of the technical details of such questions, but they should be in a position to be informed of such developments, through their in-house technical staff, through corporate advisers, or through the service provider itself, and they need to be able to understand the issues and recognize the significance of the potential impacts of changes in order to evaluate recommendations and take the most appropriate decisions.

10.13.1.2 Improvements in virtualisation

Virtualisation technologies are also a rapidly developing field and management need to be appropriately informed about progress and opportunities in order to be able to respond effectively to suggestions and recommendations for change.

10.13.1.3 Improvements in platforms and infrastructure

The technical details of the platforms and infrastructure used by the service provider are broadly speaking a matter for the provider. Management should retain an interest, however, in developments in the market in order to be able to evaluate proposals for, or announcements of, change. In particular they should be aware of the likely improvements in the quality of service and of the likely risks inherent in a platform migration, even if the service provider will be responsible for all aspects of any changes that are made.

10.13.1.4 Improvements in storage technologies

Similarly, storage technologies have been undergoing constant and rapid evolution over recent decades and management should be aware of the advantages and limitations of newer technologies in order to be able to take informed decisions.

10.13.2 Compliance developments

The domain of compliance is rarely static, with changes made to existing requirements on a regular basis and new requirements introduced in response to specific concerns or to the wider application of technology in different fields.

10.13.2.1 Changes to existing requirements

Management need to be aware of, and capable of responding to, changes in requirements with which they are already compliant. This can occur when a piece of legislation or a guideline is updated to reflect changes in technology, focus or concern. The specific challenge in this respect is for the client to ensure that the service provider does indeed take all necessary actions required of it within the timeframe appointed.

10.13.2.2 Introduction of new requirements

New requirements can apply to organisations, either because new legislation or standards have been introduced or because the organisation has changed an aspect of its business or data handling that brings it into scope for requirements that are new to it. Once again, the challenge for management is two-fold: be aware of the impact and obligations created under the new compliance standards, and ensuring that all necessary measures are taken, and can be documented, both in-house and at the service provider.

10.14 Obtaining periodic and ongoing comfort about the service provider’s security management

The final key aspect for management to consider and define is how they will obtain comfort over the security of their data. As discussed in Chapter 6, access to and use of their data will be invisible to them and so they will have to rely upon the internal controls designed and operated by the service provider.

10.14.1 Comfort obtained from third parties

Traditionally the most robust form of assurance can be provided in the form of a service auditor’s report following standards such as SAS 70 Type 2 or ISAE 3402, if such a report is available. Management must of course review any such report to ensure that it covers the whole of the period in question, the infrastructure and internal controls in question, and that any exceptions are highlighted and explained.

Alternatively management could consider sending their own auditors, or specially mandated external auditors, to perform an appropriately scoped review of the service provider’s internal controls. In practice, however, this can be very difficult to organise and is avoided by service auditors because of the disruption it can cause to their daily operations.

10.14.2 Comfort obtained through internal procedures

Alongside, or instead of, external comfort, management should consider how assurance can be obtained through internal procedures. Depending upon the level and scope of access to data and databases that technical users retain post-transfer, they might consider performing their own reviews of access logs and user access rights. They might also consider reviewing changes in data elements, highlighted through the comparison of regular data extractions, to ensure that no unexpected or unexplainable changes have been made to data.

10.15 Framework for Management Comfort

This following framework, in the form of a simple template, sets out guidance for the minimum preparatory and ongoing work that management should consider when planning and performing a data transfer project, from internal data storage onto the cloud.

Item	Objective	Description
1	Ensuring that the key concepts and the related vocabulary are understood	Develop a project background document that sets out the scope and objectives of the project and defines the key concepts and vocabulary.

Item	Objective	Description
		<p>Ensure the distribution of this document to all members of staff who will participate in, or have oversight over, the project.</p> <p>Organise information and training sessions for the project team to ensure that there is no confusion over the strategy, direction and objectives of the project</p>
2	Ensuring the adequate scoping of the project: types of data, uses of data, and storage options	<p>Develop a project scoping document that clearly sets out the scope and range of the data to be included in the project, as well as the rationale for this and a clear definition of terms.</p> <p>Ensure that all relevant sources of structured data are identified and documented, with details of the nature, volume and utility of the data that they contain. This documentation should also, for the sake of completeness, refer to sources of structured data that are not considered to be in scope, with an explanation of the decision.</p> <p>Ensure that all relevant sources of unstructured data are identified and documented. This documentation should contain details of the nature, volume and utility of the data that they contain. It should also, for the sake of completeness, refer to sources of structured data that are not considered to be in scope, with an explanation of the decision.</p> <p>Ensure that the different types of storage media that are considered to be within scope are clearly and completely catalogued. This listing should include: standard, IT-supported supports such as central databases, file servers and workstations; local and often temporary supports such as external drives; transport media such as USB keys and memory cards; and centralized and local backup media.</p> <p>Ensure that the historic and potential uses of the data are clearly documented, with focus on areas where datasets are intended for combination in order to yield additional information.</p> <p>Ensure that the quality of the data is assessed and documented. Particular attention should be paid to data that are of potential value but doubtful quality in their raw form, to data that duplicate other datasets, and to data that correspond to distinct periods, in order to develop the most useful map of the data that exist and what coverage they give.</p>
3	Understanding the risks	<p>Develop an exhaustive and practical analysis of the risks to the organisation that the data accumulation and migration project will present.</p> <p>Ensure that legal risks, such as those related to the retained responsibility for and ownership of data regardless of their location, are documented, communicated, understood and approved.</p> <p>Ensure that operational risks, such as reduced control over the availability of data, the speed of access to data, and the response to support or</p>

Item	Objective	Description
		<p>change requests, are documented, communicated, understood and approved.</p> <p>Ensure that technical risks, such as dependence on the service provided by the third-party for the availability and integrity of data and on connectivity to the service provider, are documented, communicated, understood and approved.</p>
4	Roles and responsibilities	<p>Develop clear and consistent documentation setting out the split of roles and responsibilities between local management and the service providers. This is essential to avoid both disputes in the event of problems being encountered, and incidents that occur or are aggravated because of confusion or hesitation.</p>
5	Planning the transfer	<p>Develop a complete and robust plan for all phases of the data processing and transfer, considering the overall objectives for this phase and the significant issues that may arise.</p> <p>Ensure that the data to be moved are clearly and consistently identified, documenting provenance, contents and relevance.</p> <p>Ensure that data are cleaned and restructured in a controlled and consistent manner, retaining relevant contextual information and metadata where possible but removing inconsistencies, inaccuracies and irrelevancies.</p> <p>Ensure that a standard format has been designed and communicated for the datasets emerging from the preparation process so that users can be confident of the consistent formatting and portability of data, bearing in mind that data retrieved from different legacy platforms may require specific efforts to conform to the standard.</p> <p>Ensure that a policy is in place for determining what should happen to local copies of data being transferred. This policy should include a description of the workflow and the decision-taking process.</p> <p>Ensure that procedures are in place to determine and confirm the completeness, integrity and accuracy of the data transfer process. These procedures should include ongoing tests of the data, review tests to be performed by supervisors and internal auditors, and adequate and documented approvals by relevant managers.</p>
6	Handling other data appropriately	<p>Develop thorough and consistent plans for the post-processing treatment of data not selected for transfer to the service provider.</p> <p>Ensure that any reorganisation and restructuring of the data are performed consistently and in conformity with the standards for data structures, allowing for the simplest possible future access and manipulation of the data.</p> <p>Ensure that data intended for deletion are both clearly and reliably marked, that all copies of the data are identified and deleted, and that where</p>

Item	Objective	Description
		secure deletion is required, this is carried out effectively and traceably.
7	Addressing legal considerations	<p>Policies and procedures should be in place to identify relevant legal considerations made significant by the data preparation and migration project.</p> <p>Ensure through due diligence procedures that no requirements or obligations are being contravened. The review should include specialist professional advice that is documented and consider both national and international guidelines, as well as industry-specific requirements.</p> <p>Ensure that the subjects of data have been properly informed, in a tangible and verifiable fashion, of the fact that changes are being made to data storage locations and procedures, and of any impact that this will have on them. This will include confirmation of the continuation of existing compliance activities and explanation of any new compliance requirements that are being addressed.</p> <p>Ensure that regulators and authorities have been informed, formally and in writing, of the changes made to data management conditions and of the activities undertaken to ensure continued compliance with compliance regulations. This may involve planning and performing a compliance audit within a set period in order to provide evidence to these third parties that compliance has been maintained.</p>
8	Contingency planning for problems with the transfer	<p>Policies and procedures should be in place in order to respond to any problems encountered during the transfer process that could have an impact on the success of the project or on the organisation's capacity to continue to function properly.</p> <p>Ensure that sufficient risk analysis has been performed to allow management to be comfortable with undertaking the transfer project, in full awareness of the nature and range of the risks to the organisation that the project represents.</p> <p>Ensure that sufficient business impact analysis has been performed to inform management of the likely consequences of any of the risks materializing and thus be able to take considered and appropriate actions in response.</p> <p>Ensure that appropriate monitoring activities have been designed in order to allow management to supervise the transfer process and detect any errors, omissions or anomalies before the final go-live decision is taken.</p> <p>Ensure that appropriate contingency activities have been designed to address any of the potential risk scenarios and correct any errors or mitigate any issues encountered during the transfer process.</p> <p>Ensure that the decision taking process is approved, documented and communicated before the transfer</p>

Item	Objective	Description
		<p>process commences, in order to allow project team members to know who to refer to at each stage of the process in order to receive the approvals they require, and who should be present at the end of the process to give the final approvals needed to go live.</p>
9	Awareness of problems	<p>Policies and procedures to ensure that management are informed of issues and incidents related to the transferred data will need to be designed and implemented.</p> <p>Ensure that the definition and scoping of incidents is clearly set out and communicated to ensure consistency of reporting, reaction and communication.</p> <p>Ensure that reliable procedures are in place for appropriate staff to be informed of incidents or detect them, and to escalate them if necessary. These procedures should include the logging of incidents by users in a consistent and consultable way, the highlighting of incidents identified by automated procedures, and the regular and systematic review of logs, reports and process outputs by trained and responsible staff.</p> <p>Ensure that reliable procedures are in place to identify incidents affecting the confidentiality of data.</p> <p>Ensure that reliable procedures are in place to identify incidents affecting the integrity of data</p> <p>Ensure that reliable procedures are in place to identify incidents affecting the accuracy of data</p> <p>Ensure that reporting lines are clearly defined, documented, communicated and respected. This should apply to two levels: communicating incidents to the specialist staff responsible for incident handling; and these staff communicating upwards in the contexts of escalation or periodic reporting.</p>
10	Problem tracking over time	<p>Policies and procedures should be in place in respect of the systematic and regular tracking of problems. It is essential for management to be able to perform trend analysis and identify issues that recur or which follow a discernable pattern.</p> <p>Ensure that, in common with the control objectives set out in respect of Point 9 above, that incidents are recorded consistently, accurately and completely.</p> <p>Ensure that there is consistent and timely follow-up and closure of incident reports. Typically this can be enforced by regular reviews of statistics and open items, discussed by senior management with relevant line management.</p>
11	Monitoring the quality of service	<p>Policies and procedures should be in place to ensure that the services being provided meet their requirements, correspond to the contract terms, and represent value for money. In order to achieve this, management need to obtain regular, up-to-</p>

Item	Objective	Description
		<p>date and accurate information about the quality of service.</p> <p>Ensure that regular reporting is received from the service provider in respect of the key performance indicators agreed as effective and informative markers of the quality of service.</p> <p>Ensure that a local dashboard is designed and maintained that will allow trends to be identified and potential areas of concern to be highlighted.</p> <p>Ensure that access to, and creation of, relevant statistics is possible on a regular, systematic and updated basis.</p> <p>Ensure that regular management meetings are held with the service provider at which reports can be discussed, incidents highlighted and resolved, and current and future requirements and potential issues discussed and prioritized.</p>
12	Monitoring of technical and compliance developments	<p>Policies and procedures are needed to monitor ongoing developments in order to anticipate and be able to respond to changes and new requirements.</p> <p>Ensure that technical developments, such as advances in cryptography, access methods and data storage, are monitored, analysed and discussed in order to seize opportunities and avoid issues arising.</p> <p>Ensure that compliance developments such as changes in existing requirements and the introduction of new requirements are monitored, analysed and discussed in order to seize opportunities and avoid issues arising.</p>
13	Obtaining periodic and ongoing comfort about the service provider's security management	<p>Policies and procedures should be developed for obtaining sufficient comfort over the design and operational effectiveness of the security provider's security management, in order for management to be confident that their responsibilities are continuing to be fulfilled.</p> <p>Ensure that consistent and pragmatic procedures are in place for regularly acquiring the information and confirmations from third parties that are necessary to maintain an appropriate level of assurance.</p> <p>Ensure that consistent and pragmatic procedures are in place for regularly generating or acquiring, through internal procedures, the information needed to supplement or replace the assurance received from third parties in respect of proper security management.</p>

Table 10.1 Template for management comfort

This template covers the whole of a data identification, preparation and migration project. The different sections it contains are reasonably self-contained and autonomous, however, and can be used in isolation or as a small group of sequential steps for smaller or more precisely scoped projects. Indeed, as will be discussed in Chapter 11, this is how this model has so far been partially validated in industry.

10.16 The Cloud Security Alliance Matrix

Previously mentioned in Chapter 5 in reference to the survey of cloud behaviours carried out in association with ISACA, the Cloud Security Alliance is “a member-driven organization, chartered with promoting the use of best practices for providing security assurance within Cloud Computing, and providing education on the uses of Cloud Computing to help secure all other forms of computing” [115].

It offers training and certifications to organisations deploying applications and workloads in the cloud, and provides guidance in the form of white papers, standards, and publications arising from working groups. They also publish matrices of suggested controls, including a wide-ranging Cloud Control Matrix, the third version of which was published in September 2013 [116].

This matrix follows a similar structure to the above framework, breaking down the field of internal controls in relation to cloud computing into a number of subdomains. These are:

Subdomain	Number of controls
Application and Interface Security	4
Audit Assurance and Compliance	3
Business Continuity Management and Operational Resilience	12
Change Control and Configuration Management	5
Data Security and Information Lifecycle Management	8
Datacenter Security	9
Encryption and Key Management	4
Governance and Risk Management	12
Human Resources	12
Identity and Access Management	13
Infrastructure and Virtualization Security	12
Interoperability and Portability	5
Mobile Security	20
Security Incident Management, E-Discovery and Cloud Forensics	5
Supply Chain Management, Transparency and Accountability	9
Threat and Vulnerability Management	3

The control specifications within each domain are linked explicitly to aspects of architectural relevance (such as physical security, network, and storage), the cloud service delivery models (SaaS, PaaS, IaaS), and to a number of specific methodologies and standards. These include AICPA Trust Service criteria, Cobit, ENISA guidance, FedRAMP (the Federal Risk and Management Program, a US government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services), HIPAA, ISO/IEC 27001, NIST SP800-53, and PCI DSS. There is thus a direct linkage between the suggested areas of control and the frameworks for control and compliance that are commonly in use.

In common with the framework presented in Section 10.15, the control specifications in the CSA matrix are fairly high level, setting out the objectives and the overall principles without being prescriptive about the precise form the control should take, nor how it should be operated.

The focus is entirely upon aspects of cloud computing. As such, there is some overlap with the framework presented above, particularly in respect of risk management and data access and preparation. Its scope is more restricted, however, and also focused more specifically on US standards such as FedRAMP and HIPAA.

This restricted scope is offset by the increased detail into which the matrix goes in each of the subdomains when compared with my proposed framework.

I suggest that my proposed framework can be used alongside and above the CSA matrix. Indeed, one of the design decisions behind the framework was to maintain a sufficiently high level of abstraction, almost to the extent of creating a meta-framework, that other templates or matrices could be referred to or incorporated into relevant sections as deemed appropriate. Not every aspect of the framework will be of the same significance to every organization; some aspects will prove to not to be applicable in any detail and so can be addressed quickly and simply, while others will need greater attention and the further details provided by such sources as the CSA matrix will be valuable additions to the process.

10.17 Conclusion

Scoping, planning and carrying out a combined data preparation and transfer project is a lengthy and complicated task for any organisation. Good practice and professional experience demonstrate that it is essential that the project plans are detailed, appropriate, approved, documented and communicated in order to ensure, as far as is possible, that sufficient overall control is operated over the process and that risks are addressed in an appropriate way.

The attention paid to risk analysis at an early stage in the planning ensures that all relevant risks are identified and that appropriate measures are designed in response; it also ensures that resources are not wasted in respect of inappropriate or irrelevant risks.

The detailed structure of the analysis, particularly with regard to the technical aspects of the project within the overall framework, ensures that focus is appropriately placed on the higher-level, management facets of security and operations. In developing project plans, there is a common tendency to concentrate on the details of processes and of technical issues. This can lead to situations in which inconsistent and non-optimal decisions are taken about parts of the project in isolation. Documenting everything in a single, structured and approved plan assists in avoiding such situations and in maintaining a high level of consistency and completeness.

The success of any project is of course a direct function of the quality and competence of the project managers. The framework above is not, and cannot be, a panacea to be applied blindly and without experience and judicious analysis. It can serve, however, as a framework and checklist for competent and experienced managers to focus their efforts and direct their resources.

Chapter 11 Merits and Limitations of the Proposed Solution

11.1 Introduction

“Every model is wrong, but some are useful.”¹⁴

The results of the survey presented in Chapter 5 and the analysis of risks and opportunities presented in Chapters 6 to 9 have been subjected to peer review and presentation, and have led to a chapter in a book to be published by Springer Verlag in 2014. The text of that chapter, and explanatory notes, can be found in Appendix C.

As presented in Chapter 10, the project management framework represents a solid and practical general structure that will allow project and corporate management to plan, prepare and carry out a data cleaning and migration project with the confidence that the major areas of risk had been considered and that the opportunities to design and implement control activities had been identified.

Nonetheless the solution does present a number of limitations.

11.1.1 General versus specific

In its current form, the framework is very general, applicable to all types and sizes of organisation and not specific to any industry or sector. This was a deliberate choice in order to create a framework that was as widely applicable as possible rather than be limited.

There are two main areas in which the framework could be extended and developed in order to be more specific. The first is in respect of specific aspects of the technologies in use. For example, the section in which questions of the use of encryption are discussed could be further developed to include specific objectives and requirements related to specific technologies and their particular modes of use. This could take the form of a table of current requirements and standards for key length or preferred protocols. Clearly this would need to be reviewed and updated on a regular basis as the environment evolves, with particular attention paid to developments in which particular implementations of a method or protocol had been demonstrated to contain weaknesses.

The second area in which more detail could usefully be added would be in respect of specific territorial or industry requirements, or even both. If the framework were to be followed by a financial services company in Switzerland, for example, it would be of great benefit to the users to include clear references to the most recent Swiss regulations and legislation, including the Data Protection Laws, instructions from FINMA (the

¹⁴ Box, G. E. P., and Draper, N. R., (1987), “Empirical Model Building and Response Surfaces”, p. 424

regulatory body for Swiss financial institutions), and the Code des Obligations that underpins all of Swiss commercial activities.

11.1.2 Audit and internal control perspective

The framework is designed and written from a clear controls and audit perspective, with the objective of compelling users to design and implement processes and procedures that correspond to control objectives and genuine risks.

This approach is based upon a number of assumptions about both project and organisational management.

Firstly, it assumes that they share a common vocabulary and familiarity with audit standards and practices. In practice, this may not necessarily be the case. Concepts such as control objectives, risk assessment and control activities (as set out in Chapter 2) are widely employed within audit firms and internal audit departments, but it should not be assumed that they will be familiar to all participants in a project team, especially not members from a more technical and non-audit background. Project managers and sponsors should be aware of the need to train and coach project team members and other key users affected by the project on these matters, and also take into account the costs, both direct and indirect, and the time needed to perform such training.

Secondly, it assumes that the basic principles are understood and not open to discussion. This may not be the most unreasonable assumption, given the close linkage between the framework and the industry standard best practice models referred to in Chapter 8, but at the same time the philosophy of the framework is supportive rather than prescriptive and different approaches, if argued and considered to be valid, should not be automatically discounted. It would be important to emphasise to management the importance of consistency, however. Experience of project management indicates very strongly that a particular model should be followed or not followed, but that following it half-heartedly or trying to cherry-pick what appear to be the key elements is rarely a successful approach. When such an approach is adopted, be it from a desire to cut costs or corners, a result of a lack of detailed appreciation of risks, or a lack of project management experience, the results are frequently unsuccessful, with inconsistencies in approach, gaps in processes and procedures, a lack of financial control, and an absence of overall strategy, control and identified success factors.

An example of this was alluded to in Section 10.1 above. In attempting to improve the overall level of security within an organisation, for example, it is necessary to begin with the design of an overall approach that is consistent and structured. There is a tendency, however, to attempt quick fixes by attacking obvious and visible details, such as password disciplines. Such disciplines can be introduced or made more restrictive very simply on many platforms by modifying some simple parameters, so that for example a minimum password length can be enforced, password expiry can be introduced, password complexity can be enforced, and password histories can be maintained in order to prevent users keeping and repeating simple passwords. Such measures do indeed form a significant part of good logical access controls. But if they are introduced outside the context of a structured approach to security, without proper controls over the creation, modification and deletion of user rights, without the monitoring of user activities, and without verification of how compliance is being achieved (for example, are users writing down their passwords because of difficulties in remembering them), then the measures may both be ineffective and give rise to misplaced confidence in the overall quality of logical security.

Thirdly, there is an underlying assumption in respect of both the overall and the detailed objectives behind the project. It is assumed that the objectives will automatically be as set out in the model: compliance, security, privacy, and efficiency. Realistically it should be recognized that this will not always be the case. Even if management have a philosophy and project management approach that includes the recognition of such broad objectives, practically it may be that they are concentrating more on cost saving and financial management rather than compliance and control. Equally, they may be committed to performing the strict minimum in order to ensure compliance with the letter of the law rather than making additional efforts to ensure a more complete and wide-ranging compliance.

11.1.3 Supportive rather than prescriptive

The framework is deliberately non-specific on how objectives are to be achieved. The outcomes are what are measured rather than the methods employed to reach them, with the expected outcomes being robust and documented results. This design decision was based on two main factors.

Firstly, there is a strong and well-accepted history in the field of audit and control frameworks of adopting such a non-prescriptive approach. The case of AS5 (see Chapter 8) serves as an instructive example of this: it specifies the outcomes of control measures and sets out what auditors should expect to see and be able to verify, but it does not prescribe in any way what the expected controls should be, particularly in respect of internal controls relating to the management of information systems. On this basis, a supportive rather than prescriptive framework appears to be entirely reasonable as an approach.

Secondly, there is the example given by, and experience drawn from, highly prescriptive models for internal processes and controls such as COBIT and ITIL. These methodologies contain detailed checklists and exhaustive catalogues of possible internal control activities, wide-ranging enough to be applied to almost every circumstance and control objective. Such detailed documentation is very useful to process and controls managers who need to ensure that their control objectives and identified risks are properly covered.

The danger with such detailed materials is that in practice the specific evaluation of control objectives and risks can be forgotten. In practice there is a tendency for business managers, if they are obliged to implement internal control activities, to implement the whole catalogue. This is impractical and counter-productive for a number of reasons: firstly, it is inefficient, with resources directed towards the operation of controls that may not add any value or offer any additional comfort to management; secondly, if the purpose behind the controls is not understood by the users, which is often the case if controls exist that duplicate each other or apply to unimportant or insignificant aspects of operations, it is likely that the controls will be performed badly; and thirdly, the greater the number of detailed controls that are implemented, the less likely it is that management will have the resources, time and competence to monitor their operation. Thus control activities are likely to be carried out without any benefit to the organisation, using up limited resources and potentially, as discussed in Chapter 8, providing management with an inappropriate level of comfort over the quality of internal processes and the mitigation of risks.

An additional weakness of such an approach also results from the untargeted selection of control activities. As discussed in Chapter 2, the selection of internal control activities should be a milestone on a process that begins with the definition of control objectives and the identification and evaluation of relevant risks. If an internal control catalogue is referred to indiscriminately, there is always the possibility that the planning and design phase is neglected in favour of simply selecting a wide range of controls. This could have as

consequence situations in which not only a redundant, inappropriate and duplicate control activities are in place, but also that relevant risks have not been properly identified and evaluated, with the consequence that the most appropriate control activities are not in place.

11.1.4 Project management methodology

This framework is very firmly based on a particular perception of how large projects should be managed within an organisation. This philosophy and approach is consistent with good practice as described by a number of authorities including COBIT, but it is not necessarily the only valid approach. It is possible that individual organisations could argue that such an approach does not suit them optimally for a number of reasons, such as: culture (smaller, informally-managed family businesses, for example, are less likely to favour formalised and well-documented project procedures); experience (organisations may not have any experience in performing such formal projects and be unwilling to invest in training or support); resources, both human and financial (because formality and documentation incur costs); and timings and duration (management might be unwilling to commit themselves to a longer formalised project if their assessment of risks and rewards convinces them that a smaller, quicker and less formalised project could yield satisfactory results).

In the author's experience the old adage of "failure to plan means planning to fail" applies without exception in such circumstances, and the choice not to follow a structured and formalised framework would need to be very carefully justified indeed, but organisational management will always have the final say.

11.1.5 Evolution of assumptions

Any framework and methodology is by its nature a snapshot of prevailing good practice and current technical thinking. This framework is a product of its time, reflecting the current immature state of several of the technologies in use and also drawing strongly on common trends in thinking in respect of corporate governance and internal controls over the two previous decades. The whole governance, control and audit world is undergoing constant evolution, however, and there is no certainty that what applies to a given organisation or situation today in terms of acknowledged best practice will continue to apply next year. The core assumptions behind the framework should therefore be revisited periodically in order to confirm that they continue to be valid.

In practical terms, however, it is difficult to envisage ways in which the whole current structure of good practice could be overturned sufficiently to invalidate completely any models that draw upon such structures and proven principles.

11.1.6 Direct linkage to security objectives

A common aspect of professional audit programmes is that each audit activity is directly linked to the assumptions, assertions or objectives that underlie the corresponding business activity. For example, financial audit programmes will indicate which of the management assertions are being examined. These assertions are commonly stated as being as follows:

- In respect of transactions (Income statement): Occurrence, Completeness, Accuracy, Cut-off, and Classification.

- In respect of Accounts balances (Balance sheet): Existence, Rights and Obligations, Completeness, and Valuation and Allocation.
- In respect of Presentation and disclosure: Occurrence, Rights and Obligations, Completeness, Understandability, and Accuracy and Valuation.

Similarly in the domain of the audit of internal controls, each audit activity will examine the four key control objectives of the internal control activity in question. These four control objectives are commonly known as “Caviar” from their initials: Completeness, Accuracy, Validity and Restricted Access.

Such linkage is extremely useful for both management and users responsible for processes and controls and for the auditors who are examining their operation. By considering the coverage of the objectives and assertions provided by the different subprocesses or activities within a given process, management will be able to identify any gaps in their structures. For example, for the overall process in respect of managing user access rights, management would want to see that all four of the Caviar objectives are represented at least once in respect of the whole population. From the perspective of efficiency, management would also want to ensure that, say, Completeness is not ensured by five or six control activities when two or maybe three would be sufficient.

Auditors as well rely on such linkage in order to assess the completeness of their work and ensure that their conclusions are valid in respect of the coverage of assertions and objectives.

Based on these models, a case could be made for extending the framework in Chapter 10 to include an indication of which information security objective – Completeness, Integrity or Availability, or one drawn from the wider sets of objectives discussed earlier – is ensured by each activity. Clearly this would be of little value at the highest level of the framework, where each section by necessity covers all three main objectives, but a more developed and prescriptive framework could usefully contain such detail at the level of the detailed project steps.

11.2 The feasibility of the model

This framework remains theoretical, having not been tested in a full way in practice. Elements of the framework have been tested in a variety of contexts, however, in the early stages of planning and performing the data identification and preparation, and also in the risk identification phases.

11.2.1 Data identification and preparation

The early phases have been completed at a medium-sized company that became aware that it was unable to control the data being generated by and stored within its main systems. The IT Manager had for some months been reporting capacity and performance issues to business management and requesting approval for investment in additional storage across the local network, but did not have the internal resources to research the problem in sufficient detail to identify the exact nature of the additional data accumulated month by month, nor where the data were being produced. Before approving capital expenditure, management wanted to understand the core causes of the problem and, above all, to know whether it was a situation that was likely to recur and if so with what frequency.

A data identification exercise was thus approved and scheduled, in which the contents of file servers were examined and analysed and the size of central databases was tracked over the course of several weeks.

The results were interesting from a data management perspective.

11.2.1.1 Database growth

Firstly, the central operational – non-financial – databases were growing at a rate of around 2% every month as a result of data flowing in from core activities through automated interfaces. This result surprised both operational and IT management, who had not been monitoring the quantities of data being generated and transferred. One of the action points resulting from the exercise was that the configuration of the systems and the utility of the data being produced would be reviewed in order to determine whether all these data were necessary or useful. It was generally accepted that such consistent growth in the size of databases was unsustainable in the long-term and would only lead to additional processing and performance issues as time passed.

11.2.1.2 Report generation

Secondly, a hitherto unremarkable process was identified as being at the core of the problems. This process ran periodically to generate reports of operational performance and activities, both by exception and in total, and these reports were transferred onto a fileserver to be available for consultation by relevant managers. Interviews ascertained that relevant managers were unaware of the existence of these reports, while the IT Manager was unaware of the scheduled performance of the process in question: research in the system documentation suggested that the functionality had been enabled by one of his predecessors and not widely communicated. Although the individual reports were not of great size, they were increasing over time in proportion to the size of the operational databases, and were also accumulating, unwanted and unread, on the file servers. Another action point resulting from the exercise was to review the nature and content of these reports and decide which, if any, were useful as management supports. The generation of others could be stopped.

11.2.1.3 Multiple copies of files

Thirdly, the review identified many multiple copies of word processing documents, spreadsheets and presentations stored in different locations across the network. This appeared to be the result of a well-meaning but unstructured approach to data backup in which user directories were periodically copied to different locations on the servers. The utility of this approach was debatable, because a basic tenet of backup and restore is that if users or management do not know that the backups exist or where they are located, they are of no real value. The corresponding action point was to review the age and utility of these backups, keep the useful ones, moving them to a centrally managed and monitored location, and delete the others.

11.2.1.4 Backup and archiving

An overall action point relating to all three findings concerned backup and archiving in general. It was clear that the backup media reflected the live systems that they replicated, being full of redundant and unwanted files. Management therefore undertook to review the backup and archiving strategies, identifying data that no longer changed or were of lesser immediate relevance to be archived and removed from production systems, and determining which data should be included in which backup sets and routines. The backup media themselves, used in traditional rotation style, would gradually be overwritten and any capacity problems resolved without specific and targeted action.

11.2.2 Risk evaluation and data identification

The risk evaluation and data identification phases were carried out successfully at a large international not-for-profit organisation as part of a larger and more broadly scoped review of information systems and information management strategies.

This organisation is politically very sensitive and subject to a great deal of public and media scrutiny, and is obliged to publish regular and detailed reports in respect of certain of its activities. One challenge that management had identified was that they possessed some information that was essentially public, some information that was extremely sensitive, and a lot of information that was possibly sensitive but was as yet unclassified.

At the time of my visit, in the summer of 2012, they were also considering the future strategy in respect of IT. A great deal of IT had previously been outsourced, and among options that management were considering were a complete outsourcing of IT services, with possible use of the cloud.

Management recognised that before taking any decisions about further changes to the IT environment, it would be necessary to identify the risks to which they would be exposing themselves. They therefore commissioned a project to report on two areas of concern: the risks affecting them in respect of data management, with classification and high-level suggestions for managing those risks; and the nature of the data being held within the organisation, with classification and an indication of how each class of data should be prepared.

The risk analysis phase was interesting because it was clear that there was no overall or common understanding of risk strategies and management across the organisation. Individual managers in different departments, such as legal, finance and PR, had their own accurate and appropriate viewpoints, but there was no sense of a corporate standard, if it was not the simple objective of avoiding bad publicity. Through a series of targeted interviews and workshops, the input of line managers and senior managers was obtained and consolidated so that a common understanding of risks in respect of data use, storage and transmission was developed, and these risks were classified and evaluated.

11.2.2.1 Information management

A specific finding of note was that although risks in respect of external access to sensitive information had been considered with some degree of thoroughness in the past, there was essentially no understanding of the risks presented by lax information management within the organisation's premises. Files containing sensitive legal or financial information were regularly left on tables in meeting rooms, while whiteboards were rarely cleaned after meetings and flipcharts rarely replaced. A major outcome of the risk analysis phase was thus a transformation of the security culture within the organisation, to be driven by education, training, and enforcement in the form of spot-checks.

Other action points included the formalisation of the risk management process and the documentation and communication of the results of the analysis.

11.2.2.2 Data identification and classification

The data identification phase involved two parallel streams of work. One involved the development of a data classification based on what managers knew were the main sources and formats of data in use across

the organisation, drawing on the results of the risk analysis phase to identify the most significant information flows outside the organisation and determine which of these included sensitive or confidential data.

The second stream consisted of running queries against the file servers, workstations and external data storage devices within the organisation to see what information was being stored and where precisely it was located. In common with the case previously cited, the review revealed that users stored a great deal of unstructured data in the form of spreadsheets and other documents, many of which were based on data extracted from the central systems. Many of the files were duplicated or stored in different versions in different locations across the network. The results of this stream fed into the first stream in order to ensure that the presence and use of all datasets that were considered sensitive or confidential were noted and explained.

The action points arising from this phase were to formalise and implement a clear data classification structure and to determine which data should be subjected to the most stringent security measures and prevented, if possible, from transmission outside the organisation. In addition, discipline would be reinforced in respect of file storage on site and file transfer, with particular attention paid to the use of USB keys and their use in transferring files from professional to private machines. Indeed, the use of private machines for professional work was to be strongly discouraged and corporate laptops provided to key users as an alternative.

Overall the findings of the review were received positively by senior management, who were made aware of the complexity of information security and how difficult it was to ensure and implement. Indeed, their next project in respect of IT was to commission a gap analysis based on ISO 27001, in order to determine where their major weaknesses were in respect of information security management systems, and with a view to aiming for ISO 27002 compliance in the medium term.

11.3 Future activities and developments

Models and methodologies are never static and should always be subject to review, test and improvement. This is a flexible framework that has been designed to be straightforward to modify or update in an environment that is undergoing constant and unpredictable change.

11.3.1 Walkthrough or implementation

The ideal next step would consist of a full-scale walkthrough of the framework by a large organisation that was undertaking exactly the kind of project discussed in Chapter 10, a review of data sources, a cleaning and rationalisation process, and a migration to a service provider's infrastructures.

Such an implementation would generate a great deal of useful information about the structure, the focus and the weighting of the framework. As is the case most notably with full-scale tests of disaster recovery and business continuity plans, it is only when a plan is fully walked through that its limitations and inconsistencies become apparent. Steps that appeared to follow a logical succession are shown to be interdependent in more subtle ways than had been expected, and milestones and checkpoints are shown to be more difficult to reach in practice than in theory.

In common with the tests of contingency plans, the project should be followed by a structured and systematic post-implementation review in which the success of the project is evaluated, the gaps or inconsistencies in the framework are identified and discussed, and recommendations for improvement are proposed and accepted.

11.3.2 Updates of the theoretical and structural basis

In the absence of such an implementation, the theoretical basis for the framework can always be reconsidered in order to make improvements. There are two main areas in which the framework can be improved: depth and current applicability.

Depth would involve both revisiting the individual areas already included in the framework under sections 1 to 13 to see where further pertinent guidance could usefully be given, and considering whether additional sections should be added. An example of the former would be increasing the guidance on cryptographic measures as discussed in Section 11.1.1 above; an example of the other would be adding a section to clarify the choice and significance of data formats, ensuring that the organisation does not tie itself into the use of closed or proprietary formats, say, in an environment where continued access to data over a long period and using a range of technologies is one of the most fundamental objectives of the project.

Ensuring current applicability would involve the regular and systematic review of a number of key features and key assumptions of the framework. These include the identification and perception of risks, which change over time as environments and attitudes change; the choices and implications of technologies, which also change rapidly and in potentially unpredictable ways (for example, improvements in the cost and capacity of storage was always expected and predictable, while the explosion in broadband and high speed mobile technologies was less so); and the whole compliance landscape, which is subject to perpetual changes with profound impacts as nations, groups of nations and industries update their compliance regulations in order to ensure greater protection for the subjects and users of data or for political, electoral or strategic reasons.

A recent initiative emerging out of the EU Cyber Security Strategy provides an illustration of significant current trends. This initiative has taken the form of a major survey of risk management methods and frameworks [117] in order “to identify and capture key reference documents for cybersecurity risk management, and to gather views of all EU nations and lead industries.” Essentially the organisers wish to provide the best possible guidance on risk management and information sharing, and within their structures have set up working groups tasked with reviewing existing risk management methods, frameworks and capability models, perform a gap analysis, and identify which models are likely to endure.

Although the survey is aimed at methods or standards that are or are intended to become national standards, the detailed questions that are asked in respect of each are revealing. These include:

- Could the method be considered as cross-sectorial? (i.e., could it be applied to different industry sectors?)
- Do you think it is technology neutral?
- In your opinion, is it easily applicable to SMEs?
- Does it include a link with organisations objectives or strategy?

- Does it include a way to identify critical information assets?
- Does it consider vulnerabilities?
- Does it include threat sources / actors?
- And likelihood of attack / compromise?
- Does it provide a ways for 'calculate' risk appetites / thresholds?
- Does it include a method / best practice for risk evaluation?
- And for risk reporting?
- Does it include a catalogue of mitigation mechanisms for risks?
- Does it provide alternatives for risk monitoring (audit, compliance, and governance)?
- Does it consider a supply chain perspective?
- Does it include the dependencies with external services and providers / customers?

I believe that this can in some way be viewed as a form of remote and partial validation of the framework presented in Chapter 10. Although clearly not aiming to occupy the same space as the national standards in terms of overall applicability or breadth of scope, it does correspond to many of the criteria anticipated by NISP.

Another recent publication of interest is the World Economic Forum's Global Information Technology Report 2014 [118], subtitled "Rewards and Risks of Big Data". This report "features a number of essays that inquire into the rewards and risks accruing from big data, an unprecedented phenomenon in terms of the volume, velocity, and variety of sources of the creation of new data. These essays also advise on the changes that organizations, both public and private, will need to adopt in order to manage, make sense of, and obtain economic and social value from this vast quantity of newly generated data" (p. v).

In the introduction, it is specifically noted that Big Data is such a new field that the implications of its uptake have yet to be determined, in spite of the amount of popular and industry discussion of its use. "This world of big data has also become a source of concern. The consequences of big data for issues of privacy and other areas of society are not yet fully understood" (p. ix).

The publication contains chapters on the move towards maturity of the technology, managing risks and rewards, and on building trust through regulation. It also demonstrates, with contributions from analytical bodies and industry players, how the consequences of the adoption of these technologies are starting to be recognized and that the aspects of risk management and proper controls are beginning to be addressed. Thus the choice of subject for this thesis, made in 2009, can be seen as anticipating future developments and pre-empting current industry debates.

11.4 Overall conclusion

In this thesis I have set out to identify the impacts of significant recent trends in technology on organisations that rely heavily on information systems and that have a philosophy and strategy of constantly seeking to improve the performance of their information systems environment. The three topics selected for analysis, unstructured data, big data, and the cloud, are in practice very often inextricably linked as organisations seek to maximise the competitive or organisational advantage that they can gain from the data that they own or to which they have access, at the same time seeking to reduce the operational costs of their information systems by outsourcing, either along traditional lines or by adopting cloud services.

Each of these tendencies brings with it a set of risks: strategic; operational; security-based; and compliance. In performing projects intending to take advantage of these tendencies and incorporate them into the operating environment, the management of organisations will need to identify and understand these risks, both separately and as an ensemble, and design policies and procedures that take account of those risks. As discussed in Chapter 2, there are a number of options for addressing risks: these can be mitigated, transferred, avoided or simply accepted. There can be no one-size-fits-all approach to addressing risk that can be rolled out to all organisations, because all organisations operate in different environments and have different risk appetites and profiles. The project planning framework presented in this thesis does not aim to provide a unique global response to every situation potentially encountered by an organisation that undertakes a major project of the type described; rather, it attempts to provide a framework that encourages managers to perform their own detailed analyses while being certain that they have considered all of the major risk groups and are following a structured approach.

The sheer quantity of literature and discussion reviewed in Chapters 3 and 4 bear witness to the importance that these tendencies are assuming in the modern business world. The technical and academic media have consecrated a great deal of time and space to the subjects; academic and business conferences dedicated to the topics are taking place regularly; and while this analysis did not consider the more popular end of the media spectrum, a cursory examination of the covers of the popular press demonstrates that the subject is attracting attention and provoking discussion at all levels.

Chapter 5 demonstrated how the questions were being addressed at a particular point in time. The results were in line with educated expectations: there are pressures to adopt new technologies and seize new opportunities if it appears that there are prospects of competitive advantage, greater efficiency or greater control over costs; at the same time, technical managers tend to be more reluctant to adopt early and are more suspicious of the claims made by service providers. In addition, many organisations were, at the time of asking at least, unable to identify the source, nature and location of much of the data they held. Clearly this would be a limiting factor in any move they made towards being able to leverage these data for operational purposes.

Chapters 6 to 9 described some of the constraints faced by organisations in respect of the security of data and demonstrable compliance with relevant restrictions. Emphasis was placed on the role of monitoring and audit because these are critical in being able to demonstrate the risks have been managed and that processes are under control. The complexities of cryptography were also described from a high-level and theoretical perspective: it is not unreasonable to assume that an effective data security framework will include the use of various cryptographic techniques, but choosing the most appropriate technique – or

techniques – and successfully implementing the chosen measures are non-trivial exercises. As emphasised throughout this thesis, in addition, the use of security technologies such as cryptography is not a guarantee in and of itself of effective security. Technologies are only one small part of a framework that needs to be robust, flexible and constantly reviewed and validated.

The title of this thesis emphasises the key words *control* and *compliance*. These principles underpin the whole structure of the model presented in Chapter 10 and of the analysis that precedes it. In practical terms the two principles are so closely related as to be inseparable, Control is the first objective of management at every level: the ability to know exactly what is happening within their units, functions and departments; to know that processes and procedures are operating completely, accurately and efficiently; and to have reasonable confidence that exceptions, breakdowns or anomalies will be identified and that sufficient remedial processes exist to respond to these incidents. Compliance very much follows on from this: if control is being exerted, in the right places and in a coherent and documented manner, then compliance with relevant requirements can be demonstrated. The process of acquiring certification of compliance within an organisation that already has consistent and documented internal control systems is very often essentially a process of fine-tuning, making adjustments to existing processes and practices to ensure that the details of the requirements are being respected.

Without effective control, management are handicapped in their efforts to manage their organisations efficiently and responsively. Very often the auditor will discover ad hoc control activities in place, aimed at managing what are considered to be the most critical and risky activities and processes, with a great deal of reliance on the abilities, experience, qualities and honesty of a number of key staff and managers. Depending on the size and nature of an organisation, this might be sustainable for a short period of time, but changes in the operating environment, operational activities, or in staffing, will typically expose weaknesses and allow errors, omissions and malicious actions to creep in.

Exerting effective control is difficult. As discussed in Chapter 2, management need to have a clear and complete notion of their overall control objectives, of the risks that might prevent those objectives from being achieved, and of the nature and number of the detailed internal control activities that they need to design, implement, operate and monitor in order to address those risks. Each step requires education, experience, perspective and guidance, not all of which are available to the desired degree within every organisation.

The framework presented in Chapter 10 is a framework aimed directly at permitting management to both exert control and, after verification of specific requirements, demonstrate compliance. It offers a step-by-step, top-down approach to risk management within a project, providing a level of comfort to management that the key risks have been identified and evaluated and that appropriate control activities have been implemented and performed. It cannot, for the reasons discussed above, provide a complete solution to every set of circumstances. It can, however, represent a useful tool that provides guidance and allows management to identify and consider the key issues confronting them.

As discussed in Chapter 10, the framework has been designed as an overview approach – a meta-framework in some respects - that can provide enough detail in certain circumstances for management to obtain the comfort that they feel they require. At the same time, other templates or control matrices can be introduced to provide greater detail where deemed necessary: the CSA matrix discussed in Section 10.16 provides valuable insights into a number of cloud-specific aspects of risk and control, while management

might seek detailed guidance on compliance issues from a standard-specific checklist or on technical issues from a vendor-supplied checklist. It would also be possible to further develop the framework to encompass industry-specific requirements such as Basel II and III in relation to the banking sector in general and the detailed requirements of FINMA with respect to the Swiss banking sector.

The use of this framework could be facilitated by the implementation of an interactive online version rather than reliance on a checklist open in a word processing or spreadsheet application. This could greatly facilitate the multi-user use of the framework, as well as allowing rapid and simple customisation of the content. It is possible to imagine an interactive version in which the user is guided through menus that allow the objectives of the project to be specified and parameters defined which then remove aspects that are not relevant to management's risks and control objectives, but also include additional content that corresponds to their specific requirements. The framework for a financial services company migrating a subset of its data onto a private cloud would thus look significantly different to the framework used by a small manufacturing company to identify and classify its data before consolidating them at a service provider, with the customising of the template having been semi-automated and reasonably straightforward. Updates and additional content would also be easy to distribute and apply.

In an operating environment in which cost-benefit analyses form an expected part of every decision-making process, strategic decisions in respect of information systems can be difficult to justify. A manufacture or buy choice can be supported by reasonably robust and verifiable numbers: assumptions can be questioned and the methods for arriving at particular number challenged, with calculations being re-run, but in general the analysis can be reduced to a collection of quantifiable assertions and suppositions. Strategic decisions in respect of information systems cannot necessarily be supported in the same way because of the difficulty of quantifying the impact of risks and of unpredictable future events. Management therefore need to rely on what qualitative methods they can find and try to ensure that risks are addressed in the most efficient and effective ways while taking decisions and driving projects towards completion.

References

1. British Museum - Ashurbanipal Library Phase 1 [Internet]. [cited 2014 Feb 1]. Available from: http://www.britishmuseum.org/research/research_projects/all_current_projects/ashurbanipal_library_phase_1.aspx
2. Hilbert M, López P. The world's technological capacity to store, communicate, and compute information. *Science*. 2011 Apr 1;332(6025):60–5.
3. COBIT 5: A Business Framework for the Governance and Management of Enterprise IT [Internet]. Available from: <http://www.isaca.org/COBIT/Pages/default.aspx>
4. Buneman P, Davidson S, Hillebrand G, Suciu D. A Query Language and Optimisation Techniques for Unstructured Data. Tech Rep MS-CIS 96-09. 1996;
5. AGICOA: The Rights People [Internet]. [agicoa.org](http://www.agicoa.org); [cited 2013 Jun 3]. Available from: <http://www.agicoa.org/english/about/about.html>
6. AGICOA: How it works [Internet]. [agicoa.org](http://www.agicoa.org); [cited 2013 Mar 6]. Available from: <http://www.agicoa.org/english/howitworks/howitworks.html>
7. Seiner RS. A Conceptual Meta-Model for Unstructured Data [Internet]. 2003 [cited 2014 Jan 28]. Available from: <http://www.tdan.com/view-articles/5087/>
8. Suciu D. Semistructured Data and XML. 5th Int Conf Found Data Organ FODO98 Kobe Jpn Novemb 12-13 1998. 1998;1–12.
9. Jhingran AD, Mattos N, Prahesh H. Information integration: A research agenda. *IBM Syst J*. 2002;41(4).
10. Chu E, Baid A, Chen T, Doan A, Naughton J. A Relational Approach to Incrementally Extracting and Querying Structure in Unstructured Data. 2007.
11. Buneman P. Semistructured Data. Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems. 1997. p. 117–21.
12. Franklin M, Halevy A, Maier D. From Databases to Dataspace: A New Abstraction for Information Management. *SIGMOD Rec*. 2005 Dec;34(4).
13. Raghuvver A, Jindai M, Mokbel M, Debnath B, Du D. Towards Efficient Search on Unstructured Data: An Intelligent-Storage Approach. Lisbon, Portugal; 2007.
14. Doan A, Naughton J, Ramakrishnan R, Baid A, Chai X, Chen F, et al. Information Extraction Challenges in Managing Unstructured Data. *ACM SIGMOD Rec*. 2008 Dec;37(4):14–20.
15. Maluf D, Tran P. Managing Unstructured Data With Structured Legacy Systems. *IEEEAC Follow Up*. 2007;
16. Sarawagi S. Information Extraction. *Found Trends Databases*. 2007;1(3):261–377.
17. Doan A, Naughton J, Baid A, Chai X, Chen F, Chen T, et al. The Case for a Structured Approach to Managing Unstructured Data. *CIDR Perspect*. 2009;
18. Kuechler W. Business Applications of Unstructured Text. *Commun ACM*. 2007 Oct;50(10).

19. Jowitt T. Unstructured Data Grows Unchecked, Study Says | PCWorld [Internet]. 2008 [cited 2014 Jan 28]. Available from: http://www.pcworld.com/article/153558/unstructured_data.html
20. Abadi D. Data Management in the Cloud: Limitations and Opportunities. *IEEE Comput Soc Bull Tech Comm Data Eng.* 2009 Mar;32(1).
21. Beyer K, Ercegovac V, Krishnamurthy R, Raghavan S, Rao J, Reiss F, et al. Towards a Scalable Enterprise Content Analytics Platform. *IEEE Comput Soc Bull Tech Comm Data Eng.* 2009 Mar;32(1).
22. Ceglowski M, Coburn A, Cuadrado J. An Automated Management Tool for Unstructured Data. *Proc IEEEWIC Int Conf Web Intell* 2003. 2003 Oct 13;554–7.
23. Lidwell W, Holden K, Butler J. *Universal Principles of Design.* Rockport Publishers; 2010.
24. Blumberg R, Atre S. The Problem with Unstructured Data. *DM Review.* 2003 Feb;
25. Tsai W, Wei X, Chen Y, Paul R, Chung J-Y, Zhang D. Data provenance in SOA: security, reliability and integrity. *SOCA.* 2007;l:223–47.
26. Groth P, Luck M, Moreau I. A protocol for recording provenance in service-oriented grids. *Proc of 8th international conference on principles of distributed systems.* 2004.
27. Mukherjee R, Mao J. Enterprise Search: Tough Stuff. *Queue.* 2004 Apr;37–46.
28. Duffy J. The tools and technologies needed for knowledge management. *Inf Manag J.* 2001 Jan;35(I):64.
29. Coy P. The Creative Economy. *BusinessWeek.* 2000 Aug 28;78–82.
30. Guynes C, Vanacek M. Critical success factors in data management. *Inf Manage.* 30 (1996):201–9.
31. Rao R. From Unstructured Data to Actionable Intelligence. *IT Pro.* 2003 Dec;29–35.
32. Baars H, Kemper H-G. Management Support with Structured and Unstructured Data – An Integrated Business Intelligence Framework. *Inf Syst Manag.* 2007;25:132–48.
33. Infosecurity - Infosecurity Europe 2010: Varonis explains need for unstructured data governance [Internet]. 2010 [cited 2014 Jan 28]. Available from: <http://www.infosecurity-magazine.com/view/9095/infosecurity-europe-2010-varonis-explains-need-for-unstructured-data-governance/>
34. Varonis Systems, Inc. *Managing Unstructured Data: 10 Key Requirements.* 2010.
35. Detlor B. Information management. *Int J Inf Manag.* 30 (2010):103–8.
36. Cheung CF, Lee WB, Wang WM, Wang Y, Yeung WM. A multi-faceted and automatic knowledge elicitation system (MAKES) for managing unstructured information. *Expert Syst Appl Int J.* 2011 May;38(5):5245–58.
37. Soibelman L, Wu J, Caldas C, Brilakis I, Lin K-Y. Management and analysis of unstructured construction data types. *Adv Eng Inform.* 22 (2008):15–27.
38. Baker S. *The Numerati.* Mariner Books; 2009.
39. Information Systems Audit and Control Association. *Big Data: Impacts and Benefits.* ISACA; 2013.
40. Manyika J, Chui M, Brown B, Bughin J, Dobbs R, Roxburgh C, et al. *Big data: the next frontier for innovation, competition and productivity.* McKinsey Global Institute; 2011.
41. Bollier D. *The Promise and Peril of Big Data.* Aspen Institute; 2010.
42. Bryant R, Katz RH, Lazowska ED. *Big-Data Computing: Creating Revolutionary Breakthroughs in Commerce, Science and Society* [Internet]. Computing Community Consortium; 2008. Available from: http://www.cra.org/ccc/docs/init/Big_Data.pdf

43. Skibiski G. New Deal on Data [Internet]. Sense Networks; 2009. Available from: http://www.sensenetworks.com/press/wef_globalit.pdf
44. Talmor E. Data-centric security [Internet]. 2010 [cited 2014 Jan 28]. Available from: <http://www.infosecisland.com/blogview/4249-Data-centric-security.html>
45. Kelley D. Addressing the Unstructured Data Protection Challenge. securitycurve.com; 2008.
46. Verizon Business. Security in the new information age: Striking a balance between risk and opportunity. 2010.
47. Chickowski E. Eight Steps to Securing Unstructured Data - Security news from Channel Insider [Internet]. 2009 [cited 2014 Jan 28]. Available from: <http://www.channelinsider.com/c/a/Security/Eight-Steps-to-Securing-Unstructured-Data-851366/>
48. Brink D. Securing Unstructured Data: How Best-in-Class Companies Manage to Serve and Protect. Aberdeen Group; 2009 Jun.
49. McAdams J. Top Secret: Securing Data with Classification Schemes - Computerworld [Internet]. 2006 [cited 2014 Jan 28]. Available from: http://www.computerworld.com/s/article/110510/Top_Secret_Securing_Data_with_Classification_Schemes?taxonomyId=017
50. Lorenz G. Building an unstructured data protection program [Internet]. 2010 [cited 2014 Jan 28]. Available from: <http://searchfinancialsecurity.techtarget.com/tip/Building-an-unstructured-data-protection-program>
51. aprigo.com. Preventing Data Breaches. 2010.
52. Unstructured data creating security hole - Security - Technology - News - iTnews.com.au [Internet]. 2008 [cited 2014 Jan 28]. Available from: <http://www.itnews.com.au/News/115825,unstructured-data-creating-security-hole.aspx>
53. Imperva. Five Signs Your File Data is at Risk. 2010.
54. Ely A. Beyond the Database: Protecting Unstructured Data [Internet]. 2010 Jun. Available from: <http://reports.informationweek.com/abstract/14/3477/Regulatory-Compliance/undefined>
55. Craig T, Ludloff M. Privacy and Big Data. O'Reilly; 2011.
56. Newton N, Minow F. Government Data Mining. McGraw Hill Handbook of Homeland Security. 2008.
57. Information Systems Audit and Control Association. Privacy & Big Data. ISACA; 2013.
58. Personal Data: The Emergence of a New Asset Class [Internet]. World Economic Forum; 2011. Available from: www.weforum.org/reports/personal-data-emergence-new-asset-class
59. Unlocking the Value of Personal Data: From Collection to Usage [Internet]. World Economic Forum; 2013. Available from: www3.weforum.org/docs/WEF_IT_UnlockingValuePersonalData_CollectionUsage_Report_2013.pdf
60. Lohr S. Big Data Is Opening Doors, but Maybe Too Many. New York Times [Internet]. 2013 Mar 23; Available from: www.nytimes.com/2013/03/24/technology/big-data-and-a-renewed-debate-over-privacy.html
61. Data-Driven Innovation. A Guide for Policymakers: Understanding and Enabling the Economic and Social Value of Data [Internet]. Software & Industry Information Association; 2013. Available from: www.siiia.net
62. Adolph M. ITU -T Technology Watch Report 9, Distributed Computing: Utilities, Grids & Clouds. International Telecommunications Union; 2009.
63. 2. ITU Telecommunication Standardization Bureau, Policy and Technology Watch Division. Activities in Cloud Computing Standardization. International Telecommunications Union; 2010.
64. Gantz J, Reinsel D. Extracting Value from Chaos. IDC iView; 2011.
65. Bertino E, Paci F, Ferrini R, Shang N. Privacy-preserving Digital Identity Management for Cloud Computing. IEEE Comput Soc Bull Tech Comm Data Eng. 2009 Mar;32(1).

66. Cooper B, Baldeschwieler E, Fonseca R, Kistler J, Narayan P, Neerdaels C, et al. Building a Cloud for Yahoo! IEEE Comput Soc Bull Tech Comm Data Eng. 2009 Mar;32(1).
67. Mather T, Kumaramswamy S, Latif S. Cloud Security and Privacy: An Enterprise Perspective on Risks and Compliance. O'Reilly Media Inc.; 2009.
68. Information Systems Audit and Control Association. Cloud Governance: Questions Boards of Directors Need to Ask. ISACA; 2013.
69. Staten J. What are Enterprises Really Doing in the Cloud? [Internet]. The Forrester Blog; 2011 [cited 2013 Nov 5]. Available from: http://blogs.forrester.com/james_staten/11-10-25-what_are_enterprises_really_doing_in_the_cloud
70. Security Considerations for Cloud Computing. ISACA; 2012.
71. Mell P, Grance T. The NIST Definition of Cloud Computing, US National Institute of Standards and Technology (NIST) Special Publication (SP) 800-145. US National Institute of Standards and Technology (NIST); 2011.
72. Krutz R, Vines D. Cloud Security: A Comprehensive Guide to Secure Cloud Computing. Wiley Publishing Inc.; 2010.
73. Schneier B. Cloud Computing [Internet]. Schneier on Security. 2009. Available from: https://www.schneier.com/blog/archives/2009/06/cloud_computing.html
74. Babcock C. HP On The Cloud: 'The World Is Cleaving In Two'. Information Week [Internet]. 2009 Feb 10; Available from: <http://www.informationweek.com/hp-on-the-cloud-the-world-is-cleaving-in-two/d/d-id/1076503?>
75. Stoneburner G, Hayden C, Feringa A. Engineering Principles for Information Technology Security (A Baseline for Achieving Security), Revision A; NIST Special Publication 800-27 Rev A [Internet]. NIST (National Institute of Standards and Technology); 2004. Available from: <http://csrc.nist.gov/publications/nistpubs/800-27A/SP800-27-RevA.pdf>
76. ISACA, Cloud Security Alliance. Cloud Computing Market Maturity Study Results 2012. Cloud Security Alliance, ISACA; 2012.
77. IT Governance Institute. Board Briefing on IT Governance. 2001.
78. Van Grembergen W. Introduction to the Minitrack IT Governance and Its Mechanisms. Proceedings of the 35th Hawaii International Conference on System Sciences (HICSS). 2002.
79. EUR-Lex - 31995L0046 - EN [Internet]. [cited 2013 Oct 24]. Available from: <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:EN:HTML>
80. Export.gov - U.S.-EU Safe Harbor Homepage [Internet]. [cited 2013 Oct 24]. Available from: http://export.gov/safeharbor/eu/eg_main_018365.asp
81. Safe Swiss Cloud | Real enterprise cloud with privacy [Internet]. [cited 2013 Oct 24]. Available from: <https://www.safeswisscloud.ch/en/node>
82. European Telecommunications Standards Institute. Information Security Indicators (ISI); Event Model; A security event classification model and taxonomy. ETSI; 2013.
83. Ghernaouti S. Cyberpower: Crime, Conflict and Security in Cyberspace. EPFL Press; 2013.
84. UK Cabinet Office. ITIL Home [Internet]. [cited 2013 Oct 15]. Available from: <http://www.itil-officialsite.com/home/home.aspx>
85. ISO/IEC. ISO/IEC International Standard 27000:2012 - Information technology — Security techniques — Information security management systems — Overview and vocabulary. 2012.
86. COSO. Welcome to COSO [Internet]. [cited 2013 Oct 16]. Available from: <http://www.coso.org/>
87. COSO. About Us [Internet]. [cited 2014 Mar 18]. Available from: <http://www.coso.org/aboutus.htm>
88. McNally JS. The 2013 COSO Framework & SOX Compliance: One Approach to an Effective Transition. COSO; 2013.

89. Shaw H. The Trouble with COSO. CFO Magazine [Internet]. 2006 Mar 15; Available from: <http://ww2.cfo.com/accounting-tax/2006/03/the-trouble-with-coso/view-all/>
90. COSO. Guidance on Monitoring Internal Control Systems. AICPA; 2009.
91. Racz N, Weippl E, Seifert A. A frame of reference for research of integrated GRC. In: de Decker B, Schaumüller-Bichl I, editors. Communications and Multimedia Security, 11th IFIP TC 6/TC 11 International Conference, CMS 2010 Proceedings. Berlin: Springer; p. 106–17.
92. IT GRC Software Solutions - MetricStream [Internet]. [cited 2013 Oct 28]. Available from: http://www.metricstream.com/solutions/it_grc.htm
93. ISO/IEC. ISO/IEC 38500: 2008, Corporate governance of information technology. ISO/IEC; 2008.
94. How the SEC Protects Investors, Maintains Market Integrity, and Facilitates Capital Formation (Securities and Exchange Commission) [Internet]. [cited 2013 Oct 28]. Available from: <http://www.sec.gov/about/whatwedo.shtml>
95. Auditing Standard No. 2 [Internet]. [cited 2013 Oct 28]. Available from: http://pcaobus.org/standards/auditing/pages/auditing_standard_2.aspx
96. ITGI. IT Control Objectives for Sarbanes-Oxley: The Role of IT in the Design and Implementation of Internal Control Over Financial Reporting. ITGI; 2006.
97. Official Source of PCI DSS Data Security Standards Documents and Payment Card Compliance Guidelines [Internet]. [cited 2013 Oct 28]. Available from: https://www.pcisecuritystandards.org/security_standards/index.php
98. admin.ch - Droit fédéral [Internet]. [cited 2013 Oct 29]. Available from: <http://www.admin.ch/opc/fr/classified-compilation/19920153/index.html>
99. admin.ch - Droit fédéral [Internet]. [cited 2013 Oct 29]. Available from: <http://www.admin.ch/opc/fr/classified-compilation/19930159/index.html>
100. admin.ch - Droit fédéral [Internet]. [cited 2013 Oct 29]. Available from: <http://www.admin.ch/opc/fr/classified-compilation/19110009/index.html>
101. admin.ch - Droit fédéral [Internet]. [cited 2013 Oct 29]. Available from: <http://www.admin.ch/opc/fr/classified-compilation/20071826/index.html>
102. PFPDT - Certification de produits/systèmes en matière de protection des données [Internet]. [cited 2013 Oct 29]. Available from: <http://www.edoeb.admin.ch/dokumentation/00153/00262/00265/index.html?lang=fr>
103. Préposé fédéral à la protection des données et à la transparence. Etat de la protection des données dans le monde. Confédération suisse; 2013.
104. Préposé fédéral à la protection des données et à la transparence. Explications concernant l'informatique en nuage (cloud computing). Confédération suisse; 2011.
105. Comité de la Chambre fiduciaire. Norme d'audit suisse: Vérification de l'existence du système de contrôle interne (NAS 890). 2007.
106. Rijmen V, De Cock D, Smart N. Recommended cryptographic measures: Securing personal data. European Union Agency for Network and Information Security (ENISA); 2013.
107. Gürses FS. Multilateral privacy requirements analysis in online social network services. KU Leuven; 2010.
108. Leistikow R, Tavangarian D. Secure Picture Data Partitioning for Cloud Computing Services. 27th International Conference on Advanced Information Networking and Applications Workshops (WAINA), 2013. 2013.
109. 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 [Internet]. European Union, Official Journal L 281; 1995. Available from: <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2013:173:0002:0008:en:PDF>

110. Proposal for a regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation), COM (2012) [Internet]. European Commission; 2012. Available from: http://ec.europa.eu/justice/data-protection/document/review2012/com_2012_11_en.pdf
111. Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications) [Internet]. European Union, Official Journal L 201; 2002. Available from: <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32002L0058:en:HTML>
112. Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009 amending Directive 2002/22/EC on universal service and users' rights relating to electronic communications networks and services. European Union; 2009.
113. Commission Regulation (EU) No 611/2013 of 24 June 2013 on the measures applicable to the notification of personal data breaches under Directive 2002/58/EC of the European Parliament and of the Council on privacy and electronic communications [Internet]. European Union; 2013. Available from: <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2013:173:0002:0008:en:PDF>
114. Ghernaoui S. Sécurité informatique et réseaux. 4th edition. Dunod; 2013.
115. CSA. Cloud Security Alliance Membership [Internet]. [cited 2014 May 5]. Available from: <https://cloudsecurityalliance.org/membership/>
116. CSA. Cloud Controls Matrix v3 [Internet]. CSA; [cited 2014 Apr 24]. Available from: <https://cloudsecurityalliance.org/download/cloud-controls-matrix-v3/>
117. EU Cyber Security Strategy. Survey of Risk Management Methods, Frameworks and Capability Maturity Models for the EU Network Information Security Platform. ENISA; 2014.
118. Bilbao-Osorio B, Dutta S, Lanvin B. World Economic Forum - The Global Information Technology Report 2014 - Rewards and Risks of Big Data [Internet]. World Economic Forum; 2014 [cited 2014 Apr 28]. Available from: http://www3.weforum.org/docs/WEF_GlobalInformationTechnology_Report_2014.pdf

Glossary

Accuracy - one of the four information processing objectives, ensuring that every transaction or modification is processed correctly and accurately.

Assurance - the process of obtaining and using accurate and current information about the efficiency and effectiveness of policies and operations, and the status of compliance with the statutory obligations, in order for management to control an organization's activities.

Asset - in accounting and auditing terms, an economic resource. Anything tangible or intangible that can be owned or controlled to generate value.

Audit - a planned and documented activity intended to determine adequacy and compliance with established procedures or standards.

Audit comfort – the support needed to draw conclusions provided by audit procedures.

Audit evidence - the documentary and other elements obtained to support the assertions of management and staff and to demonstrate the operation of processes and controls or the existence of transactions or other activities.

Audit procedures – the various tests of details, analytical procedures, confirmations and other activities performed to gain audit evidence.

Authentication - the act of confirming a claim attesting to identity.

Authorisation - the function of specifying access rights to resources.

Availability - one of the three main information security objectives, allowing data and resources to be accessible to and useable by appropriate users.

Big data – the accumulation of datasets from different sources and of different types that can be exploited to yield insights. ISACA provides a definition as datasets that are too large or too fast-changing to be analysed using traditional relational or multidimensional database techniques or commonly used software tools to capture, manage and process the data in a reasonable elapsed time

Cloud computing – a computing paradigm in which highly scalable computing resources, often configured as a distributed system, are provided as a service through a network.

Compensating control - a control activity designed to reduce the risk that an existing or potential control weakness will result in a failure to meet a control objective.

Completeness - one of the four information processing objectives, ensuring that all relevant transactions or modifications are processed once and once only.

Compliance - the process or state of conforming to a rule, such as a specification, standard, policy or law.

Confidentiality - one of the three main information security objectives, ensuring that data and information remain private and that only authorised users have access.

Control - the management function of ensuring that processes are operating appropriately and that errors are identified and corrected.

Control objectives - a statement of the desired results or purposes to be achieved by implementing control procedures in and around a particular activity.

Corrective control - a control activity established to remedy control or processing problems

Design effectiveness – determining that internal control procedures are appropriate in respect of the control objective they are intended to achieve and the risk they are intended to address.

Detective control - a control activity designed to identify errors or inconsistencies in input, processing or operations.

Encryption - the process of encoding messages or information in such a way that only authorized parties can read the contents.

Identification - the act or process of stating or otherwise indicating a claim purportedly attesting to a person or thing's identity.

Integrity - one of the three main information security objectives, ensuring that data and information remain complete and free from tampering.

Internal controls – the activities performed to ensure that policies and procedures are implemented and operated consistently and effectively, allowing errors and omissions to be prevented or detected.

Key control - in a framework of internal controls, a control activity of particular importance, the failure of which would have a significant impact on the quality of the overall control environment.

Operating effectiveness – that internal control procedures are functioning as designed, regularly and consistently.

Preventive control - an internal control activity designed to prevent an error, omission or other unwanted event from taking place.

Restricted access - one of the four information processing objectives, ensuring that systems and data are protected from unauthorised modification or deletion.

Risk - the potential of losing something of value, weighed against the potential to gain something of value.

Risk analysis - the evaluation and measurement of risks in order to be able to prioritise them and scope appropriate responses.

Security - the degree of resistance to, or protection from, harm, applicable to any vulnerable and valuable asset.

Security triad – the three key information security objectives of Confidentiality, Integrity and Availability.

Service auditor - an auditor who reports on controls of a service organization that may be relevant to a user organization's internal control, frequently as it relates to an audit of financial statements.

Third party assurance - the assurance obtained over the operations and controls in place at a service provider through an audit performed by an independent, third party auditor.

Threat - an action, device or event that threatens the integrity and security of data and information stored within, and communicated between, information systems.

Threat agent - an individual or group that can manifest a threat.

Unstructured data – data sitting outside structured and organized data repositories such as relational databases.

Validity - one of the four information processing objectives, ensuring that only genuine, appropriate and approved transactions and modifications are recorded in information systems.

Vulnerability - a flaw or weakness in security or in controls that could be exploited to gain access to, modify, copy or delete information.

List of Appendices

Appendix A - Additional Illustrations

Appendix B - Survey Statistics

Appendix C - Book Chapter

Appendix D - Published Articles

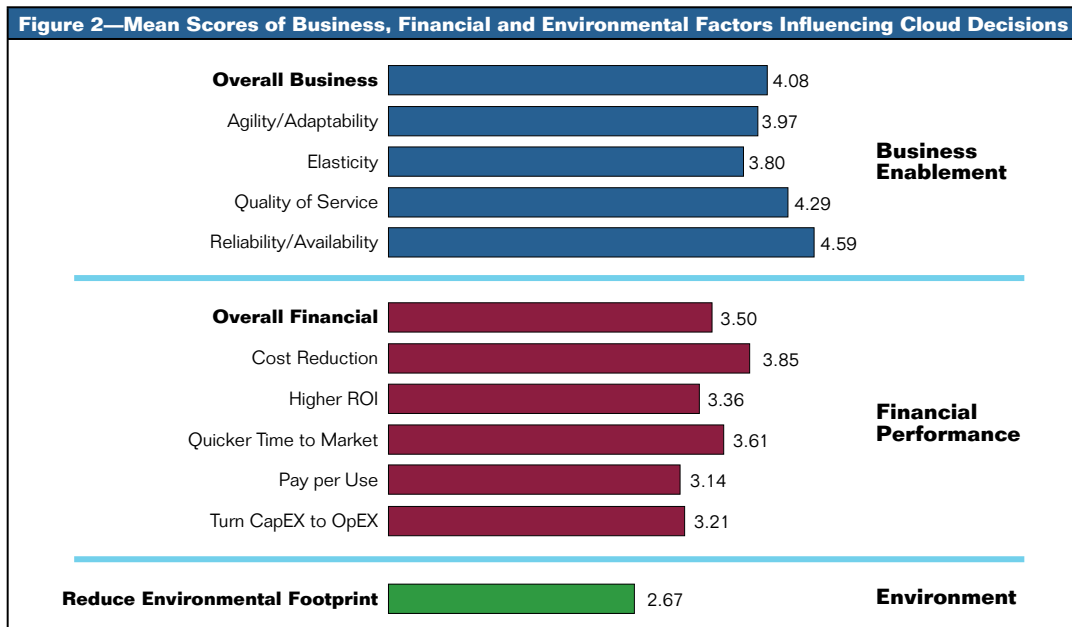
Appendix A - Additional Illustrations

This appendix contains illustrations that provide additional information and background to elements discussed in Chapter 5 and Chapter 7. These illustrations are drawn from the source publications and if necessary modified for presentation here.

1 Figures and Tables from the ISACA Survey (Chapter 5)

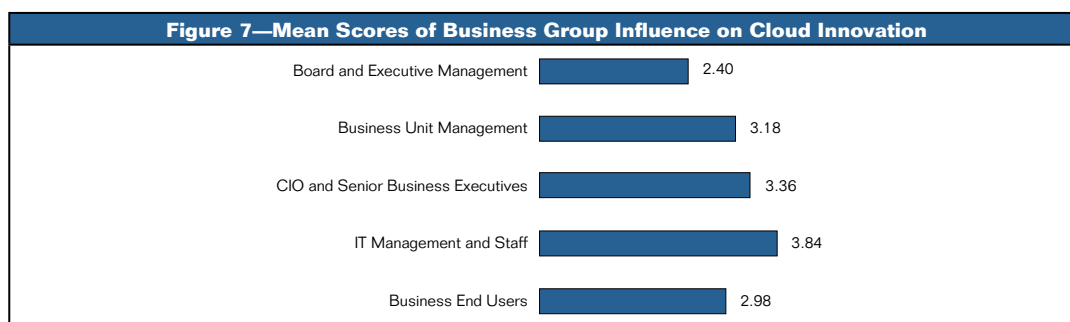
These figures provide the details of the survey results that are discussed in the second part of Chapter 5.

1.1 Mean Scores of Business, Financial and Environmental Factors Influencing Cloud Decisions – from p. 11.



“Responses were recorded on a range of 0 to 5, where 0 indicates that the factor has no influence. A response of 1 indicates a minimal level of influence while 5 indicates the highest level of influence in making the cloud decision” (p. 10).

1.2 Mean Scores of Business Group Influence on Cloud Innovation – from p. 14



1.3 Positive and Negative Influences on Cloud Adoption and Innovation – from p. 17

Table 4—Positive and Negative Influences on Cloud Adoption and Innovation

Positive Influence on Cloud Adoption/Innovation	Mean Score	Rank	Negative Influence on Cloud Adoption/Innovation	Mean Score	Rank
Cost management	3.77	1	Information security	4.22	1
Agility	3.75	2	Data ownership/custodian responsibilities	4.12	2
Time to market	3.73	3	Legal and contractual issues	4.04	3
Efficiency	3.65	4	Regulatory compliance	4.01	4
Productivity	3.61	5	Information assurance	3.77	5
Business unit demand	3.55	6	Longevity of suppliers	3.44	6
Resilience	3.52	7	Contract lock-in	3.42	7
New technology	3.46	8	Performance standards	3.30	8
Customer demand	3.42	9	Disaster recovery/business continuity	3.25	9
Technical resources	3.37	10	Performance monitoring	3.21	10
New markets	3.33	11	Technology stability	3.10	11
Summary Mean	3.56		Summary Mean	3.62	

1.4 User and Provider Perspectives on Negative Influences on Cloud Adoption and Innovation – from p. 18

Table 6—User and Provider Perspectives on Negative Influences on Cloud Adoption and Innovation

Negative Influence on Cloud Adoption/Innovation	User Perspective		Provider Perspective	
	Mean Score	Rank	Mean Score	Rank
Information security	4.23	1	4.21	1
Data ownership/custodian responsibilities	4.18	2	4.09	2
Regulatory compliance	4.05	3	3.99	4
Legal and contractual issues	4.04	4	4.05	3
Information assurance	3.87	5	3.73	5
Contract lock-in	3.50	6	3.39	7
Longevity of suppliers	3.50	7	3.41	6
Disaster recovery/business continuity	3.47	8	3.16	10
Performance standards	3.40	9	3.25	8
Performance monitoring	3.32	10	3.16	9
Technology stability	3.14	11	3.09	11
Summary Means	3.76		3.51	

1.5 Risk Components in the Confidence Barometer – from p. 21

Table 11—Risk Components in the Confidence Barometer	
Risk Component	Mean Score
Technical risk relative to cloud computing is considered a part of technical risk management.	2.71
Business risk relative to cloud computing is considered as part of enterprise risk management.	2.30
Technical risk relative to cloud computing is considered as part of business unit risk management.	2.25
Board and executive management make business decisions concerning cloud computing based on a realistic understanding of risk.	1.65
Overall	2.22

1.6 Role Definition Components in the Confidence Barometer – from p. 22

Table 12—Role Definition Components in the Confidence Barometer	
Role Definition Component	Mean Score
IT staff roles and responsibilities are clearly defined and coordinated.	2.28
Supplier and user roles and responsibilities are clearly defined and coordinated.	2.12
Roles and responsibilities between business units and IT are clearly defined and coordinated.	2.08
Overall	2.16

1.7 Requirements Components in the Confidence Barometer – from p. 22

Table 13—Requirements Components in the Confidence Barometer	
Requirements Component	Mean Score
Service level requirements are clearly defined.	2.23
Business and technical requirements are clearly defined and incorporated into contracts.	2.02
Overall	2.12

1.8 Performance Monitoring Components in the Confidence Barometer – from p. 22

Table 14—Performance Monitoring Components in the Confidence Barometer	
Performance Monitoring Component	Mean Score
Service providers effectively monitor and report against service level requirements.	2.02
Users monitor provider performance against service level agreements.	1.88
Overall	1.95

1.9 Problem Resolution Components in the Confidence Barometer – from p. 26

Table 23—Problem Resolution Components in the Confidence Barometer		
Problem Resolution Component	Mean Score	Rank
Continuity/availability	2.89	1
Performance	2.53	2
Problem management	2.42	3
User-supplier relationship	2.35	4
Solution integration	2.31	5
Security/assurance	2.25	6
Contracts	2.07	7
Regulation and legislation	2.06	8
Overall	2.37	

1.10 Continuity and Availability Components in the Confidence Barometer – from p. 26

Table 24—Continuity and Availability Components in the Confidence Barometer	
Continuity and Availability Component	Mean Score
Availability	3.06
Disaster recovery/business continuity	2.84
System outages	2.77
Overall	2.89

1.11 Performance Components in the Confidence Barometer – from p. 27

Table 25—Performance Components in the Confidence Barometer	
Performance Component	Mean Score
Service performance	2.84
Performance standards	2.58
Monitoring performance	2.55
Performance metrics	2.41
Performance assurance	2.40
Availability of expert support from suppliers	2.37
Overall	2.53

1.12 Security and Assurance Components in the Confidence Barometer – from p. 29

Table 32—Security and Assurance Components in the Confidence Barometer	
Security and Assurance Component	Mean Score
Concerns for multitenancy	2.42
Information security	2.41
Testing and assurance	2.31
Data ownership/custodian responsibilities	2.18
International data privacy	1.90
Overall	2.25

1.13 User and Provider Perspectives on Security and Assurance Components – from p. 30

Table 33—User and Provider Perspectives on Security and Assurance Components			
Security and Assurance Component	Overall Rank	User Rank	Provider Rank
Concerns for multitenancy	10	12	11
Information security	12	7	17
Testing and assurance	18	18	18
Data ownership/custodian responsibilities	22	22	23
International data privacy	25	25	26

1.14 Confidence Perspectives of Different Disciplines on Security and Assurance Components – from p. 30

Security and Assurance Component	Overall Rank	Business Rank	Security Rank	Technology Rank
Concerns for multitenancy	10	13	15	8
Information security	12	11	14	13
Testing and assurance	18	22	10	21
Data ownership/custodian responsibilities	22	23	23	19
International data privacy	25	25	27	25

1.15 Regulation and Legislation Components in the Confidence Barometer – from p. 31

Regulation and Legislation Component	Mean Score
Regulatory compliance	2.34
Government regulations keeping pace with the market	1.80
Overall	2.06

2 Tables of Incidents and Vulnerabilities from the ETSI Model (Chapter 7)

These tables summarise the classifications and descriptions of incidents and vulnerabilities described in the ETSI Model and referred to in Chapter 7.

2.1 Intrusions and external attacks (Category IEX)

	Type of attack	Means
1	Unauthorized access to organization's information system or inter-professional network.	Both human motivation and use of technical means.
2	Various unauthorized actions against organization's servers or applications	Both human motivation and use of technical means
3	External installation of unauthorized software programs on organization's workstations or servers	Primarily use of technical means
4	Organization's information system remote disturbance	Primarily use of technical means
5	Social engineering attacks	Primarily use of social engineering
6	Verbal act of stress	External attacker without physical action
7	Physical intrusion or action	Personal and physical intervention of attacker on an organization's technical capabilities
8	Illicit activity carried out on the public Internet, which harms an organization	External and almost out of the organization's control

2.2 Malfunctions (Category IMF)

	Type of attack	Means
1	Unavailability caused by natural disaster	Cause external to organization
2	Physical failure	Accidental origin
3	Unavailability due to environmental failure	Origin stemming from a unforeseeable environmental disturbance or breakdown
4	Malfunction due to abnormal activity	Due to unforeseen deviation from normal usage patterns
5	Accidental destruction or removal of sensitive data or equipment	Due to human weaknesses
6	Accidental modification of sensitive data	Due to careless computer usage
7	Accidental leakage of protected data	Due to human error or carelessness at software development and configuration level
8	Programming error leading to system malfunction	Due to human errors or software development mistakes
9	Configuration error leading to system malfunction	Due to human errors or mistakes in system configuration

2.3 Deviant internal behaviours (Category IDB)

	Type of attack	Means
1	Usurpation of identity (or user impersonation)	Malicious action, curiosity, or ease of use
2	Usurpation or abuse of rights (or privileges)	Wide range of cases, from strong motivation using technology to moderate information thirst-related motivation to very moderate motivation due to the sense of impunity and of all-mightiness (or to carelessness and ease of use)
3	Other sheer malicious behaviours	Generally strong motivation to enrich oneself or to harm one's organization
4	Organization disturbance	Related to the general workplace atmosphere
5	Personal attack on organization's personnel	Loss of self-control with aggressiveness level that can be extreme, or very motivated means of pressure to obtain an advantage at all costs
6	Reckless or careless action or behaviour	Deviant behaviours that can lead to critical risks

2.4 Behavioural vulnerabilities (Category VBH)

	Type of attack	Means
1	Illicit or dangerous protocols used	Great variety of situations involving low or high level communication methods
2	Internet illicitly accessed	Various dangerous methods used to prevent user tracking and escape controls
3	File illicitly transferred between the organization and the outside world	Often related to confusion between professional and private worlds
4	Workstation used without complying the required security tools, configurations and rules	Great variety of deviant practices jeopardizing workstations (widely recognized as the weakest link) and subsequently possibly the entire information system
5	Password illicitly handled or managed	Vulnerabilities remaining amongst the most frequent causes of security incidents
6	Authentication illicitly handled or managed	Can lead to very serious situations given the trust placed in such measures
7	Access rights illicitly granted	Issue to be monitored closely to avoid too many drifts damaging trust in organization overall security
8	Central systems or applications irrelevantly handled	Great variety of deviant practices concerning either administrators or plain users
9	Weakness exploited through social engineering methods	Human weaknesses exploited to make possible security incidents underway
10	Miscellaneous	All other vulnerabilities that cannot be categorized above

2.5 Software vulnerabilities (Category VSW)

The content of this category is extremely diverse. No classification is proposed for this category within the ETSI model, as such a classification would only be replicating the well-established and widely-accepted CWE identification scheme (detailed in NIST SP 800-126 Revision 2).

Annual surveys (such as those performed by the SANS Institute) generally list the most significant weaknesses, categorizing them into high-level categories such as:

- Insecure interaction between components;
- Risky resource management;
- Porous defences.

2.6 Configuration vulnerabilities (Category VCF)

The contents of this category are extremely diverse. No classification is proposed for this category within the ETSI model, as such a scheme would be replicating the CCE (Common Configuration Enumeration) list (detailed in MITRE CCE Version 5).

Standard secure configurations are available for major operating systems (Unix-like and Windows), browsers (IE, Chrome), Internet servers (WebLogic), applications (Office) or cross-platform best practices. Based on these extensive lists, it is a straightforward task to derive configuration vulnerabilities that can be also considered as non-conformities against the best current practices. Configuration vulnerabilities related to fixed or wireless network security should be added to these lists.

2.7 General security vulnerabilities (technical & organizational) (Category VTC and Category VOR)

	Type of attack	Means
1	Security organization	Is the responsibility of the top security positions
2	Governance	Make sure security governance is effective and, if possible, efficient
3	Development and testing	Make sure that the initial steps in projects address IT security relevantly
4	Security operations	Directly related to continuous auditing
5	Incident detection and management process	At the core of cyberdefence and SIEM approaches
6	Vulnerability or weakness management process	Part of cyberdefence and SIEM approaches
7	Auditing process	Make sure auditing process is running appropriately
8	Miscellaneous	All other vulnerabilities that cannot be categorized by one of the above

Appendix B - Survey Statistics

Results of the survey – raw data in tabular form

Respondent number	1	2	3	4	5	6 (5A)	7	8 (7A)	9 (7B)	10	11
Size (1-50, 51-100, 101-500, 500+)	500 500	51 51	500 500	500 500	1 1	1	500 500	500	500	51 51	51 51
Ownership (Pub, Pri, Oth)	Pub Pub	Pri Pri	Pri Pri	Oth Oth	Oth Oth	Oth	Pub Pub	Pub	Pub	Pri Pri	Oth Oth
Section A											
Q1	Y	N	Y	Y	Y	Y	Y	Y	Y	Y	Y
Q2	Y	N	Y	N	Y	Y	Y	Y	N	N	N
Q3	Y	N	Y	N	Y	N	Y	Y	Y	N	N
Q4	Y	N	N	N	N	N	Y	N	N	N	N
Q5	Y	N	N	N	N	N	Y	N	N	N	N
Q5A	Y						Y				
Q6	Y	N	Y	N	Y	N	Y	Y	Y	N	N
Q6A	Y		N		N		Y	N	N		
Q7	Y	N	N	N	Y	N	Y	Y	Y	N	N
Q7A	Y				N		Y	N	Y		
Q8	Y	N	N	N	N	N	N	N	N	N	N
Section B											
Q9	Y	Y	Y	Y	N	N	Y	Y	Y	Y	Y
Q10	N	N	N	N			N	N	N	N	N
Q11	N	N	N	N			N	N	N	N	N

Appendix B

Respondent number	12	13	14	15	16 (15A)	17 (15B)	18	19 (18A)	20	21 (20A)	22
Size (1-50, 51-100, 101-500, 500+)	1 1	1 1	101 101	500 500	500	500	500 500	500	500 500	500	1 1
Ownership (Pub, Pri, Oth)	Oth Oth	Pri Pri	Pri Pri	Pub Pub	Pub	Pub	Oth Oth	Oth	Pub Pub	Pub	Pri Pri
Section A											
Q1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	N
Q2	Y	N	Y	Y	Y	Y	N	Y	Y	Y	N
Q3	N	N	N	Y	Y	Y	N	Y	Y	Y	N
Q4	N	N	N	Y	Y	Y	N	N	Y	Y	N
Q5	N	N	N	Y	Y	Y	N	N	Y	N	N
Q5A				Y	Y	Y			Y		
Q6	N	N	N	Y	Y	Y	Y	N	Y	Y	N
Q6A				Y	N	N	N		Y	Y	
Q7	N	N	N	Y	Y	Y	N	N	Y	Y	N
Q7A				Y	Y	Y			Y	Y	
Q8	N	N	N	Y	N	N	N	N	Y	Y	N
Section B											
Q9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Q10	N	N	N	Y	Y	Y	N	N	Y	Y	N
Q11	N	N	N	Y	Y	Y	Y	Y	Y	Y	N

Appendix B

Respondent number	23	24	25	26	27	28	29	30	31	32	33
Size (1-50, 51-100, 101-500, 500+)	1	51	51	101	101	101	500	500	1	1	51
	1	51	51	101	101	101	500	500	1	1	51
Ownership (Pub, Pri, Oth)	Pri Pri	Pri Pri	Pri Pri	Oth Oth	Oth Oth	Pri Pri	Pri Pri	Pri Pri	Pri Pri	Oth Oth	Oth Oth
Section A											
Q1	N	N	N	Y	Y	Y	Y	Y	N	Y	Y
Q2	N	N	N	Y	Y	Y	Y	Y	N	N	N
Q3	N	N	N	N	Y	Y	Y	Y	N	N	N
Q4	N	N	N	N	N	N	Y	Y	N	N	N
Q5	N	N	N	N	N	N	N	N	N	N	N
Q5A											
Q6	N	N	N	N	N	N	Y	Y	N	N	N
Q6A							N	N			
Q7	N	N	N	N	N	N	Y	Y	N	N	N
Q7A							N	N			
Q8	N	N	N	N	N	N	N	N	N	N	N
Section B											
Q9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Q10	N	N	N	N	N	N	N	N	N	N	N
Q11	N	N	N	N	N	N	N	N	N	N	N

Appendix B

Respondent number	34	35	36	37	38	39	40	41	TOTALS					
Size (1-50, 51-100, 101-500, 500+)	51	500	500	51	51	51	101	1	8	5				
	51	500	500	51	51	51	101	1	10	11				
Ownership (Pub, Pri, Oth)	Pri	Oth	Pub	Pri	Pri	Oth	Pri	Pri	Pub	5				
	Pri	Oth	Pub	Pri	Pri	Oth	Pri	Pri	Pri	18				
									Oth	11				
Section A														
											Yes	No	Yes %	No %
Q1	Y	Y	Y	Y	N	N	N	Y	32	9	78%	22%		
Q2	N	Y	Y	Y	N	N	N	N	22	19	54%	46%		
Q3	N	N	Y	N	N	N	N	N	17	24	41%	59%		
Q4	N	N	N	N	N	N	N	N	9	32	22%	78%		
Q5	N	N	N	N	N	N	N	N	6	35	15%	85%		
Q5A									6	0	100%	0%		
Q6	N	N	N	N	N	N	N	N	14	27	34%	66%		
Q6A									5	9	36%	64%		
Q7	N	N	N	N	N	N	N	N	12	29	29%	71%		
Q7A									8	4	67%	33%		
Q8	N	N	N	N	N	N	N	N	4	37	10%	90%		
Section B														
Q9	Y	Y	Y	Y	Y	Y	Y	Y	39	2	95%	5%		
Q10	N	N	Y	N	N	N	N	N	6	33	15%	85%		
Q11	N	N	N	N	N	N	N	N	7	32	18%	82%		

Appendix C - Book Chapter

This appendix contains the text of a chapter to appear in a reference book published in 2014.

Publisher: Springer-Verlag GmbH Berlin Heidelberg

Title: Modelling and Processing for Next Generation Big Data Technologies and Applications

Editors: Fatos Xhafa, Leonard Barolli, Admir Barolli

This chapter, entitled 'Big Data, Unstructured Data and the Cloud: Perspectives on Internal Controls', was inspired by the paper presented at the WAINA 2013 conference in Barcelona and included as article number 5 in Appendix D. The editors of the book requested an expanded and more wide-ranging discussion of internal controls over data as a topic of interest to their readers.

The chapter draws upon Chapters 5 to 9 of this thesis in discussing the results of the Swiss survey on attitudes towards data and the cloud and in setting these in the context of internal control, audit and compliance requirements.

Big Data, Unstructured Data and the Cloud: Perspectives on Internal Controls

David Simms, Haute Ecole de Commerce, University of Lausanne, Switzerland

david.simms@unil.ch

Abstract. The concepts of cloud computing and the use of Big Data have become two of the hottest topics in the world of corporate information systems in recent years. Organizations are always interested in initiatives that allow them to pay less attention to the more mundane areas of information system management, such as maintenance, capacity management and storage management, and free up time and resources to concentrate on more strategic and tactical issues that are commonly perceived as being of higher value. Being able to mine and manipulate large and disparate datasets, without necessarily needing to pay excessive attention to the storage and management of all the data that are being used, sounds in theory like an ideal situation. A moment's consideration reveals, however, that the use of cloud computing services, like the use of outsourcing facilities, is not necessarily a panacea. Management will always retain responsibility for the confidentiality, integrity and availability of its applications and data, and being able to develop the confidence that these issues have been addressed. Similarly, the use of Big Data approaches offers many advantages to the creative and the visionary, but such activities do require an appropriate understanding of risk and control issues.

1 Introduction

This chapter will set out the risks related to the management of data, with particular reference to the traditional security criteria of confidentiality, integrity and availability, in the contexts of the wider use of unstructured data for the creation of value to the organization and of the use of Big Data to gain greater insights into the behaviours of markets, individuals and organizations. Of particular interest are the questions of identifying what data are held internally that could be of value and of identifying external data sources, be these formal datasets or collections of data obtained from sources such as social networks, for example, and how these disparate data collections can be linked and interrogated while ensuring data consistency and quality.

Much of the current debate around big data technologies and applications concerns the opportunities that these technologies can provide and where issues of security and management are addressed; there is a tendency for these to be considered somewhat in isolation. This chapter will set out the risks, both to the owners and the subjects of the data, of the use of these technologies from the perspectives of security, consistency and compliance. It will illustrate the areas of concern, ranging from internal requirements for proper management to external requirements on the part of regulators, governments and industry bodies. The chapter will discuss the requirements that will need to be met in

respect of internal control mechanisms and identify means by which compliance with these requirements can be demonstrated, both for internal management purposes and to satisfy the demands of third parties such as auditors and regulators.

The chapter will draw upon the author's wide experience of auditing the IT infrastructures of organizations of all sizes in describing the processes for identifying relevant risks and designing appropriate control mechanisms. The chapter will contain discussion of the standard frameworks for implementing and assessing controls over IT activities, of information processing and security requirements, objectives and criteria, and of monitoring, testing evaluating and testing control activities. The author will also apply his insights into the strategies being followed, and initiatives being considered, by a range of organizations and corporations in order to illustrate how risks are changing as technologies change and how control activities need to develop in response. This analysis will be based upon the results of a survey performed within Switzerland in 2011 and 2012 that set out to establish an overview of how organizations viewed and understood both cloud technologies and the nature and value of the data that they held, and what impact these concepts and opportunities would have on their strategies, policies and procedures.

2 Preliminary concepts and definitions

Many of the terms used in this chapter are reasonably recent coinages and definitions can still be flexible and varied.

For the purposes of this chapter, we will follow Bernard Marr [1] with the definition that "Big data refers to our ability to collect and analyze the vast amounts of data we are now generating in the world. The ability to harness the ever-expanding amounts of data is completely transforming our ability to understand the world and everything within it." The fundamental idea is of the accumulation of datasets from different sources and of different types that can be exploited to yield insights.

According to an article by Mario Bojilov in the ISACA Now journal [2], the origins of the term come from a 2001 paper by Doug Laney of Meta Group. In the paper, Laney defines big data as data sets where the three Vs—volume, velocity and variety—present specific challenges in managing these data sets.

Unstructured data as a concept has been identified and discussed since the 1990s but finding a definitive and all-encompassing definition of what this might be is surprisingly difficult. Unstructured data have been defined vaguely positively by Manyika et al [3] as "data that do not reside in fixed fields. Examples include free- form text (e.g., books, articles, body of e-mail messages), untagged audio, image and video data. Contrast with structured data and semi-structured data" which provides a definition but one which is rather more in respect of characteristics that the data do not possess rather than what they are.

To complicate matters slightly, Blumberg and Atre [4] discussed the basic problems inherent in the use of unstructured data a decade ago. They wrote "The term unstructured data can mean different things in different contexts" and that "a more accurate term for many of these data types might be semi-structured data because, with the exception of text documents, the formats of these documents generally conform to a standard that offers the option of meta data" (p. 42).

Cloud computing is described by NIST as “a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g. networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction” [5], while the Cloud, according to Manyika et al, is “a computing paradigm in which highly scalable computing resources, often configured as a distributed system, are provided as a service through a network.” For corporate users, the immediate impact of using cloud services is not necessarily clearly distinguishable from that of using traditional outsourced services, with services and/or data being available through an external network connection. The key distinction is that unlike conventional outsourcing, where typically the service provider contracts to store, process and manage data at a specific facility or group of facilities, in the context of the cloud the data could be stored anywhere, perhaps split into chunks, with no external visibility over how that was organized.

For the individual, private user, the cloud, represented by such services as Dropbox, iCloud or Google Services, for example, is exactly as its name suggests, a virtual and distant and slightly opaque facility by which services are provided without significant identifying features or geographical links.

Internal Controls are the mechanisms – the policies, procedures, measures and activities – employed by organizations to address the risks with which they are confronted. To define a little further, these activities fall into the framework of control objectives are the specific targets defined by management to evaluate the effectiveness of controls; a control objective for internal controls over a business activity or IT process will generally relate to a relevant and defined assertion and provide a criterion for evaluating whether the internal controls in place do actually provide reasonable assurance that an error or omission would be prevented or detected on a timely basis [6].

The traditional triad of information security objectives consists of Confidentiality, Integrity and Availability [7].

- Confidentiality is the prevention of the unauthorized disclosure of information, providing the necessary level of security.
- Integrity is the prevention of the unauthorized modification of information, assuring the accuracy and integrity of information.
- Availability is the prevention of the unauthorized withholding of information or resources, ensuring the reliable and timely access to the information for authorized users.

To give a concrete example of how control objectives and internal controls fit together, an organization might be concerned about unauthorized access to its data. The control objective might be to ensure the confidentiality of the data, and one of many possible control activities might be to perform regular reviews of the appropriateness of user access rights at the application level, to ensure that there are no redundant or unallocated accounts. Clearly in a well-controlled environment there will be a number of internal controls relating to each activity and each objective, and management will draw their comfort from the combined effectiveness of these controls.

At the same time there are a number of other objectives in respect of information security that may be of greater or lesser significance for different organizations, depending on the data they hold and process and the sectors in which they operate. These areas, which are included in standards such as those published by the ISO [8] and NIST [9], among others, include:

- Authenticity / authentication: having confidence as to the identity of users, senders or receivers of information;
- Accountability: ensuring that the actions of an entity can be traced uniquely to that entity.
- Accuracy: having confidence in the contents of the data;
- Authority: the means of granting, maintaining and removing access rights to data;
- Non-repudiation: making it impossible for the person or entity who has initiated a transaction or a modification of data to deny responsibility after the event;
- Legality: knowing which measures are appropriate for the legal frameworks within which an organization is operating.

From a point of view of completeness, it should also be mentioned that there are a number of other classifications of information security criteria. In 2002, for example, Donn Parker proposed an extended model encompassing the classic CIA triad that he called the six atomic elements of information [10]. These elements are confidentiality, possession, integrity, authenticity, availability, and utility. Similarly, the OECD's Guidelines for the Security of Information Systems and Networks, first published in 1992 and revised in 2002 [11], set out nine generally accepted principles: Awareness, Responsibility, Response, Ethics, Democracy, Risk Assessment, Security Design and Implementation, Security Management, and Reassessment.

3 State of the art

Cloud computing is a paradigm in computing that has emerged from the spread of high-speed networks, the steady increase in computing power, and the growth of the Internet. The interconnection of resources that may be separated by significant distances is allowing users access to what appears to be a single resource. As a result, there is an opportunity to outsource resource-intensive or resource-specific tasks to service providers who will deliver the service for a fee, often based on consumption, rather than developing and maintaining the infrastructure and competences in-house.

Neither of the central ideas in this model, distributed computing or outsourcing, is new. Distributing resources within and across networks has been performed since the days of mainframe computers, where data were entered on remote terminals and processed centrally, while outsourcing of computing activities and services has taken place for many years for strategic, operational and financial reasons.

In technical literature reference is often made to cloud computing, grid computing and utility computing; there is no real consensus over whether there is a distinction to be made between the three and, if so, whether the distinctions are clear or are rather subtle. In general, though, the models are identified by features such as ubiquitous access, reliability, scalability, virtualization, the exchangeability of and independence from location considerations, and cost-effectiveness [12].

A key area of growth is that of the types of service that can be offered by providers. Historically providers typically offered remote data processing, storage and systems management as major services, but the newer paradigm is to group offerings together as types of services: Software as a Service (SaaS), Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Service-Oriented Architecture (SOA). The underlying principle is that every facet of computing activities can be offered as services to consumers to be paid for on a usage basis. Computing is offered as a utility with billing based on consumption, with a

corresponding shift towards IT spending being classified as expenditure on a service rather than infrastructure investment [13].

This paradigm has been likened to the electrical power grid, both in the impact it has on social and industrial development and because of its nature: elements of the grid can be owned and operated by different entities in different locations and might not share any physical characteristics, but the users, those who pay for the service, will only see a single interface or point of contact and will consume a consistent and homogenous service [14].

3.1 The advantages of cloud services

The advantages of such services for certain users are, at least in principle, clear. The users need not worry about maintaining the application infrastructure, with all the operating system, database and application patches and upgrades that are necessary. Nor need the users worry about data management practices such as backups or transfer between machines and environments. Both individuals using office-style applications and remote storage services and large corporations using large-scale applications remotely will see the benefits of avoiding questions of ongoing software compatibility and upgrade paths. Essentially these services are seen as an opportunity to avoid having to deal with certain aspects of the complexity involved in managing a technological infrastructure.

For corporate users there are also the advantages related to the use of scalable architectures. Instead of making potentially significant capital investments in hardware and software that might never be fully exploited and thus represent an unnecessary cost to the business, management will be tempted to subscribe to the model of being able to request and exploit additional resources when required, and for as long as required, on something like a Pay-As-You-Go basis. This has the potential to allow a much more accurate and dynamic management of costs while still permitting absolute flexibility in the access to and use of resources.

The advantages for service providers of operating in this sector also seem to be well established. Once a data center has been established and the reasonably fixed costs of operation have been established and incorporated into the business model, the variable costs of service provision and marginal costs related to the acquisition of additional clients or providing additional services or additional capacity to existing clients should be straightforward to manage and reasonably simple to recover (and exceed, of course) through appropriate pricing. Thus the provision of such services can be viewed as a potentially lucrative business, with important initial investment but steady and permanent future revenue streams.

Conceptually the techniques of cloud computing are not significantly different to those of traditional outsourcing of IT services. Many service providers offer outsourced data processing, system management, monitoring and control services and these services are very popular with many organizations who either do not wish to have internal IT services and competences for organizational reasons or prefer to outsource for financial or logistical purposes.

Key elements of any outsourcing agreement are the contract terms and the service level agreements. With these terms, it is clear to both parties which services are being provided, what resources are being made available, how these resources are being managed, and how the quality of service can be evaluated and managed.

Many structured cloud services will provide such terms and conditions but might not be able to specify every element of interest to the customer, such as the precise location where data are stored or processed.

It is in the nature of any kind of outsourcing or service provision activity that both parties will wish to maximize their revenues and benefits from the agreement while at the same time accepting the minimum level of responsibility for addressing the risks involved and for handling any issues that arise. In the context of cloud computing, customers need to pay particular attention to the clear definition of roles and responsibilities in order to try to avoid situations of blame-shifting and cost avoiding should problems subsequently arise.

A key driver for the use of cloud facilities is costs: organizations might not wish to tie up capital in IT infrastructure that might not be used to capacity and which might only have a short life before obsolescence, when there might exist the possibility of renting services as a regular P&L charge. Along with the resource and competency questions, which of course also incur ongoing costs and can be expensive to update or replace, this has long been a prime mover for outsourcing services. Experience in the domain of outsourcing, however, reveals that the cost savings may not always be as significant as hoped for. As mentioned above, the prudent management team will look to ensure that its systems and data are secure and reliable by implementing additional internal controls to generate evidence that there are no weaknesses in the service provider's controls that can be or are being exploited. Typically these internal controls will take the form of data and transaction analysis to identify exceptions, or the appointment of specialist staff that can manage the relationship with the service provider and ensure that trends are identified, that service levels are maintained, and that issues are identified, reported and rectified. Very often the costs associated with implementing and operating such additional internal controls in an effective and robust manner are such that they can eat into the margins created by the whole outsourcing initiative.

3.2 The disadvantages of cloud services

If the advantages for users of cloud services, as set out above, appear to be obvious, then the disadvantages are equally clear, particularly if viewed from a management and control perspective. The management of organizations always retains responsibility for the security and availability of their systems and data, in particular for the three key attributes of confidentiality, integrity and availability. Ensuring that these attributes are understood, managed and evaluated has traditionally been a significant challenge in systems management, even when systems and data are kept within a well-defined and efficiently policed perimeter. Once this perimeter is extended outside the sphere of direct control of relevant management, the challenge becomes increasingly difficult.

In an article published in the Gartner blogs [15], Thomas Bittman argues that the widely-used analogy of cloud computing to provision of electricity or water supplies is not particularly illustrative, for two reasons: "(1) Computing is a rapidly evolving technology, and (2) Service requirements vary widely for computing. Electricity production and distribution hasn't evolved much since the invention of AC that made distance distribution possible. How many forms of water are needed around the world? It's H₂O – maybe it can be potable, purified, or come at a special temperature, but it's still pretty basic stuff."

His analogy is that of transmitted music: radio broadcasts of music began in 1916 and phonograph cylinders were used for storage, with the idea being that people would not wish to bear the expense of storing their own copies of music when it was available over the

airwaves. But technologies have advanced, both for the broadcast and transmission of music and for the local storage and consumption. Individuals continue to be prepared to pay a premium, for infrastructure, material and content, for a quality and rapidity of service.

The choice between broadcast and local access to music is the same as the cloud computing question: it will depend on a number of factors and requirements that are constantly evolving and ultimately become a question of the best balance between costs, risks, and quality. A key element in the cloud computing choice will be assessing the development of requirements and adopting a strategy that will maximize the Return on Investment, however that is calculated.

Such analogies do reach their limits, however, because they do not consider the specific nature of the service provided. When plugging in a laptop in a foreign hotel room, the consumer broadly speaking does not care who has provided the electricity and the identity of the provider has no impact on the use of the service. Using cloud services is fundamentally different, though, for as soon as a provider is storing data or performing processing for a client, there emerge questions of the confidentiality and integrity of the systems and data as well as the availability of services.

A further aspect in which the analogy breaks down is in the area of the difficulty of changing approach once decisions have been taken. There is no question of not continuing to use electricity to power devices, but there will always be questions about the most efficient way to have access to and use IT infrastructure and resources. Moving from one cloud solution to another is likely to prove as complicated, if not more, than moving from one traditional outsourcing provider to another.

3.3 Data storage and the cloud

The opportunity to exploit the storage potential of the cloud can feasibly be viewed as an encouragement to organizations to place less priority on good practices in relation to data management. As storage space increases, so does the tendency simply to store all data rather than to classify, prioritize, archive and delete them as appropriate, and there is nothing in the principles of cloud computing to reverse this tendency.

It is therefore reasonable to envisage situations in which individuals and organizations simply do not know what data they have placed in the cloud. There may be multiple copies of the same documents and datasets. There may be inconsistent and incoherent datasets. There may be quantities of unstructured data – data extracted from central systems and databases that are not in standard, easily recognized and easily classified structures – that by their nature have not been classified and evaluated.

Of course, the question of unstructured and unmanaged data is not unique to the cloud: it exists in virtually all IT environments. Where it becomes particularly significant in the cloud environment, however, is as a consequence of the lack of visibility the data owners have over their data. Access to in-house, and properly secured outsourced data collections, can be defined, logged, and reviewed, at least in theory given sufficient resources and sufficient motivation on the part of the data owners. Once the data are in the cloud, however, the owners have very little visibility and control over the security of their data [16]. If this weakness is exacerbated by an absence of real knowledge of what data are out there, the risk of data loss and disclosure is increased.

The classical model for many businesses and other large organizations is to have one or several centralized key systems in which all of the organization's financial and operational activities are recorded. Typically the data used in key applications and upon which users depend are stored in structured, centralized and at least theoretically well-controlled databases.

There is also, however, widespread use of non-centralized data repositories. Individual departments or users may have their own specific applications that do not fall into the overall organization-wide systems landscape and thus are subject to different standards and procedures for management and control. This is not to say that these systems and data are necessarily badly controlled, simply that management cannot necessarily be certain that standard policies and procedures are being applied.

3.4 The use and frequency of unstructured data

In the modern business environment great use is made of unstructured data in a variety of contexts. Typically users will extract data from central databases, often those underlying ERP systems, and then manipulate these data in a variety of ways as part of their business activities. Such data are thus often found in spreadsheets, user-built databases and desktop or departmental server-based applications. These data can also be found in text files, pdfs, and even multimedia formats, depending on the use to which they have been put.

From an internal controls perspective the presence and use of unstructured data can pose numerous problems in respect of the confidentiality and integrity of data, two thirds of the famous "CIA" triad of information security objectives (the third being availability). When data are located in a structured central database, they can, at least in theory, be controlled, managed, verified and secured. Once extracted from the database, though, they can easily escape the internal control environment [17]. They can be used for decision taking without the assurance that they are still current, complete or valid. They can also, depending on the security measures in place and the efficiency with which these are enforced, leave the organization easily, typically on the USB media that have become ubiquitous, on laptop hard drives, or even as attachments to emails.

4 Problems, issues and challenges

There are several underlying problems related to the management of multiple datasets and of collections of unstructured data.

The fundamental problems related to the management and control of data stored in the Cloud are not dissimilar to those encountered when using outsourcing facilities or third-party service providers. Simply stated, as soon as data or applications are moved outside the perimeter of integrated and monitored internal controls, management have less control over their data.

Typically in an outsourced environment there are a number of ways for the organization purchasing the services to acquire at least a reasonable level of control and assurance. These can be described as obtaining comfort from audit procedures performed by the organization itself or by an audit entity specifically mandated to audit on its behalf; by obtaining a service auditor's report for the environment in question; or by implementing, performing and monitoring the success of a series of internal control activities.

4.1 Self-performed audit procedures

The first of these methods is the most direct and is identical in form and process to the use of internal audit functions to evaluate and report upon business operations that are carried out in-house, as well as the use of the findings of external audits, where internal controls are often evaluated in the context of a controls-based audit. The audit team works with the organization's management to define scope and timeframe, documents the business procedures, identifies the key control activities, assesses these for both **design effectiveness** (whether as designed the control activities do address the risks to which they relate and the control objectives to which they correspond) and **operating effectiveness** (whether the control activities are actually working as designed, with necessary inputs being received, the control being performed, and appropriate conclusions being drawn and actions undertaken). Based on the results of these audit activities, management will be in a position to draw informed conclusions on whether or not internal controls are working as intended.

The right to perform such audits will need to be specified in the contract for service provision.

The difficulties of performing such reviews in an outsourced environment are many. Firstly, the organization may not have an appropriately skilled, experienced or available audit department – it is frequently the case that internal audit functions tend to have greater expertise in, and requirements to focus on, financial management issues such as the use and disposal of assets, Value For Money, and purchase and inventory management. Detailed technical skills in the field of information systems are more likely to be found in larger and more IT-dependent organizations.

Secondly, it might be prohibitively expensive to mandate a third party to perform the necessary audit on their behalf.

Thirdly, it might be logistically difficult to identify all the areas in which processing or data storage or IT management tasks are performed at the service provider, especially if multiple sites in multiple locations are used. This will add to the complexity of scoping and scheduling such an audit.

Fourthly, but far from being the least significant issue, there is the impact of being audited on the service provider. In order to respond to audit questions, staff and documentation need to be made available, and then resources need to be provided to assist with the testing of individual controls, extracting system information and reports and explaining their contents. This can have multiple impacts on the service provider, taking up human and technical resources, distracting staff from their daily activities, and using office space. If a service provider were to allow all of its clients to pay audit visits independently and at times that suited them, it is not difficult to imagine that this could cause significant disruption to the provider's activities.

4.2 Third-party (service) audit procedures

Many service providers thus prefer the second method of allowing their clients to obtain comfort over internal controls: the service audit report. In this situation, the service provider itself mandates an external auditor to review its internal controls and provide an opinion on the effectiveness and efficiency of these controls [18]. This report is then made available to the service provider's customers who can incorporate its findings into their own evaluation of the control environment.

The advantages of such an approach for the service provider are clear. They meet their requirements to provide reliable information about their control environment, while avoiding the inefficiencies of having multiple audit teams visiting their premises and having to provide repeated briefings, explanations and copies of documentation. If the audit partner chosen has a size and structure that allows consistency of team membership and a minimum of rotation, there will also be the advantage of familiarity year to year with business practices and documentation standards as applied by the service provider, which will also increase the efficiency of the audit process.

The report issued by the auditor can take many forms, but for many years the de facto international standard to be followed was the Statement on Auditing Standards No. 70: Service Organizations, commonly abbreviated as SAS 70.

This standard specifies two kinds of reports, named Type I and Type II. A Type I service auditor's report includes the service auditor's opinion on the fairness of the presentation of the description of controls that had been placed in operation by the service organization and on the suitability of the design of the controls to achieve the specified control objectives (thus corresponding to the concept of design effectiveness as discussed above). A Type II service auditor's report includes the information contained in a Type I service auditor's report and also includes the service auditor's opinion on whether the specific controls were operating effectively during the period under review. Because this opinion has to be supported by evidence of the operation of those controls, a Type II report therefore also includes a description of the service auditor's tests of operating effectiveness and the results of those tests.

SAS 70 was introduced in 1993 and effectively superseded in 2010 when the Auditing Standards Board of the American Institute of Certified Public Accountants (AICPA) restructured its guidance to service auditors, grouping it into Statements on Standards for Attestation Engagements (SSAE), and naming the new standard "Reporting on Controls at a Service Organization". The related guidance for User Auditors (that is, those auditors making use of service auditors' reports in the evaluation of the business practices or financial statements of organizations making use of the facilities provided by service providers) would remain in AU section 324 (codified location of SAS 70) but would be renamed Audit Considerations Relating to an Entity Using a Service Organization. The updated and restructured guidance for Service Auditors to the Statements on Standards for Attestation Engagements No. 16 (SSAE 16) was formally issued in June 2010 and became effective on 15 June 2011. SSAE 16 reports (also known as "SOC 1" reports) are produced in line with these standards, which retain the underlying principles and philosophy of the SAS 70 framework. One significant change is that management of the service organization must now provide a written assertion regarding the effectiveness of controls, which is now included in the final service auditor's report [19].

Internationally, the International Standard on Assurance Engagements (ISAE) No. 3402, *Assurance Reports on Controls at a Service Organization*, was issued in December 2009 by the International Auditing and Assurance Standards Board (IAASB), which is part of the International Federation of Accountants (IFAC). ISAE 3402 was developed to provide a first international assurance standard for allowing public accountants to issue a report for use by user organizations and their auditors (user auditors) on the controls at a service organization that are likely to impact or be a part of the user organization's system of internal control over financial reporting, and thus corresponds very closely to the old SAS 70 and the American SSAE 16. ISAE 3402 also became effective on 15 June 2011.

The importance of understanding the related guidance for user auditors (which also applies, by extension, to user management) is critical. When a service audit report is received, the reader need to follow a number of careful steps in order to be certain that the report is both useful and valid before beginning to draw any conclusions from it.

These steps include:

1. Confirming that the report applies to the totality of the period in question. This is of particular importance when using a service audit report in the context of obtaining third-party audit comfort for a specific accounting period, but also applies to more general use of the report. If management's concern is over the effectiveness of internal controls for the period from 1 January to 31 December 2012, say, and the audit report covers the period from 1 April 2012 to 31 March 2013, how valid is it for management's purposes and how much use can they make of it? In the simplest of cases, a prior period report will also be available that will provide the necessary coverage, but in other circumstances this may not be the case and management will have to turn to other methods to acquire the comfort that they need. These could include obtaining representations from the service manager that no changes had been made to the control environment during the period outside the coverage of the audit report and that no weaknesses in internal control effectiveness, either design or operational, had been identified during that period. Other methods could involve the performance and review of internal controls within the client organization, a subject to which we will return below.
2. Confirming that all the systems and environments of operational significance to the organization were included in the scope of the audit report. Management will need to have an understanding of the platforms and applications that are being used for their purposes at the service provider, including operating systems, databases, middleware, application systems and network technologies, and be able to confirm that these were all appropriately evaluated. Very often in the case of large service providers a common control environment will exist, under which they apply identical internal control procedures to all of their environments: if the auditors have been able to confirm that this is the case, then it is not inappropriate for them to test the internal controls in operation around a sample of the operating environments, and management can accept the validity of their conclusions without needing the confirmation that their particular instance of the database, for example, had been tested. Should the coverage of the audit not meet management's requirements, again management would need to evaluate the size and significance of the gap and consider means by which they could obtain the missing assurance.
3. Understanding the results of the work done and the significance of any exceptions or weaknesses noted by the auditors. Generally speaking, if the work has been performed to appropriate standards and documented sufficiently, and if the conclusions drawn by the auditors are solidly based on the evidence, this step should be reasonably straightforward. Management should, however, guard against skipping to the conclusion and, if the report contains it, the service provider's management attestation, and blindly accepting the absence of negative conclusions as being sufficient for their purposes. Each weakness in the design or the operation of controls should be considered, both individually and cumulatively, to identify any possible causes for concern. This is because it is possible that weaknesses identified during the audit, considered to be insignificant within the overall framework of internal controls could potentially be significant in respect of the specific circumstances (use of systems and combination of technologies, for example) of one particular customer.

4.3 Other procedures

Should such an independent audit report not be available, or only provide partial coverage of the control environment or the period in question, and should it in addition be impossible or impractical to the organization to perform their own audit procedures at the service provider's premises, a third way of obtaining comfort over operations is needed. This method can be summarized as consisting of two groups of activities: internal controls performed within the organization over the activities of the service provider, or audit procedures designed to evaluate the completeness, accuracy and integrity of the service provider's processing and outputs.

Firstly, internal controls can be designed that allow local management to monitor and evaluate the activities of a third party. These can take the form, for example, of procedures to track the responses of the service provider to requests for changes: if there is a process by which the customer asks the service provider to change access rights to an application or a datastore to correspond to the arrival or departure of a member of staff, management can track these change requests and the responses of the service provider in order to ensure that the correct actions have been undertaken.

Very often the contract terms with the service provider will include regular meetings and reporting mechanisms through which the provider will present status updates, usually in the form of progress against KPIs and lists of open points, and management can ask questions and ensure that everything is under control. These meetings can form the central points of control activities for management, as a structured means of ensuring that they are monitoring the performance of the service provider in a regular and consistent way, observing long-term trends and identifying anomalies.

Secondly, audit procedures can be designed along similar principles to the above, based on the expected results from the service provider's activities. An example of this would be extracting at the period end a list of user access rights to the organization's applications and comparing these to expectations, expectations based on management's understanding of the access requirements and on the instructions given during the year to create, modify or delete access rights. If the rights correspond, this can provide a layer of assurance to management that both the service provider's internal procedures and controls are operational and that access to their applications and data is being appropriately managed.

With an appropriately selected range of internal controls and audit procedures, management can therefore obtain a certain level of comfort over the existence and quality of the control environment in place at the service provider.

4.4 Timescales and logistical concerns

It is important to recognize that the process of achieving the confidentiality, integrity and availability of data is likely to be lengthy. Data growth is one of the most challenging tasks in IT and business management, with increasing quantities of data being generated in-house within organizations and being acquired from external sources. It is therefore essential to develop a structured plan for overall data management within with the steps necessary for data accumulation, security and transfer can be carried out in a systematic and reliable manner.

In practical terms, management needs to take a number of operational decisions at an early stage in the planning process. Even before decisions can be taken about which data should

be retained within an organization's infrastructure and which should be outsourced, more fundamental decisions need to be taken in respect of how, where, and by whom data will be cleaned, structured, collated, error checked, completed, or even clearly marked for deletion or a separate archiving process.

The "how" aspect will almost inevitably be addressed by a combination of automated tools and manual intervention. Software will be necessary for processing and transforming large quantities of data, but human intervention will be essential for defining and implementing parameters, reviewing output, and making ongoing decisions.

The "where" aspect throws up a question that anticipates, rather, the questions of data security that the move towards outsourcing and third-party storage also throws up. Arguments could be made on the grounds of costs and logistics that the data preparation process should be performed offsite at a third-party site, on purpose-built infrastructure and away from the organization's internal networks. This would be in order, for example, to prevent excessive strain on resources caused by intensive processing and by large quantities of data passing across the network. Such a solution would, however, introduce the additional complication of ensuring adequate security over the data once the datasets have left the organization's security perimeter.

The "by whom" aspect will concern the use of internal resources, insofar as they can be spared, and external resources. Depending on the amount and the nature of the data to be processed, it may be appropriate to bring in resources from outside to perform aspects of the work. Internal resources will always be necessary, however, both from IT to provide technical input, and from the business side as users who understand where the data come from, what they represent, and how they relate to each other. A common failing in any data cleaning or migrating exercise is to view it as a purely technical procedure, whereas in practice the input from experienced and knowledgeable data owners and end-users is critical.

The timescale for such a project will depend on several factors, including the quantity and the nature of the data to be prepared, the availability of resources, and the priority set by senior management for the process. Experience of such projects would indicate that management should be setting their expectations in terms of months rather than weeks, however, and that if data quality and security are really expected, there is no scope for cutting corners.

5 Proposed approach and solutions

In order to gain an independent perspective on how the questions of cloud facilities and unstructured data were affecting organizations in Switzerland, the author performed a survey in 2011 and 2012.

5.1 Background to the survey

50 questionnaires were sent out to contacts in 43 organizations across Switzerland, using the lead author's business experience to identify correspondents across a range of industries and sectors of activity who would be likely to respond. Completed questionnaires were received from 34 organizations, with three sending two responses and two sending three.

The organizations that declined to respond did so on the grounds of confidentiality, not wishing to divulge details of their IT strategy or approach to security to a third party.

The organizations were selected in order to provide a wide cross-section of the range of IT environments and attitudes to the management of data and the use of new technologies. Of the organizations that did respond, five were publicly owned, eighteen privately, and the remaining eleven were public administrations or NGOs. The breakdown by organization size also shows variety: eight with less than fifty employees; ten with between fifty-one and one hundred; five with between one hundred and five hundred; and eleven with more than five hundred.

The rationale behind sending multiple questionnaires to the same organization was to attempt to discover whether there would be cases of poor communication of strategy or developments within those organizations. It is the author's experience of large corporations in particular that there can be differences in the understanding of the overall strategy, the firm objectives, the initiatives undertaken and the impact on users between top management, middle management and IT management, for example.

5.2 The Survey

The questions in the survey were split into two sections, dealing with data security within the organization and with data storage in the cloud, as follows:

Table 1. Survey questions

Section A

Q1	Is there an overall policy concerning data management, security and retention?
Q2	Is there awareness at the level of senior management and/or security management of the concept of "unstructured data"?
Q3	Has there been any assessment of whether the organization should be concerned, from security or efficiency perspectives, about the existence of unstructured data?
Q4	Has the organization established any guidelines on the classification of data according to their sensitivity, age and relevance?
Q5	Has there been any structured attempt to identify and quantify the data stored around the organization outside centralized databases?
Q5A	If so, was this a manual process?
Q6	Does the organization have policies on the extraction, use and storage of data by end-users?
Q6A	If so, is compliance with these policies monitored and enforced, and how?
Q7	Are there policies and/or restrictions in place of the use of removable storage media for file transfer or storage?
Q7A	If so, is compliance with these policies monitored and enforced and how?
Q8	Does the organization have any mechanisms for determining whether there have been breaches of security or confidentiality in respect of its sensitive data?

Section B

Q9	Is the organization using, or planning to use, cloud computing services for the storage of data?
Q10	If yes, have policies and guidelines been drawn up for the nature of data that can be stored in this way?
Q11	If cloud computing is being used, is the organization’s approach based on the centralized management, monitoring and retrieval of data, or do departments and/or individuals retain responsibility for their data?

5.3 Survey results

Table 2. Survey results

Section A				
	Yes	No	Yes %	No %
Q1	32	9	78%	22%
Q2	22	19	54%	46%
Q3	17	24	41%	59%
Q4	9	32	22%	78%
Q5	6	35	15%	85%
Q5A	6	0	100%	0%
Q6	14	27	34%	66%
Q6A	5	9	36%	64%
Q7	12	29	29%	71%
Q7A	8	4	67%	33%
Q8	4	37	10%	90%
Section B				
Q9	39	2	95%	5%
Q10	6	33	15%	85%
Q11	7	32	18%	82%

5.4 Interpretation and analysis of the results

The results demonstrated two major trends in respect of unstructured data. The first was that although 78% of organizations reported having designed and implemented an overall policy concerning data management, security and retention, the exceptions being overwhelmingly small organizations with informal internal control structures, there was little systematic follow-up in terms of the management of unstructured data or in terms of managing and monitoring the use of data by users. 54% reported that senior management were aware of the issue of unstructured data; but only 41% reported that a risk analysis had been carried out to evaluate their exposure; 22% reported that guidelines for the classification of data had been developed and published; and only 15% reported having carried out a structured attempt to identify and quantify the data stored outside centralized databases. In each case it was a large multinational company that had undertaken such an initiative; interestingly, each one reported that the process had been largely manual.

From the perspectives of internal controls and good corporate management, many of these numbers are worryingly low. That more than three quarters of organizations report the existence of an overall policy in respect of data management is a reasonable starting point, but the lower numbers in respect of detailed data management indicate that few organizations had progressed further than the big picture, high-level aspects of data management at the time this survey was performed.

In respect of the management of user activities related to data handling, 34% reported having policies on the extraction, use and storage of data by end users, and only 36% of these were able to report that compliance with these policies was monitored and enforced. Only 29% of organizations reported having policies or procedures in place concerning the use of removable storage media for file transfer or storage, while only 67% of this small subset were able to describe compliance mechanisms. Finally, only 10% of organizations were able to describe the existence of mechanisms for determining whether breaches of security or confidentiality in respect of sensitive data had occurred.

These rather low numbers suggest that organizations will have a significant amount of work to do in collating, cleaning and preparing data. The responses suggest an overwhelming absence of effective internal control to date over data usage, and a lack of certainty over the nature, quality and integrity of the datasets in use around organizations, factors that will automatically increase both the amount of scoping work that needs to be carried out and the quantity of detailed cleaning and tidying of data.

In the cases where more than one response was received from an organization, it was noted that end users were less aware of the existence of strategies and policies than senior management, a situation that is consistent with the author's long experience of auditing large organizations. A key difficulty in the implementation of internal controls within an organizations relates to ensuring that details of the control objectives and activities, their importance, their relevance, their nature and their follow-up, permeate sufficiently through the organization so that controls are operated effectively and efficiently, improving processes rather than hindering them, and with evidence of their performance and output being available for timely review and, if necessary, corrective actions.

In respect of the use of cloud computing facilities for the storage of data, 95% of the responses reported that the organization was using or was planning to use such facilities. The reasons most frequently given were cost management, flexibility, and a desire to streamline IT activities to concentrate on more value-added activities in-house. Of these

organizations, however, only 15% had drawn up policies and guidelines concerning the nature of data that could be stored in this way, and only 18% were adopting an approach based on the centralized management, monitoring and retrieval of data rather than leaving it to departments or individuals to manage.

Once again, the wide absence of overall policies and guidelines and the tendency as reported to adopt potentially uncoordinated approaches to project scope and management runs counter to established good practice in respect of internal controls. Without clear and enforced structure, inconsistency is likely to become a significant barrier to success. In addition, delegating down to departments or individuals increases the risks of decisions being taken without sufficient skills, experience or perspective.

Overall it was possible to draw a clear distinction between large and small organizations and publicly and privately owned ones in respect of their approach to structured and formalized control environments. It was also possible to identify with high accuracy the responses received from publicly-owned corporations and those from organizations subject to other strict and demanding controls requirements such as Sarbanes-Oxley or local industry or environment-specific regimes. It was also possible to identify organizations that had significant internal audit or internal controls functions, or that had been alerted to the risks involved in going through a process of discovery in the context of a legal dispute.

6 Summary evaluations and lessons learned

The tentative conclusions that can be drawn from this very specific and targeted survey are the following.

Firstly, the importance of having a structured and documented approach to data management and security is widely understood among the organizations surveyed, with a particularly positive attitude towards risk management and compliance from larger organizations and those subject to definite compliance regimes because of their ownership or industry. How this understanding actually translates into positive measures designed to ensure compliance is another question, however, and IT security managers in particular reported seeing greater enthusiasm for establishing policies within their organizations than for implementing and complying with the necessary procedures. There was also a question of priorities and resources raised by smaller organizations that did not feel that such policies corresponded to or were a part of their core daily activities.

Secondly, the concept of unstructured data and the particular challenges posed by such data is reasonably widely understood, but there has been little activity outside large publicly-owned corporations to address the issue in any systematic way. In general IT departments had a good grasp of the nature and impact of the matter, and the subject was frequently raised by internal and external auditors and by legal advisers, but it was rarely considered to be a subject of great priority by senior management.

Thirdly, even if thought has been given within some organizations to managing employee access to and extractions of sensitive data, in the majority of cases monitoring is weak and compliance cannot be ensured. Generally speaking, users who have access to data stored in centralized databases tend to have the ability and the opportunity, and frequently the encouragement, to extract those data and use them for analytical or reporting purposes. Once the data have been extracted, access controls around them are usually weaker, often

being restricted to network or workstation access controls, and even these restrictions reach their limits once data are copied onto portable devices such as USB keys.

Fourthly, very few organizations are in a position of being able to detect reliably whether their security has been breached or their data compromised, even when all their systems and data are hosted and managed internally. Indeed, in respect of both security breaches and employee misuse of data, it was reported that incidents were typically identified either by chance or on the basis of information received, rather than on the basis of regular and reliable compliance measures. Of course, in practice being able to design, implement and monitor the operation of such control activities is frequently non-trivial and requires competence, resources and careful planning. Whether organizations would be able to apply such controls to outsourced data, or be satisfied by the monitoring and reporting services provided by their cloud service provider, is the same question with an added layer of complexity.

Fifthly, the idea of cloud computing as a financially attractive option for outsourcing a number of traditional IT activities including data storage is widespread and there is a great deal of enthusiasm for it across a wide range of organizations, but this enthusiasm is not yet being widely and systematically backed up by detailed risk assessments and careful consideration of the approaches needed to identify, classify, manage and monitor the data being transferred into the cloud.

The comments provided alongside the answers also provided useful information. In particular, several respondents referred to the different types of cloud that are beginning to exist: in certain industries such as financial services it would be unthinkable to use cloud services in which confidential data might be stored outside Switzerland or for which it would be difficult to obtain adequate audit comfort over key concerns, but the use of some kind of industry-specific Swiss cloud, perhaps set up as a joint venture, with appropriate controls and safeguards in place, might be conceivable.

Several respondents also flagged up the importance of the proper management of backup media, which of course need to be subject to the same policies and procedures for data management and security as live datasets. From the perspective of the management who will retain the responsibility for ensuring the availability and reliability of system and data backups for good practice and going concern reasons, and for the auditors who will be verifying this, it will be a challenge to identify exactly which data are backed up where, how this is managed from a security and availability perspective, and what the timelines, sequences and interdependencies would be for restoring part or all of a missing dataset.

Within all businesses there is constant pressure to reduce costs and cloud computing could be seen as an effective method of managing and reducing costs, particularly in the short-term. If there are no significant in-house IT systems, a case will always be made for reducing to a bare minimum, or even eliminating entirely, the IT function, thereby reducing staff costs alongside the operational costs of monitoring and maintaining systems.

The decision to choose cloud computing services is not one to be taken lightly. For individuals, the use of personal services in the cloud (such as Facebook, gmail and Dropbox) is, or rather should be, a matter of a calculated assessment of risks and benefits, for the potential negative impacts of breaches in security, for example, can be significant. For businesses and other organizations, the same concerns apply but on a larger scale. For all organizations that have a responsibility to keep their, and others', data secure and

confidential, the decision can only be taken after a detailed analysis of how they will obtain, and continue to obtain, the necessary comfort that this is the case. If they cannot build into their own procedures, into enforceable contract terms, and into audit plans, the means of confirming the confidentiality, integrity and availability of their systems and data, they should not consider externalizing it.

In addition, organizations should not forget the immediate costs and efforts involved in moving onto the cloud. Datastores need to be identified, classified, cleaned up and archived, and serious technical and operational decisions are needed to determine what data will sit where. This will frequently be a project of a significant size requiring expertise, resources and input from a number of people across the business who understand the business, the systems, the data, and their use.

Experience of traditional outsourcing suggests that it is very easy for organizations to overestimate the cost savings generated by a move towards service providers and to underestimate the amount of internal competence and dedicated management required to make a success of such initiatives. As long as the organization relies on its data and retains responsibility for all aspects of its business from a regulatory perspective, it will need to ensure that its management of the relationship with its service providers and its access to critical operational information are both adequate and appropriate. Typically this will require retaining or recruiting skilled, experienced and reasonably senior staff to liaise with and monitor the performance of the service provider.

The long-term consequences of opting for a cloud solution also need to be examined. Once on the cloud, systems and data are likely to stay there, and feasibly with the same provider. What begins as a simple and cost-effective solution to a small problem could develop into a long-term strategic commitment with little scope for alteration.

7 Further research issues

It will be necessary to monitor the development of the use of cloud services for data storage from a controls perspective in order to see how such use develops, whether it becomes widespread and whether it does become a factor in assessing the quality of internal controls. Perhaps practical and effective operating procedures will be developed and become standard very rapidly, so that the subject does not become a major concern for controls managers, or perhaps the take up of the technologies will not be great, because of perceived controls issues or questions of cost or access.

It will also be interesting to study the impact of the uptake of such services on audit opinions and compliance reports. This will surely be a major driver in the development and the use of these services: if organizations find themselves subject to adverse comments from regulators or external auditors, that will necessarily cause a slow-down in adoption. On the other hand, if clean reports are issued and no concerns are raised, take up will only be encouraged.

8 Conclusion

In common with all IT-related projects, initiatives in respect of data collation and accumulation and in respect of data migration and transfer require a great deal of planning, strategic awareness and effective controls in order to ensure that the key security objectives defined by the organization continue to be met. Both managing unstructured data and

managing the migration onto, and ongoing monitoring of, cloud services, present countless opportunities for the loss of the confidentiality, integrity and availability of data, to cite once again just the three most famous security objectives.

The use of large datasets for competitive advantage is highly tempting for many organizations from a tactical perspective, while the use of cloud services is attractive for a number of reasons, financial, operational and strategic. The senior management of organizations tempted by such initiatives should be aware, however, that neither type of project can be successfully completed overnight, and that they should be prepared to provide the necessary resources, guidance, oversight and supervision to ensure that the advantages obtained through the initiatives are not outweighed by decreased security, increased costs, or reduced comfort from internal controls.

9 References

- [1] Marr B (2013) Is This the Ultimate Definition of "Big Data"?. <http://smartdatacollective.com/node/128486>. Accessed 2 September 2013
- [2] Bojilov M (2013) Big Data Defined. In ISACA Now. <http://www.isaca.org/Knowledge-Center/Blog/Lists/Posts/Post.aspx?ID=299>. Accessed 2 September 2013
- [3] Manyika J, Chui M, Brown B, Bughin J, Dobbs R, Roxburgh C, Hung Byers A (2011) Big data: the next frontier for innovation, competition and productivity. McKinsey Global Institute, Washington DC
- [4] Blumberg R, Atre S (2003) The Problem with Unstructured Data. DM Review
- [5] National Institute of Standards and Technology (2011) The NIST Definition of Cloud Computing. Special Publication 800-145, NIST, Gaithersburg
- [6] Committee of Sponsoring Organizations of the Treadway Commission (2009) Guidance on Monitoring Internal Control Systems. AICPA, New York
- [7] Krutz R, Vines D (2010) Cloud Security: A Comprehensive Guide to Secure Cloud Computing. Wiley Publishing Inc, Hoboken
- [8] Internal Standards Organization (2005) ISO/IEC 27001 Information Technology - Security Techniques – Information Security Management Systems. ISO/IEC, Geneva
- [9] National Institute of Standards and Technology (2006) Information Security. Special Publication 800-100, NIST, Gaithersburg
- [10] Parker D (2002) Toward a New Framework for Information Security. In Bosworth S, Kabay M.(eds) Computer Security Handbook, 4th edn. John Wiley & Sons, New York
- [11] Organisation for Economic Co-operation and Development (2002) OECD Guidelines for the Security of Information Systems and Networks: Towards a Culture of Security. <http://www.oecd.org/sti/ieconomy/15582260.pdf>. Accessed 8 January 2014
- [12] Adolph, M (2009) Distributed Computing: Utilities, Grids and Clouds. ITU-T Technology Watch Report 9, ITU, Geneva
- [13] Gantz J, Reinsel D (2011) Extracting Value from Chaos. IDC iView.
- [14] Information Systems Audit and Control Association (2013) Big Data: Impacts and Benefits. ISACA, Chicago.
- [15] Bittman T (2009) A Better Cloud Computing Analogy. Gartner Blogs. http://blogs.gartner.com/thomas_bittman/2009/09/22/a-better-cloud-computing-analogy/. Accessed 8 January 2014
- [16] Information Systems Audit and Control Association (2012) Security Considerations for Cloud Computing. ISACA, Chicago
- [17] Ghernaouti-Hélie S, Tashi I, Simms D (2011) Optimizing security efficiency through effective risk management. Paper presented at the 25th IEEE International Conference on Advanced Information Networking and Applications (AINA 2011), Biopolis, Singapore, 22-25 March 2011.
- [18] American Institute of Certified Public Accountants (2012) Quick Reference Guide to Service Organizations: Control Reports. AICPA, New York.
- [19] American Institute of Certified Public Accountants (2011) Service Organizations: Reporting on Controls at a Service Organization Relevant to User Entities' Internal Control over Financial Reporting Guide. AICPA, New York.

10 Index of terms

- Assurance – the process of obtaining and using accurate and current information about the efficiency and effectiveness of policies and operations, and the status of compliance with the statutory obligations, in order for management to control an organization’s activities.
- Audit comfort – the support needed to draw conclusions provided by audit procedures.
- Audit procedures – the various tests of details, analytical procedures, confirmations and other activities performed to gain audit evidence.
- Big data – the accumulation of datasets from different sources and of different types that can be exploited to yield insights.
- Cloud computing – a computing paradigm in which highly scalable computing resources, often configured as a distributed system, are provided as a service through a network.
- Design effectiveness – that internal control procedures are appropriate in respect of the control objective they are intended to achieve and the risk they are intended to address.
- Internal controls – the activities performed to ensure that policies and procedures are implemented and operated consistently and effectively, allowing errors and omissions to be prevented or detected.
- Operating effectiveness – that internal control procedures are functioning as designed, regularly and consistently.
- Security triad – the three key information security objectives of Confidentiality, Integrity and Availability.
- Unstructured data – data sitting outside structured and organized data repositories such as relational databases.

Appendix D - Published Articles

This appendix contains the relevant publications from 2009 onwards that contributed to the development of the central ideas of this thesis. Each article underwent peer review for acceptance at an international conference, and was presented and discussed in workshops as part of each conference.

1. Information Security Optimization: from theory to practice (2009)

Published in: Availability, Reliability and Security, 2009. ARES '09. International Conference on
Date of Conference: 16-19 March 2009
Page(s): 675 - 680
E-ISBN: 978-0-7695-3564-7
Print ISBN: 978-1-4244-3572-2
INSPEC Accession Number: 10699292
Conference Location: Fukuoka, Japan
Digital Object Identifier: 10.1109/ARES.2009.106

2. Reasonable Security by Effective Risk Management Practices: from theory to practice (2009)

Published in: Network-Based Information Systems, 2009. NBIS '09. International Conference on
Date of Conference: 19-21 Aug. 2009
Page(s): 226 - 233
E-ISBN: 978-0-7695-3767-2
Print ISBN: 978-1-4244-4746-6
INSPEC Accession Number: 11007640
Conference Location: Indianapolis, IN, USA
Digital Object Identifier: 10.1109/NBIS.2009.66

3. Protecting Information in a Connected World: a question of security and of confidence in security (2011)

Published in: Network-Based Information Systems (NBIS), 2011 14th International Conference on
Date of Conference: 7-9 Sept. 2011
Page(s): 208 - 212
ISSN: 2157-0418
E-ISBN: 978-0-7695-4458-8
Print ISBN: 978-1-4577-0789-6
INSPEC Accession Number: 12306821
Conference Location: Tirana, Albania
Digital Object Identifier: 10.1109/NBIS.2011.38

4. Optimizing Security Efficiency Through Effective Risk Management (2011)

Published in: Advanced Information Networking and Applications (WAINA), 2011 IEEE Workshops of International Conference on

Date of Conference: 22-25 March 2011

Page(s): 613 - 619

E-ISBN: 978-0-7695-4338-3

Print ISBN: 978-1-61284-829-7

INSPEC Accession Number: 11976707

Conference Location: Biopolis, Singapore

Digital Object Identifier: 10.1109/WAINA.2011.93

5. Structured and Unstructured Data in the Cloud: a Swiss Perspective on Readiness and Internal Controls (2013)

Published in: 2013 27th International Conference on Advanced Information Networking and Applications Workshops (WAINA)

Date of conference: 25-28 March, 2013

Page(s): 643-648

ISBN: 978-1-4673-6239-9

Conference location: Barcelona, Spain

DOI Bookmark: <http://doi.ieeecomputersociety.org/10.1109/WAINA.2013.63>

(1) Information security optimization: from theory to practice

David Simms, *PricewaterhouseCoopers SA, Lausanne, Switzerland / Faculty of Business and Economics – University of Lausanne*

Abstract— Organizations face a significant challenge in designing and implementing appropriate information security measures. There are many sources of guidance on good and best practice relating to platforms, architectures and industries, but this guidance needs to be interpreted in the context of the specific risks faced by the organization, the desire to mitigate those risks, and the requirements for user friendliness, system performance and system availability driven by the user community. The process of identifying, justifying, implementing and maintaining the correct balance between security and ease of access for authorized users requires careful consideration at a number of phases, including the assessment of risks, the identification of appropriate standards, the definition of policies and the education of users, and organizations also need to implement mechanisms for the regular and effective review and update of the measures taken. This paper discusses the issues involved in implementing an optimized information security policy, the common pitfalls encountered by organizations in this respect, and presents an outline framework for such implementations.

Index Terms— good practices, information security, risk management, security policy optimization

I. Introduction

THE requirement for effective and appropriate information security within organizations of all sizes can be the starting point for a significant effort in research and policy development. Very often, however, such efforts fail to deliver a suitable solution, because no coherent framework for the design and operation of controls is identified and more attention is paid to specific measures to address individual points than to ensuring that the ensemble of measures and controls put into place meets the organization's requirements.

In section II of this paper some of the commonly encountered steps undertaken in implementing an appropriate information security structure are discussed, with the background to each stage described and examples given of typical situations encountered. These steps are drawn from practical experience rather than a specific formal framework, and apply to organizations of a wide range of sizes and natures.

In section III a high-level framework for meeting information security requirements in a practical way is set out, with consideration of some of the benefits of the model. These are illustrated by case studies drawn from the author's professional experience.

Section IV outlines the major drawbacks and limitations of the framework.

II. Typical procedures encountered in industry for defining an organization's approach

A. Identification of standards and external requirements

Depending on the size and nature of an organization and the sector in which it operates, there may be published standards with which it needs to comply in the area of internal controls including formation

security. Typical industry standards include the requirements set out for financial services companies within Switzerland [1], and similar requirements have been defined by and within public bodies such as cantonal administrations [2]. Such standards tend to be expressed in terms of objectives rather than specific technical guidance, however, leaving the question of how they will be responded to and implemented to the organizations themselves and, very frequently, to their auditors and other professional advisors.

Similar fixed objectives are included in the implementation of compliance with legislation such as the Sarbanes-Oxley Act 2002, which introduced the requirement for documented and demonstrably effective internal controls over business processes and information system management within the context of financial reporting. Once again, these frameworks and accompanying standards such as Auditing Standard No 5 (AS5) [3] are more focused on the objectives and the end results than on the mechanisms of achieving compliance.

On a more technical level, a number of standards for general information system management and specific security management have been developed and are widely used in practice. These include: CobiT (Control Objectives for Information and related Technology) [4], a framework of best practices created by the Information Systems and Control Association (ISACA) and the Information Technology Governance Institute (ITGI) in 1992 and revised several times since then; the Information Technology Infrastructure Library (ITIL) [5], which was originally developed under the auspices of the UK Government's Central Computer and Telecommunications Agency; and the family of international standards such as ISO/IEC 27001 and 27002 [6]. These standards provide both objectives for management, at an overall and a detailed level, and also offer specific recommendations for individual controls.

Finally, there is widespread technical information available concerning the security of individual platforms and applications. Many manufacturers and vendors provide detailed best practice information, while audit and consulting companies also have clearly defined standards for what they consider to be appropriate security measures.

Frequent complaints from those charged with corporate governance is that these sources of information provide too much unhelpful high level guidance on objectives, and that the detailed guidance for practical policies and controls consists of a lengthy list of suggestions that would be impractical to implement and unworkable in practice.

B. Risk Assessment

The critical component of the implementation of any information security framework is the performance of an appropriately-scoped risk assessment. Ideally this should be an iterative process involving input from several functions within the organization; often such a risk assessment, if performed at all, falls into the responsibility of the IT function, but in order to address all issues surrounding the organization the working group should include representatives of several functions. These include: IT, who know what information is stored and in what format; business lines, as the data owners, who can define which data is required for daily transactions and can define the sensitivity and confidentiality of each dataset and application; and legal and compliance, who can provide technical input into the regulatory and compliance frameworks within which the organization operates and explain external requirements over restricting access to systems and data.

The risk assessment should be regularly updated in order to address changes in the technical, regulatory and operating environments of the organization. A frequent failing in the compliance process is that reasonable conclusions have been drawn at a particular moment in time but that the area has not subsequently been readdressed and changes in the internal environment, such as the implementation of new systems as a result of changing business requirements, or in the external environment, such as a modification in the data protection legislation in the territory, have not been taken into account.

C. Performance and user accessibility considerations

Once the risk assessment has been performed and validated, consideration needs to be given to the impact of security policies and individual security measures on the usability and performance of corporate systems. Effective information security management can very often be viewed at its most basic level as the challenge of finding the correct balance between robust security measures and providing an acceptable working environment for legitimate users.

The difficulty of finding such a balance can be illustrated with a number of examples. Firstly, there is the question of the overall security parameters for key systems and their impact on the usability of these systems. The classic example from industry of this is the IBM AS/400 platform, which allows the configuration of a single parameter, QSECURITY, as the base setting for system security. This parameter can take any one of five values from 10 to 50, with 10 being the weakest setting and 50 the strongest, providing for the detailed logging of user activities and the use of system functions. In business environments where such logging is both desirable and necessary, such as certain financial services, such a strict setting is of great use to system managers and compliance functions. In mainstream manufacturing environments, however, where this architecture has historically been widely employed, many organizations have found that such a high setting has a negative effect on performance, slowing down the rate of processing of transactions and also generating sizable log files of little interest in the context of ongoing operations. Thus such organizations have the tendency to set this parameter to 30 or 40, which offers a reasonable baseline for system security, and introduce other measures of control at key points in the operating cycles.

Similarly, the question of encryption requires careful consideration. The impact on system performance of implementing full encryption of data flows needs to be measured against the benefits that it would bring, especially in operating environments where users are accessing data from remote sites over slower communications frameworks than the corporate LAN. Another situation where this question arises is in the encryption of hard disks in laptop computers. This is a widespread precaution taken by many organizations to prevent unauthorized access to sensitive applications and data, but this approach also requires the creation of appropriate procedures for resolving problems legitimate users encounter in accessing their systems, especially when travelling away from the organization's centre.

A third consideration that is commonly encountered is in relation to the login process, particularly for remote connections over Virtual Private Networks. A very common model is that of two-factor authentication, where for example a remote user identifies himself with a combination of a personal factor, such as a password, and a technical factor, such as an ID card or token. Such techniques can provide a reasonable level of security over remote connectivity, but consideration needs to be paid to the cost and effort required in implementing the infrastructure to support such technologies, and to the potential for lost efficiency and frustration on the part of remote users should difficulties be encountered in using these technologies to gain access to resources.

A fourth area that commonly needs to be taken into consideration concerns the allocation of local access rights to users' machines. As part of good IT governance, many organizations strongly restrict what users can do on their desktop and laptop computers, often with the intention of preventing the installation of unauthorized or unlicensed software, or the deletion or circumvention of corporate software or security measures, such as anti-virus software or forced routing through proxy servers. A consequence of such an approach, however, can be increased difficulties in remote technical support for users in situations where their PC cannot connect to corporate network remotely for remedial actions to be undertaken.

D. Definition of the approach to information security and the creation of policies

After the completion of the risk assessment and the identification of operational and organizational factors that could have an impact on the strategy, the strategy for addressing information security can be defined. Broadly speaking, there are two underlying philosophies to the management of security; one is to forbid all access unless specifically required, the other is the permit access unless specifically

forbidden. In most business contexts, the former approach is the most commonly adopted, but the latter approach can also be appropriate in certain environments. An example of a scenario in which access rights would by default be restricted to a minimum would be a corporate Enterprise Resource Planning system, such as SAP, where in a good practice context users would need to demonstrate a clear business reason to have access to each transaction. An example of a scenario where access rights could be attributed widely in the first instance could be a software development environment where wide access rights would be useful for file sharing and workflow purposes and there would be no sensitive or secret information held within that system.

Documentation of the strategy and of the information security policies is critical and care should be taken to create a framework of documentation that provides clear definitions of roles, responsibilities and terms used, and that also provides the right level of detail for users. A common and effective approach is to create a brief, overview document summarizing the philosophy behind the organization's approach to security, a document defining terms used throughout the framework, and separate documents specifying the configurations, parameters and related user and management activities for each system or platform. The question of definition of terms is highly important but often neglected and this can lead to inconsistencies in interpretation and non-performance of key activities. For example, specific instructions for managing and monitoring "critical systems" might not be followed unless the definition of "critical" has been appropriately defined and communicated.

E. Implementation and dissemination of policies

Broadly speaking, there are two key audiences for information security policies and the importance of the domain, and the means of communication to each, differ significantly.

The first is the IT function itself, which has responsibility for the implementation of the policies in terms of the configuration of systems and applying any necessary technical changes. Typically communication to this function, and in particular the security function within the IT team, is reasonably straightforward, especially as in many cases these staff members will have been involved in the development of the strategy and the identification of appropriate measures.

The second audience is the organization's user base, and this can be further subdivided into what can be termed the standard users and the users with management responsibilities. Such responsibilities typically involve the ownership of applications and data and the performance of the necessary control activities related to such ownership, such as the approval of changes in user access rights and the regular review of user access rights and system parameters.

For standard users, a practical mechanism for communicating security policies and increasing awareness involves a combination of formality and informality. On the formal side, the security policy can be made a key document to sign upon joining the organization and regular reminders can be implemented in the forms of annual update training and login messages to corporate systems. Informal means can include links to policies on an intranet site or file servers, or non-compulsory update quizzes and employee communications. A fundamental element within all of these activities is to emphasize that all users have a role to play and a responsibility towards the organization in terms of security.

Users with specific security responsibilities require more targeted training and communication, as their effective and efficient performance of security tasks can have a significant impact on the operational quality of the security environment. Typically these users will require specific training on the performance of their control activities, the importance of these activities, and the way exceptions are to be handled. Frequently the regular performance of these controls will be measured as a Key Performance Indicator.

F. Updates to the strategy and policies

As mentioned in section B above in reference to the Risk Assessment, the Risk Assessment, security strategy and security policies are all living documents that require regular and educated updates to remain useful and valid.

As part of good corporate governance, the procedures for regularly updating and reviewing company-level strategy and documents such as these should be clearly defined and responsibility assigned. Typically this process will be annual.

Good corporate governance also suggests that the review and update of such frameworks be addressed as part of the process for managing significant changes to the infrastructure of an organization, the development and implementation of new applications or other systems, or changes to the operating environment of the organization. The requirement to review and update disaster recovery and business continuity plans in response to changes in systems or the business environment is well recognized, and the same reasoning applies to information security management.

Auditors often talk of the maturity of frameworks of internal controls [7] and this is another element that can be taken into account when reviewing the continued validity of the security policies and framework. If a number of key controls are well bedded in and reliable, the reviewer can consider how much comfort they give over the whole control framework and identify areas to introduce additional controls or eliminate or downgrade existing ones.

III. A Framework For Implementing Appropriate Good Practice

To be useful in practice, any framework designed to allow organizations to implement good information security needs to be generally applicable. There are large differences in the requirements, risk profiles and available resources between an international group with a sophisticated technical infrastructure and a small private business with perhaps one member of IT staff, and attempting to follow blindly a one-size-fits-all approach would clearly be inappropriate. There are, however, a number of common themes, some of which are set out in section II above, and these can be drawn together to form a reasonable basic framework for addressing this issue.

The framework could take the following form, with the roles and responsibilities for each activity being clearly defined and real-life examples given of successful and less successful approaches.

A. Definition of overall philosophy

This should be performed by the senior management of the organization as a whole, at least at the approval phase. In many organizations this question is not considered and the outcome is a series of disjointed, separate security measures and control structures, introduced over a period of time by successive IT staff in response to specific incidents or audit reviews, but not designed to reflect a systematic approach to security management. If the organization possesses a Chief Security Officer or an IT Security Manager, it should fall into that job description to propose an overall approach and have this validated and supported by senior management.

B. Analysis of external requirements

In order to develop the most appropriate approach, the organization needs to be aware of the legislative, regulatory and compliance regimes within which it operates. Typically the organization's legal and compliance function, or specialist external advisors, would play a major role in this analysis.

This area can be problematic in circumstances where new requirements are introduced and interpretation is necessary. A good example of this is the introduction of the internal control requirements for corporations affected by Section 404 of the Sarbanes-Oxley Act: in the first year of compliance efforts, many corporations spent significant time and resources discussed the precise impact of the new legislation both internally and with their external auditors and legal advisors.

C. Risk Assessment and Control Objectives Setting

The definition of the measures to be implemented needs to follow a systematic pattern by which the

risks in relation to information security faced by the organization are clearly and rigorously set out, and control objectives defined which are designed to mitigate those risks. From the control objectives, detailed control procedures and policies can be defined.

This process is fundamental because it ensures, if performed properly, that the actual measures put into place do correspond to the real risks to which organizations are exposed.

After the introduction of the Sarbanes-Oxley Act, the businesses affected were obliged to document and critically assess their systems of internal controls. In numerous cases it was found that some of the controls in place were of little value: some were simply not operational; some did not meet the control objectives; and some duplicated other controls and were therefore redundant. In readdressing the framework of controls from the starting point of a rigorous risk assessment, companies were able in many cases to reduce the number of control activities in place in the field of information security and better target the controls towards actual risks.

D. Consideration of operational issues

IT security management needs to anticipate the impact of possible policies and measures on the operations of the organization, both from a technical perspective and that of the daily activities of users of the information systems. Experience shows that procedures that are perceived to waste resources or time, or are not of clearly discernable value to staff, are often circumvented, ignored, or complied with to the letter rather than to the spirit of the instruction, making the policy or control ineffective and potentially weakening the whole structure of security.

One example of this occurred in the application support department of a large business, where development staff were required to request access to the production environment through a formal process and this access was granted for a limited time. Many development staff found this to be a burden and so persuaded a compliant security administrator to create a user account with wide and permanent rights in the production environment. This account was then shared by several members of the development team. When corporate management became aware that unauthorized changes had been made to a key application system, the subsequent investigation was unable to identify the member, or members, of staff responsible, because this shared user ID had been used. Thus the whole control framework around changes to applications had been undermined.

Another example occurred in the finance department of a large company which had clear procedures in place for the regular review of user access rights and placed the responsibility for these reviews on the heads of departments within the business lines. Several such department heads saw no operational value in performing such reviews and hence neglected them. This came to light when management investigated a series of unusual transactions and discovered that the perpetrators had excessive access rights to the main application system; these rights had accumulated over time as a result of the retention of rights following job changes and the cloning of existing user accounts for new users. Careful performance of the regular reviews would have detected these anomalous access rights far sooner.

E. Creation of policies and procedures

The creation of policies and procedures needs to be performed with a focus on clarity, simplicity, and the intended audience. Those documents aimed at a non-technical audience, such as the standard user base, should in particular be written in an explanatory tone, seeking to inform and convince as well as simply set out the policy. In many organizations this process will be a joint effort between IT, Human Resources and the business lines, all of whom have contributions to make.

F. Design of internal controls

Internal controls should be designed to meet the control objectives determined in part C above but also to provide the most effective and efficient combination of activities from which to gain the assurance that risks are being mitigated. In particular, control designers should aim to implement the

most efficient mixture of preventive (real-time) and detective (after the event) controls, based on the risks; there should also be an appropriate mixture of automated and manual controls, as typically automated controls are more reliable and less time-consuming, but a well-designed manual control such as a periodic review or reconciliation can cover multiple objectives; the frequency of the operation of controls should be carefully considered, to ensure that exceptions would be identified and handled on a timely basis; and care should be taken to avoid the implementation of redundant controls that cover the same risk unnecessarily.

Control design should be performed by staff who have a good understanding of the process in question and of its relation to other technical and business processes.

G. Education and training

Experience shows that education about and training in information security are key elements in the implementation of effective security measures. Such activities can range from simple awareness-raising initiatives to encourage staff members to consider the impact of their actions and attitudes on the organization's efforts to maintain effective security, to more formal education sessions aimed at training users on specific aspects or techniques of security. This is another area where IT staff will often work closely with Human Resources and business management to define the requirements of users and the best methods of communicating the key messages.

Effective awareness-raising initiatives commonly seen in businesses include night-time sweeps through premises to seize potential sources of information, such as computers, USB keys, CDs and paper files, that have been left on desks or in public areas overnight, messages shown on-screen during the login process, and regular mail messages from security staff referring to cases of information security failure reported in the media and encouraging staff to bear such stories in mind when handling potentially sensitive information. The European Union agency ENISA (European Network and Information Security Agency), created in 2004, publishes regularly on such questions [8].

H. Review of compliance and effectiveness

It is a truism that policies, procedures and controls are worthless unless management implements and performs systematic reviews of compliance with the policies and of the effectiveness of the controls. Commonly such reviews will be performed by an internal audit function, but often the IT function will be involved in the process in order to provide data and information.

In order to perform such reviews effectively, the organization requires staff with the technical skills to interpret the information available, the business knowledge to understand the significance of the control activity, and the experience and judgment to draw appropriate conclusions.

An example of where all three skills are necessary is in the evaluation of the performance of a control activity in which on a monthly basis a member of IT staff reviews a list of system parameters to confirm that they remain in compliance with the set policy. The technical skills required are self-evident: the report in question consists of a list of parameters with the name and value that need to be recognized and understood. The need for wider business knowledge is also clear, as the reviewer should understand the potential impact of any deviations or exceptions on the operations of the business or on the rest of the control framework. Where the experience and judgment is significant is in assessing the actual performance of the control, as the mere presence of the control owner's signature on the monthly report does not in and of itself demonstrate that the control has been properly performed. The reviewer needs to know how to confirm that the control was indeed performed as designed, with exceptions noted, followed up and rectified if appropriate.

I. Regular review and update

The whole question of information security management within an organization should be subject to periodic, critical review and update. Changing circumstances within and outside the organization, in terms of technology, organization, objectives and the operating environment, mean that the framework

for security management needs to evolve constantly if it is to remain valid and effective. Management needs to assign appropriate resources and appropriate priority to the review and update procedure in order to ensure that it is properly performed.

An old joke in the auditing profession, told wryly from experience, concerns the “dust test”. The experienced auditor, on receiving the hard-copy policies and procedures requested at the beginning of the audit, will first run a finger along the top of the binder to see how much dust has accumulated upon it. This is a reliable indicator of the organization’s attitude towards reviewing its documentation.

IV. Drawbacks and Limitations

This framework is designed to be applied generally across a range of sizes and natures of organizations, but there are a number of evident drawbacks and limitations. The most significant are the following:

1) Resources

Designing and implementing a good quality information security framework can require significant investment in resources by an organization. The methodology outlined in this paper would be more easily applied in an environment where the IT function is large enough to support technical staff dedicated to security and where other departments, such as Human Resources and the business lines, had sufficient capacity to allow participation in the evaluation, design, implementation and operating phases of the framework. The assessment of compliance and effectiveness would also require the presence of an internal audit function, the availability of suitably qualified other staff, or a budget allowing such resources to be brought in from outside.

2) Technical knowledge and experience

The framework does assume the availability within the organization of staff with sufficient knowledge and experience to perform the different phases of the development and implementation. As described, the framework requires the critical evaluation not only of technical, information security matters, but also of legal and operational requirements. The risk assessment and control design phases also require specific skill-sets and experience that might not be widely available in smaller organizations

3) Senior management awareness and buy-in

As with all organization-wide strategies and initiatives, the implementation and optimization of effective information security measures can only be achieved with the appropriate support and understanding of senior management. The necessary investment of time and resources, support for the roll-out, and impetus for update and review, require significant awareness on the part of the most influential decision-makers and policy-setters so that the staff assigned to the project have the encouragement and correct environment to perform their tasks effectively.

V. Conclusion

The implementation and optimization of reliable information security measures is a subject that can require a great deal of expertise, energy and resources to perform properly. No one-size-fits-all framework will be appropriate for all organizations, but by following a set of reasonable, standard principles in a structured way, many organizations are able to define and meet their basic requirements in this respect.

ACKNOWLEDGMENT

The author would like to thank his employers, PricewaterhouseCoopers SA, Lausanne, Switzerland, for their ongoing support of his extra-curricular activities undertaken at the University of Lausanne, under the direction of Professor S. Ghernaouti-Hélie.

REFERENCES

- [1] Ordonnance sur les banques et les caisses d'épargne (OB 952.02), paragraph 44
- [2] Loi sur le guichet sécurisé unique (LGSU, 150.40), République et Canton de Neuchâtel, 2004
- [3] PCAOB, "Auditing Standard No. 5 – An Audit of Internal Control over Financial Reporting that is Integrated with an Audit of Financial Statements and Related Independence Rule and Conforming Amendments", May 24 2007
- [4] ITGI, "CobiT 4.1", 2007
- [5] Office of Government Commerce, "ITIL v3", 2007
- [6] ISO/IEC, 2005
- [7] PricewaterhouseCoopers, "Internal Controls: The Key to Accountability", 2005
- [8] ENISA, "A Users' Guide: How To Raise Information Security", 2006

(2) Reasonable Security by Effective Risk Management Practices: From Theory to Practice

Solange GHERNAOUTI-HÉLIE
Faculty of Business and
Economics
University of Lausanne
sgh{@}unil.ch

David SIMMS
Faculty of Business and
Economics
University of Lausanne
David.Simms{@}unil.ch

Igli TASHI
Faculty of Business and
Economics
University of Lausanne
Igli.Tashi{@}unil.ch

Abstract

In this period of grave economic uncertainty, organizations have to manage increasingly complicated situations in an environment that is subject to massive and rapid evolution. A solely intuitive approach to risk management is no longer sufficient when considering the need to optimize investments in relation to security. It is necessary to find the often difficult balance between the cost of risks and their mitigation so as to ensure that organizations can realize their objectives in a reasonable and durable manner.

The aim of this paper is to demonstrate the strong need for information system owners and managers to rely upon an effective risk analysis methodology for decision making related to efficient security measures in order to enhance business performance.

We present a methodological approach. A case study and description of real-life experiences are presented in order to illustrate the applicability of this approach.

Keywords — *Uncertainty, Risk and Security Management, Information Systems and Security Governance, Information systems conformance issues, Sustainable development, Responsibility, Complexity Management, Multidisciplinary approach, Real life case study.*

1. Introduction

In this period of grave economic uncertainty, organizations have to manage increasingly complicated situations in an environment that is subject to massive and rapid evolution. A solely intuitive approach to risk management is no longer sufficient when considering the need to optimize investments concerning security. It is necessary to find the difficult balance between the cost of risks and their mitigation so as to ensure that organizations can realize their objectives in a reasonable and durable manner.

The current, turbulent state of the economic environment is partly due to an absence of appropriate and consistent risk management within businesses that aimed to generate increased profits while either ignoring risks or assuming that risks were under control.

The aim of this paper is to emphasize the role and value of information, the pertinence of decision-taking processes, and the implications of information technologies and good communications, in the overall framework of corporate and organizational governance and in the management of financial risks?

2. Information Systems Governance and Balance of Interests

There are many possible definitions of governance [1], [2], [3], [4], [5]; for the purposes of this paper we are using it to describe a mechanism that allows us to take good decisions in the context of a more or less complex given environment. The tools associated with good governance are thus those that support the ability to take good decisions, in particular by implementing a system of efficient guidance.

The good governance of risk management allows us to master the difficulty of finding a judicious balance between the principles of elementary precautions and taking risks that can be evaluated, mitigated and managed, in an objective and systematic manner.

This approach will permit us to establish a well balanced framework to:

- Ensure continuity;
- Ensure flexibility, allowing rapid reaction to inevitable disturbances;
- Avoid the irreversible;
- Optimize performance;

having taken care to identify the characteristics of the different resources in question and the relative strategic values that have to be preserved.

The amount of protection a resource requires depends on the value that is assigned to it. This value can vary according to the players and its usage, and over time. The effort to be made in analyzing risk degradation depends on the assessed relative values of the resources.

In terms of protecting resources, one can, as an analogy, refer to the declaration from the Rio Summit [6] that reminds us that “it is necessary to limit, frame or avert those actions that are potentially dangerous, without waiting for the danger to be scientifically proven”.

This sets out the foundation of the precautionary principle, which says it is preferable to avoid action when the consequences of that action could be major and irreversible and are impossible to predict with scientific certainty.

The principle of precaution specifically applies therefore to potentially grave but uncertain risks that threaten potentially irreversible consequences, especially when poorly identified. This brings us back to uncertainty, to the quality of the information and the decision-taking process, and to some questions:

- How to take decisions in a context of uncertainty?
- As of when does one consider that the knowledge available is uncertain, or on the contrary sufficiently certain?
- How to apply the precautionary principle? Sometimes this will be a matter of simple common sense, but caution must be exercised as a simplistic, overly cautious approach based on prudence could lead to ignoring alternative economic activities that would increase competitiveness?
- How to identify and describe the probability of a risk occurring, as well as the potential gains if avoided, when our knowledge is uncertain?
- How to evaluate the capacity to cover a risk, in insurance terms, so as to allow its reversal or financial compensation?
- How to distinguish between real dangers and those generated by fear? It can be a challenge to differentiate between real risks, fear, alarmist rumors or intellectual fraud, and for managers to identify reliable tools for taking good decisions, in a context of uncertainty.
- How to build the capability to manage great uncertainties – the inherent difficulty in all prevention exercises?

Even if in practice it is impossible to master all of these uncertainties, it is possible to evaluate their consequences, assuming that the effects of the evolution of the various exogenous (external) and endogenous (internal) parameters on the performance of the enterprise can be modelled.

Such a model must be simple and be based on a few key factors, but always built with the help of efficient governance. And, of course, this governance must not omit the principle of prevention.

3. The principle of prevention and risk opportunity strategies

The principle of prevention applies specifically when the risks can be clearly identified. It thus is a preventive measure, a question of attacking the problem at source rather than a post-hoc corrective measure.

Prevention can only really be implemented if our understanding of the existing mechanisms in place allows an accurate estimation of the damage and if proportional preventive action is possible [7], [8], [9].

It should also be remembered that prevention is encapsulated in the precautionary principle because it defines the totality of the tasks destined to avoid the given threats in the short term, or to limit and reduce the risk of impacts in the longer term. Hence it covers the adoption of effective measures to avert risks of serious and irreversible damages even in the absence of certainty.

Knowledge of the value of the resources is essential because it decides whether or not a preventive action is viable and thus drives all security actions.

The underlying principle is to reduce risks step by step, until the level is reached where

every additional step will cost more than the discounted savings.

The best plan is useless if not acted upon and the feasibility of its implementation is as important as its objectives.

A significant factor that needs to be appropriately considered when addressing questions of security, risk, prevention and precaution is fear. Fear does not facilitate the success of a security plan and although it plays a largely underestimated key role, emotional reactions need to be avoided.

Notably, since the events of 11 September, we have experienced an increase in discussions about fear where risks are concerned, which are echoed in more general fears such as criminality, delinquency, bio-terrorism, and pandemics.

It is difficult to disassociate fear from real danger. It is, however, important to seize the ambivalence of fear, both a primitive emotion and a complex social characteristic. It is implied in a good number of behaviors: flight from real dangers, anxiety in the face of imaginary dangers.

Fear can lead to avoidance strategies which are entirely valid from a security point of view as awareness is raised. Calculation of the risk drives protective and preventive actions but can also, however, generate paralysis of the economic and social development of the enterprise or company.

As a disturbing but also a regulating factor, fear is not limited to a negativity that is largely uncontrollable, it works as a vector that pushes towards the need to find good tools, and helps in taking decisions based on maximum knowledge in spite of incertitude.

It should not be forgotten that in some contexts there exist advantages in orchestrating and instrumentalising fear, so that solutions are selected that are far from meeting the needs of those concerned.

Fear concerns the manager personally, with regard to his or her civil responsibilities as well as financial losses or the corporate loss of image.

On top of this we have to force ourselves to respond rationally and with discrimination. This brings us back to the importance of determining the value of assets, of our comprehension of the environment to be protected and the relationships between the entities of which it is composed that themselves command the dynamics of its operations, and thus the governance of the whole.

There is no universal solution that answers the specifics of each situation; the solution will depend, amongst other things, on:

- the collective mobilization of all the players; and
- the provision of a template that is simple, that integrates and consolidates.

The risk, up to now seen as a negative element to be suffered, tends to become an

integrated element of contemporary society, as has been emphasized by Ulrich Beck in 1986, in his reference work "Society at Risk" [10].

This integration translates itself into an evolution of mentalities and behavior, using the mastering of risks as a requirement for greater security. Modern technologically advanced societies and their citizens no longer accept catastrophes as simple, unavoidable destiny (examples here would be nuclear accidents such as Chernobyl, chemical accidents such as AZF, or the construction of buildings that are not earthquake resistant).

This integration is also a general search for those responsible in times of crisis that ends up in demands for indemnification as soon as losses and damage are identified.

Furthermore, there exists today a common requirement for more and more security, while denying or downplaying the reality of uncertainty and risk. The requirement for maximal security is omnipresent.

Scientific and philosophical thinking of the past centuries allowed the belief that would be possible to attain absolute security, to efface uncertainty and to completely master risk. Today, computer systems, with their inherent complexity have allowed the re-invention of fear, with the distinction that it is not only nature that generates risks or major catastrophes, but also science and technology. We are far from the optimism of the 19th century.

Some technology-linked risks are all the more menacing for being global, due to the interdependence of our infrastructures [11]. They exceed the understanding of any given person or institution and demand a certain amount of cooperation between private and public players at both the local and international level, thus generating even more complexity to the governance of security.

The cultural evolution of the apprehension of risks is visible in its analysis and treatment, highlighting the diversity of our management methods that themselves closely depend on our knowledge and evaluation thereof. Risk also becomes an opportunity for economic development and a competitive factor, not only for the suppliers of information security solutions, but also for example for architects, for urban planners or civil engineering companies who try to counter the effects of a tsunami or potential earthquakes due to tectonic plate movement.

Thus, depending on the players and the framework of their intervention, the representation of risk takes on specific dimensions with socio-economic constructions, with the perception of an opportunity to bring progress, well-being, performance, competitiveness and financial revenue. In brief, the management of opportunity risks.

We know of methodological, deterministic and probabilistic procedures that bring little information concerning the dangers over time of the management of complexity and incertitude. Some methods mix the empirical with expertise and do not necessarily allow good risk management and rarely provide reasonable security due to their immense complexity and the difficulty for management to interpret and use their results. This highlights the challenges of simplicity and the appropriation of expert knowledge for the

corporate players.

Concretely the answer to good risk management is embedded in four main skills. With:

1. Engineering techniques and scientific skills,
2. Legislative and legal skills,
3. Governance and managerial skills,
4. Insurance skills;

one will first try to determine, the amount of necessary risk that would be economically and socially acceptable. That is to say, the amount of risk to which one is prepared to be exposed, given the discounted advantages and the level of confidence in the evaluation and control methods.

Now is the time to mention the courage of some decision makers who take up the challenge of mastering risks and who take responsibility in an economic context where the performance pressures are exacerbated.

4. An analytical approach to risk management to obtain a reasonable level of security.

This approach is not a ready made recipe coming from a checklist, but rather a methodological framework that is sufficiently generalized and simple to be applied to all functions/professions/domains within an enterprise, while sufficiently detailed to allow an efficient analysis of the risks, to adequately determine the measures to be taken to treat them and to promote a culture of addressing risk within the heart of the enterprise.

This methodological framework has to be able to address all risks, whatever their nature (financial, environmental, operational, related to information technology and systems) in a uniform manner, integrated across all functions of the enterprise. The principal goal of such a method is to push organizations towards adopting a risk managing culture [12], [13] and it proposes objective methods to opt for a level of security that is neither too overwhelming nor too weak, to avoid an over investment in security, while ensuring sufficient protection. Thus it contributes to mastering the security budgets that the majority of relevant studies and analysts repeatedly predict will need increasing, without addressing the costs of the reasonably efficient and effective solutions implemented.

4.1 An integrated approach of risk and security aspects

First of all, one can easily deduce that there are two main approaches to be performed in order to reach an assurance level regarding the safety of organizational values: a risk-based approach and a security-based one. These two approaches address the same overriding objective, the safety of the organization's valuable assets, but each one is driven from different mid-level objectives. Conceptually, the risk management process plays an ever-increasing role in security risk strategies [14], as it is the first input to the second stage which is effective Information Security. In order to achieve the goal of effectiveness in the domains

of security and assurance, both approaches have to be integrated.

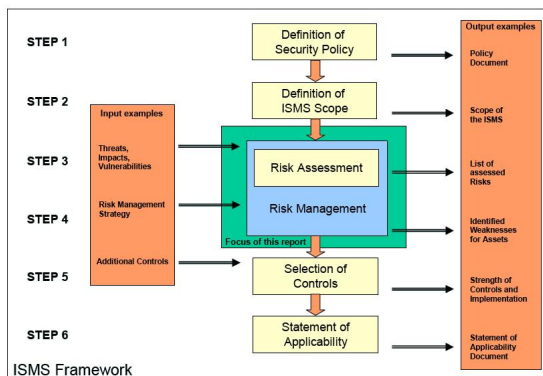


Figure 1: ISMS framework integrating both processes of Risk and Security Management

The aim of a Risk Management process is, generally speaking, to identify, evaluate and mitigate risks by choosing appropriate measures for protection [15], [16], [17].

The Information Security process supports this stage by making sure that the security measures in hand are correctly implemented, and they are not only effective but also efficient [18], [15], [19]. In fact, the Information Security process ensures that the protection measures are implemented in an appropriate and applicable way and that they are consistently and effectively applied. In addition, it is within the scope of the Information Security process to ensure that every organizational, technical and human change is considered and the process is updated regularly.

Risk Management will outline the appropriate security controls to be operated. The security management ensures that the security controls are operated in the most effective and efficient way. Different and complementary competencies are clearly required to fulfill the requirements of both processes.

This method of proceeding will respond to the first requirement regarding the establishment of a reasonable security level. Starting with the Security Policy as shown in Figure 1 means that all security controls, procedures and activities will be driven by organizational objectives, stated within the policy, rather than by compliance standard claims. By the same deductive reasoning, the coherence criterion is also achieved.

4.2 Embracing a governance approach

To move further towards our final objective of reasonable security a governance approach, rather than a standard-based implementation approach is needed for multiple reasons.

Governance implies supervision and controlling operational activities, thereby placing the criterion of improvement in the core of the system. Governance is the process of establishing and maintaining a framework and supporting management structure and

processes to provide assurance that, information security strategies [1]:

- Are aligned with and support business objectives
- Are consistent with applicable laws and regulations through adherence to policies and internal controls
- Provide assignment of responsibility

The term Information Security Governance describes the process of how Information Security (IS) is addressed at an executive level [20].

In the context of the management of information systems and the implementation and maintenance of reasonable information security, governance in our model can be split down into four key objectives.

The first of these is continuity, ensuring that services can recover and continue should a serious incident occur. A key element of continuity is in ensuring that there is no loss of processing, so that should a restoration of services be required they are restored at the same status as before the loss of service – or that users are appropriately informed of the extent of the loss of processing and the actions to be undertaken to remediate the situation.

Secondly, resilience is the ability to provide and maintain an acceptable level of service in the face of errors, interruptions to communications and processing, and deviations from normal operations.

Thirdly, availability is the provision of access to functioning systems. Its definition and calculation in the IT context is often the subject of differing views between the IT and business functions, because a simple calculation of the uptime of a system, for example, may be numerically accurate but fail to reflect the fact that the single period of downtime occurred on the final day of the month, preventing the business from performing a number of key procedures.

Finally, management is concerned with optimizing performance. This can be less technical but a fundamental part of governance, seeking as it does to gain the maximum utility from the resources in place while managing costs and still enforcing appropriate controls over security and system management.

Governing IS security means that a given organization possesses:

- an IS risk management methodology
- an IS strategy
- an effective IS organizational structure
- some IS policies
- a complete set of security standards
- an institutionalized monitoring processes
- some process to ensure continued evaluation and update of security policies, standards, procedures and risks

In that way, the decision making process becomes more transparent, structured, and comprehensive, in order to better respond to the complex and polymorphic characteristics

of risk.

4.3 Reaching an efficient guidance stage

As stated below the efficient guidance stage is function of two variables:

1. Standard precautions, which could be defined as a baseline security level, everybody has to implement.
2. Risk appetite appropriate to a given organization evolving in a given context under given objectives.

Reaching an efficient stage means thus gaining knowledge of the different standards, best practices or current practices in use. This is relation to the first variable but in practice cannot be considered to be sufficient, as it might only allow the organization to achieve the same level of insecurity as others do.

Secondly, in order to define the risk appetite, a deep knowledge of security capabilities is required in order to be sure that we are taking risks we are proficient in managing. It implies a top-down understanding of the existing mechanisms and a whole view of the totality of tasks performed within the organizational perimeter.

This requires a formal analysis system guided by the general vision of the organization regarding opportunities, risks and security countermeasures.

4.4 Reasonable and durable security assurance

Such a formal analysis could be performed in the circumstances where an organization has already deployed a set of well-documented control objectives, controls and policies. This is based upon a specific method imposing a rigorous structure of the processes. The assurance level will be reached when a structured system is built up, when all components of the system are mastered.

The first stage requires a structure regarding Information Security in order to evaluate every function of it. In order to gain assurance in the system, a formal structure, dissecting all the components of the system is needed.

The Information Security Level of a given organization will derive from the performance quality of each one of these activities. Taken together, the measured performance of each activity will allow an overall evaluation of whether the global Information Security Level has been satisfactorily attained.

By considering security as a process, and quality as an inherent feature and as a degree of excellence, thus the quality of Information Security is the degree to which a set of inherent characteristics fulfils requirements. In our case, it will have the look of the PDCA model including:

1. Management responsibility
2. Resource management

3. Product realization
4. Measurement analysis and improvement.

Last, but not least, organizational requirements have to be evaluated in order to assure that the Information Security System in place performs in an efficient (reasonable) way. This could be performed by measuring or assessing the security system maturity, which indicates the extent to which a specific process is defined, managed, measured, controlled and effective; this improves, according to [21], the predictability, the control and the process effectiveness.

5. A case study

A large manufacturing company with operations in more than twenty countries and significant IT centers in ten different locations became aware that it had very little consistency in its approach to risk assessment, security assessment and practical security measures, partly as a result of different infrastructures and application environments in different territories, and partly as a result of the historical ad hoc nature of the development and implementation of policies, procedures and internal controls. It used the opportunity provided by the introduction of the Sarbanes-Oxley Act to re-evaluate its security management from the top down and to implement and reinforce a structured approach to protecting its assets.

The first phase of the project involved a detailed review of the overall corporate risk assessment. By necessity this was driven by the need to evaluate the risks that the group encountered in respect of its financial reporting obligations, but the project was deliberately set a wider scope to incorporate elements such as operational, legal, environmental and reputational risk, and detailed questions relating to the utilisation and management of IT systems and data. The phase was led by a high-level team from corporate finance and compliance and obtained significant input from specialists in the legal and technical fields, with particular attention being paid to local requirements that were not necessarily visible from the corporate centre.

The second phase consisted of the definition of a series of control objectives designed to mitigate the risks previously identified. The objective here was to create a common structure that would be applicable across the whole group, without reference to the individual circumstances and application environments at each data centre.

The third phase involved the definition of specific control activities designed to implement the control objectives previously identified. Again, the aim was to provide as standard a template of controls as possible so that the activities performed in each data centre and by each infrastructure or application ownership group were as consistent as possible, but it was understood that complete consistency would be difficult to enforce and so some latitude was permitted to take account of local circumstances, as long as any deviations were documented, justified, and approved by the centre.

A critical reflection during this phase concerned the cost-effectiveness and the efficiency of the controls being designed. Management wanted to obtain an appropriate level of comfort and assurance over security without incurring excessive costs in the provision of new hardware and software, in training, or in employing new resources to perform the control activities in question. In addition, management wanted to ensure that the procedures implemented would be of appreciable added value to the organization and not tie up valuable resources in what would be seen as purely administrative, form-filling exercises. Accordingly, the control design process involved careful reflection on the nature and coverage of the control activities to ensure that sufficient comfort could be obtained from the effective operation of the smallest possible number of high-level, high quality controls, and preferably those that generated minimal overheads to implement and operate.

The next phase concerned the implementation of the controls, which involved the presentation of the controls matrices and supporting documentation to control owners, targeted training for those staff, and the acceptance of deadlines for the effective implementation of the new or formalized activities.

Some three months after the structured control matrices for security were implemented, management began to evaluate the operational effectiveness of the controls. As a first step, local management was asked to confirm their compliance with the new standards and indicate any areas of difficulty in implementing or operating the controls. Then corporate internal audit performed reviews for each site, looking for evidence of the operation of controls. A final phase saw the external financial auditors review the operation of controls for a sample of sites. The results of all three steps were amalgamated into a reporting package that was discussed and reviewed by senior IT management, corporate management, and the CIO's office.

As a result of these reviews and of the feedback obtained from control owners, a number of modifications were made to the overall structure of controls and to individual control activities, in order to clarify, simplify, or make them more effective and efficient. Senior management also set up a procedure for the annual review of the entire process, starting with the risk assessment and drilling down towards individual controls, in order to ensure that significant risks continued to be addressed and that resources were being focused in the correct directions, all the while complying with the corporation's stated objectives for maintaining an appropriate level of security.

6. Conclusion

The notion of survival brings to mind the long-term preservation of resources and, by analogy, sustainable development.

If we consider the question of sustainable development: what has changed today is man's ability to understand the gravity of the situation and to demonstrate its mechanisms,

combined with the previously missing technical capacity to analyze and understand the elements of sustainable development and their interactions. Sustainable development is based on the preservation of resources of the space capsule that our planet is, and their value for future generations.

The objective is to preserve the resources of our enterprise, to make them multiply, to appreciate their value relative to their potential loss and to not sacrifice the criteria of economic efficiency to the hurdles of short-term profitability.

We should not economize on resources for risk management and the establishment of efficient and sustainable security solutions as these are critical steps when addressing problems with disastrous consequences.

This questioning of the enterprise's sustainable development leads us back to considerations concerning the notions of good governance.

7. References

- [1] P. Bowen, J. Hash, and M. Wilson, Information Security Handbook: A Guide for Managers (NIST Special Publications 800-100) National Institute of Standards and Technology, 2006. [Online] Available at <http://csrc.nist.gov/publications/nistpubs/800-100/SP800-100-Mar07-2007.pdf>
- [2] C. P. Grobler, "A Model to assess the Information Security Status of an organization with special reference to the policy Dimension", M Phil (IT) Thesis, Computer Science Department, University of Johannesburg, 164, 2003.
- [3] IFAC, Enterprise Governance: Getting the Balance Right International Federations of Accountants, Professional Accountants in Business Committee (PAIB), 2004. [Online] Available at <http://www.ifac.org/MediaCenter/files/EnterpriseGovernance.pdf>
- [4] ITGI, Information Security Governance: Guidance for boards of Directors and Executive Management, 2nd Edition IT Governance Institute, 2006. [Online] Available at http://www.itgi.org/template_ITGI.cfm?template=/ContentManagement/ContentDisplay.cfm&ContentID=24384
- [5] ISF-std. The standard of Good Practice for Information Security Information Security Forum, 2007.
- [6] W. R. C. Latin_America_and_the_Caribbean, "RIO DE JANEIRO COMMITMENT: REGIONAL PREPARATORY MINISTERIAL CONFERENCE OF LATIN AMERICA AND THE CARIBBEAN FOR THE SECOND PHASE OF THE WORLD SUMMIT ON THE INFORMATION SOCIETY " 2005. [Online] Available at <http://www.itu.int/wsis/docs2/regional/declaration-rio.pdf>
- [7] T. Peltier, "Risk Analysis and Risk Management " Information Systems Security vol. 13 (4), pp. 44-56, 2004.
- [8] IEEE-Std. 1700, IEEE P 1700: Information System Security Assurance Architecture (ISSAA) Standard, IEEE, USA 2008.
- [9] AIRMIC, ALARM, and IRM, "A Risk Management Standard," The Institute of Risk Management, The National Forum for risk Management in the Public Sector, The Association of Insurance and Risk Managers, London, UK 2002. [Online] Available at http://www.theirm.org/publications/documents/Risk_Management_Standard_030820.pdf
- [10] U. Beck, La société du risque : Sur la voie d'une autre modernité. Paris, France: Flammarion, 2008.
- [11] F. Cohen, IT security governance Guidebook with security program metrics. Boston, USA: Auerbach Publications, 2006.
- [12] B. v. Solms, "Information Security - The Third Wave," Computers & Security, vol. 19 (7), pp. 615-620, 2000.
- [13] K.-L. Thomson and R. v. Solms, "Towards an Information Security Competence Maturity Model," Computer Fraud & Security, vol. 2006 (5), pp. 11-15, 2006.
- [14] M. Johnson and J. Spivey, "ERM and the Security Profession," Risk Management, vol. 55 (1), pp. 30-35,

- 2008.
- [15] ISO-Std. ISO/IEC TR 13335-1, Information Technology - Guidelines for the management of IT Security - Concepts and models for IT Security, International Organization for Standardization (ISO), Switzerland, 1996.
 - [16] ISO-Std. ISO/IEC 27005:2008, Information technology - Security techniques - Information Security Risk Management, International Organization for Standardization (ISO), Switzerland, 2008.
 - [17] ISO/TC-Std. 31000:2008, Risk Management - Principles and guidelines on implementation (draft), International Organization for Standardization (ISO), Switzerland, 2008.
 - [18] ENISA, "A Users' Guide: How to Raise Information Security Awareness," European Network and information Security Agency, Heraklion, Greece 2006. [Online] Available at http://www.enisa.europa.eu/doc/pdf/deliverables/enisa_a_users_guide_how_to_raise_IS_awareness.pdf
 - [19] ISO-Std. ISO/IEC 15408:2005, Information technology — Security techniques — Evaluation criteria for IT security — Part 1: Introduction and general model, International Organization for Standardization (ISO), Switzerland, 2006.
 - [20] S. Posthumus and R. v. Solms, "A framework for the governance of information security " Computers & Security, vol. 23 (1), pp. 638-646, 2004.
 - [21] ISSEA. Systems Security Engineering Capability Maturity Model (SSE-CMM), International Systems Security Engineering Association (ISSEA), 2003.

(3) Protecting information in a connected world: A question of security and of confidence in security

Solange Ghernaouti-Hélie, David Simms, Igli Tashi,
Faculty of Business and Economics, University of Lausanne

Abstract—The infrastructures and services related to information and telecommunications are crucial constitutive elements of our society. These elements require not only ICT security but also confidence in that security.

This paper explores at a macroscopic and integrated level the main challenges, obstacles and constitutive elements that contribute to building confidence in information security.

The aims of this paper are to identify some key technological elements and impacts that drive the development of information security from a long-term risk management perspective; point out the complexities in applying security to vulnerable ICT environments; clarify the need to master indicators and methodologies that contribute to identifying and applying good risk and security management practices and confidence in these; show how legislation can contribute to limiting or enhancing information security effectiveness; analyse the relationships, differences and interactions between security, confidence and compliance; and discuss how the audit process and the input of auditors can help in building confidence.

Index Terms—Networked society, information security, risk management, security doctrine, confidence, compliance, audit, complexity, vulnerability, dependency

I. Introduction – networked society at a global scale and technocivilisation

Nowadays it is impossible to address the questions of security and of confidence in information security that concern us without first considering the global context in which these questions arise. This context is the information society.

The protection of information in an interconnected digital world is an issue for all of us because we have only relatively recently stumbled into an information society. We are surrounded by technology from our births to our deaths. Digital technologies are a major factor in our civilisation and, whether we like it or not, information technologies and the Internet have brought us firmly into the Era of Technocivilisation.

IT is everywhere, increasingly capable, increasingly invisible, but at the same time increasingly intrusive. Telecommunications networks, using and including mobile telephone technologies, are allowing increasing access to communications, all the time, everywhere, with everyone and also with anyone. There exists a whole constellation of networks and relation-forming services. Our society is a networked society on a global scale in which micro- and macro-communities, both temporary and permanent, can exist. Telecommunications allow the constant definition of links that are formed and broken as a result of the contacts between people and machines and of the exchanges between entities that may, or may not, know each other. A dynamic with an impact on the lives of everybody is created by these interactions.

These links connect people to each other and to their environment. They allow people to live and to

give meaning to existence. They reflect the abilities to associate, to interact, to enter into communication, to build, to share, to shape the present. These links are a defence against isolation and also the driver of economic and social development in our society.

Everything is a link and everything is a story of links. Whether they are visible or invisible, these links exist and carry a certain value. As an example of this, the “Facebook” social network has based its business model on the exploitation of these links.

Confidence can be seen as a chain of trust, which, as in all chains, is only as viable and reliable as its weakest link. System users, system administrators and system evaluators are all confronted with the traditional problems of being able to identify the individual links and assess their reliability.

II. A global approach to a complex problem

A. People and technology

The French philosopher and sociologist Edgar Morin noted that the inability to see the whole picture and to connect to everything leads towards a lack of solidarity and to irresponsibility [1].

There is therefore a high risk that people will be restricted to being docile consumers in both their private and professional lives. That is also true when dealing with security solutions issues. If so, it is no longer the case that people are using cyberspace; it is cyberspace that is using people.

Technology is not neutral. This is not a consequence of the good or bad uses to which it is put, but rather of its capacity to remove power from some users, to give power to others, and to change reality for all. [2]

This is exactly what information security solutions do.

Some suppliers provide security in obscurity, some actors create, master and control security solutions, and some others use them without any real understanding or expertise.

B. Vulnerabilities, dependencies and risks

Let us discuss briefly the foundations of the networked society on a global scale: a framework that can be summarised in three words: vulnerabilities, dependencies and risks. ICT infrastructures are widely interconnected and interdependent and use the same Internet technologies to create a uniform digital world. As a result of its nature and its complexity, this world is fragile.

There exist numerous vulnerabilities, among which we can mention technical, organisational and managerial issues. Technology and its implementations can be subject to natural catastrophes (such as earthquakes), to accidental human errors, or to malicious actions (deliberate errors or criminal activities).

Attacks on information technologies can have a considerable ability to cause damage and can have dramatic consequences for the organisations and individuals who fall victim to them. Of course, this is also true for all ICT incidents regardless of their origins.

Hunting down hackers and the creators of malware, like the implementation of technological barriers, is not sufficient to provide effective protection against attacks on ICT installations. The annual reports of groups monitoring ICT disasters continue to bear witness to the increasing number and complexity of IT attacks, as do the events that regularly become newsworthy.

The dependency of modern society and infrastructures on ICT is already immense and is continuing to grow. Even forty years ago the only connection that the average citizen had with information technologies was perhaps in their dealing with large banks or areas of the public administrations that had begun to implement cumbersome, back-office mainframe systems to perform some data processing and storage tasks. Nowadays, information technologies permeate all aspects of life and even those citizens who do not actively utilize these technologies themselves (for example, in choosing or being unable to use a PC or a mobile phone) are still fundamentally dependent on the technologies that have been implemented by the businesses, organisations and administrations with which they interact.

As a consequence of this, all citizens are confronted by the risks related to the use of information technologies. These risks apply to the integrity, confidentiality and availability of the systems and the data stored and manipulated within them, and it is hard to envisage a realistic modern life (apart from perhaps a completely secluded agrarian lifestyle) in which it would be possible for an individual to avoid exposure to and reliance upon these systems. Financial transactions, travel organization, standard administrative activities such as taxation and car registration, all involve the inclusion of data and the recording of transactions in large, multi-user systems.

Organisations thus face a number of threats of varying severity and the risks of breakdown, of deterioration and of financial or reputational losses are genuine. This situation can impact the performance, competitiveness and profitability of the organisation and on occasion threaten its viability. As a simple matter of fact, the systematic use of ICT will introduce a non-negligible operational risk.

III. Protecting and defending ICT infrastructures: a question of risk management

There is therefore a need to know how to protect and defend the infrastructure and all sensitive information and information flows, a requirement that can become extremely difficult because of the great potential within ICT to monitor, steal, publish or destroy data.

Mastering risks has always been a strategic element in corporate competitiveness. So what has changed?

The goal of risk management has always been related to the necessity of finding the difficult balance between the cost of risks and their mitigation so as to ensure that organizations can realize their objectives in a reasonable and durable manner.

We have witnessed a cultural change in the perception and the understanding of risks over recent decades.

Risk, when seen in the past as a curse or a calamity simply to be suffered, tends to become a element integrated into contemporary societies, as the German sociologist Ulrich Beck pointed out in his 1986 book "Risikogesellschaft" ("Risk Society") [3]. According to Beck, we are passing from an industrial society, in which the central problem concerns the sharing of riches, to a society based on the sharing of risks.

To put it another way, risk is no longer an external threat, but rather a constituent part of society.

Today it is not the size of the risk that is changing, but its "scientification", which means that it is no longer possible to discharge responsibilities by blaming nature.

This shift in the perception of risk has shown itself in the evolution of mentalities and behaviours. This has brought an understanding of the need for reliable risk assessment methodologies and for confidence in the people and partners involved in the evaluation of risks.

The subject of mastering risks is a requirement of greater security. It is also fundamental to the general search for responsibilities (or for the guilty parties) in crisis situations and finally the requirement for compensation when there are reports of losses or damages.

In the future, there will be a strong demand for more and more security, for the assurance of security, while at the same time there will be a tendency for the realities of uncertainty and risk to be increasingly ignored and even denied.

Nowadays information systems and their complexity have allowed the reinvention of fear. It is not only nature that can cause risks and major disasters, but also science and technology. We can see it, for example, in Japan with the nuclear problems resulting from the earthquakes and subsequent tsunamis (March 2011).

Computers and telecommunication networks are also the basis of risks. To cite two examples there have been the cases of the Stuxnet malware, detected in July 2010, that affected Supervisory Control And Data Acquisition (SCADA) systems in nuclear installations in Iran, Indonesia, India and Pakistan, and the cases of the theft of company sensitive data that have occurred in several banking institutions in recent years.

Certain risks deriving from technology are even more threatening when viewed on a global scale because of the interdependence of infrastructures and of their global impact, creating even more complexity in security governance. Organisations are being required to manage increasingly complicated situations in an environment that is constantly evolving.

If certain organisations are fairly aware of their real needs in respect of information security, others are more sensitive about talking of fear, the fear of insecurity, which is regrettable. Because even if fear can contribute to selling security solutions, it is not on its own sufficient for addressing appropriately security requirements which, most of the time, are badly identified and documented. This process sometimes results in costly, wasted investments.

For businesses, the first activity related to risk management is often to identify the level of risk that would be economically and socially acceptable and that would allow management to ensure business continuity and avoid the irreversible. This is the risk to which management accepts exposure, based on the perceived advantages and their level of confidence in their means of assessment and control.

Good governance of risk management allows the identification of a judicious balance between the principles of elementary precautions and taking risks that can be evaluated, mitigated and managed, in an objective and systematic manner. It will also allow management to ensure flexibility, allowing rapid reaction to inevitable disturbances and the optimization of the enterprise's performance.

A purely intuitive approach to risk management is no longer sufficient when trying to optimise security-related investments. There is a difficult balance to be found between the cost of the risk, of its reduction, of reasonable and long-lasting comfort, and the achievement of the organisation's objectives. Decision-takers should rely upon indicators in which they should have confidence, on reliable information from within and from outside their organisation [4].

A useful example of the requirement for, and use of, reliable information in this context could be found in the situation of a large corporation that is trying to quantify its risks related to data security. One means of doing this would be to analyse the logs from various parts of its IT infrastructure, such as firewalls, networks, databases and applications, to try to evaluate the number, type, frequency and impact of attacks on those infrastructure elements.

Clearly, in order to be able to undertake this exercise with any assurance that the results would be valid and meaningful, management would first need to ensure that the logs were complete and valid. This means verifying that the systems have been configured to record all relevant and significant items, that the logs have not been tampered with in any way, and that the logs are retained in an accessible format until they can be appropriately analysed.

Furthermore, management will need to be confident about the quality, consistency and timeliness of the analysis of the logs. It will be essential to have qualified and experienced staff performing the analysis in a structured and coherent way, knowing which items to investigate further and how to perform and document such investigation.

The organization will of course also be relying on the quality and consistency of information received from outside. Knowing which activities to highlight as potentially dangerous will depend on receiving timely and accurate information, analyses and warning from software and infrastructure vendors, security analysts and other third parties. If potential weaknesses have not been identified, assessed and communicated, the organization setting out to review its potential exposure is already at a disadvantage.

At the end of the process, only on the basis of such well-founded and well-designed analysis will the final results reported to senior management be sufficiently reliable to allow informed and rational decisions to be taken.

IV. A question of confidence

The problem raised in this discussion is that of confidence in the people and means to be applied to managing ICT security risks. This problem also applies, however, to confidence in all organisations, partners and tools involved in this field. In a connected world, dependencies and confidence are strongly linked.

By definition, the staff responsible for information systems and security, as well as users with significant responsibilities, require high levels of access to potentially critical and sensitive ICT resources. Occupying such positions requires a high level of personal integrity on the part of the individual staff members, but their organisations are also obliged to implement and operate strict monitoring of, and detailed controls over, their activities, in appropriate proportion to the risks that their access rights pose to the systems that they manage or use. The growth in criminal cases with internal origins that point towards the complicity of staff should motivate organisations to consider this question very carefully and not let themselves be exploited by unscrupulous people.

Information security is not the result of the juxtaposition of security solutions but should be driven by a political willingness, and by a process of continuous management based on an adequate organisational structure, on specific human skills, and on the commitment of all staff that are involved. [5]

As is the case for information security, confidence should be created, maintained and adjusted in an ongoing and appropriate process.

At this time, there are no absolute means of managing confidence, merely internal control procedures that can be implemented to contribute to the process. Confidence, like information security, is never definitively achieved.

One of the constitutive elements of trust seems to be the knowledge the organization has about their values. It has been mentioned that organizations should be aware of their real needs and one of the approaches is to have a deep knowledge of the value stored in data. The need for effective data classification is widely understood but is little performed, for reasons that include lack of resources, lack of prioritization, and the simple difficulty of defining and applying models that are sufficiently robust, that cover the organisation's requirements and provide flexibility without being too unwieldy [6]. Another common failing in such projects arises from the interactivity between systems and data elements, potentially giving rise to situations in which individual data elements may not be seen as being particularly significant but when combined with others may then become important from, for example, a disclosure perspective. Developing and applying models that can take such situations into account can be a complex and costly task, the success of which cannot be easily measured.

In a culture of increased sharing, it is also necessary to ensure that critical information is not communicated through new vectors such as social networks. Traditional security models have concentrated on traditional means of controlling information flow, originally on paper, then on bulky supports, and more recently in electronic form using widely-available and easily transportable supports such as USB devices or via electronic mail, but increasingly attention is being paid to the confidentiality implications of the use of networking sites such as Facebook.

V. The doctrine of confidence with regard to security and compliance

The whole art of security is based on the means of finding the optimal solution that meets control requirements and provides a pragmatic answer to the question of who watches the watchmen.

It is no longer a question of solely addressing the classical criteria for IT security – availability, integrity, and confidentiality of data – but also of meeting the control requirements set out in legislation and the trust requirement required by the market.

Legislation has thus become an endogenous factor to consider in relation to security and is both a driver and a constraint.

A genuine security doctrine should be developed that will ensure security, economic intelligence, monitoring and legal compliance, relying upon an integrated and multidisciplinary holistic approach, which will contribute to delivering confidence to the end-user.

Confidence, compliance and security are three related concepts but far from synonymous, although there is a tendency to confuse them in some business contexts. If we take the example of compliance with legislation such as the Sarbanes-Oxley Act of 2002 [7], we see that as part of the standard reporting package, management and the external auditors will confirm under Section 404 the existence of an appropriately scoped, designed and functioning system of internal controls. This system of internal controls will by definition include controls over and around the information systems, and these controls will include measures designed to ensure the integrity, accuracy, availability and confidentiality of the data and applications that constitute those information systems. Thus compliance has been certified and users of the published accounts will have a tendency to have confidence that the information systems are secure.

This confidence, however, is not necessarily justified. Firstly, audit reports are always limited in scope and in the amount of assurance that they can provide. Best practice and auditing standards ensure that reasonable assurance can be drawn from the validation of internal controls, but the limitations of testing mean that absolute assurance over the operation of controls cannot be obtained [8]. Secondly, the evaluation is made on the basis of the information available at a moment in time. Even if an organization is extremely well controlled and is diligent at evaluating, testing and applying security-related patches to its systems in a reasonable timeframe, it is possible that an unpublished and therefore unpatched security weakness is present in its systems and is being exploited without anybody being aware of this. Thirdly, internal control systems are always subject to being circumvented by human actions, be they accidental or malicious. As soon as users have access to systems and data, there will be a risk of data leakage, and preventing data leaving organizations on electronic supports such as USB devices or handheld devices with data storage capacities is very difficult in environments in which users have the habit of performing data transfers as part of their daily functions. And fourthly, even a clean report on internal controls, including those around security, does not signify that there are no concerns around the management or operation of security. All the report states is that no significant weaknesses, using “significant” as meaning potentially having an impact on the financial statements, the overall control environment or the overall audit opinion, had been identified [9].

As a result of these factors, it is clear that even a formal statement of compliance and an implied confirmation of the existence of effective security measures cannot automatically be used as a basis for

well-founded confidence in the quality, appropriateness and consistency of the security measures in place.

VI. Auditing and auditors

In this section we will be concentrating on the role of the external auditors of organisations, and particularly those charged with reviewing the financial statements. Many organisations will also employ an internal audit function, either in-house or outsourced, and will also mandate external groups to perform specific reviews, often of technical and security-related areas. The findings of such reviews are typically confidential and retained within the organization, however, and play no role in determining the level of confidence in the organization of the public or of partners, and so these activities will not be discussed in this paper.

The role of the external auditors has been under constant discussion over recent years, partly as a result of the financial scandals which provided a great deal of the motivation for the Sarbanes-Oxley legislation and partly as a result of increasing concern over the quality and extent of audit work and the reliance that is placed on audit reports. The language and structure of audit reports have evolved over the years to clarify both the limitations of the work and the restricted audience for the reports [10], but even when those caveats are taken into account, legitimate users of the reports continue to have concerns about the completeness and reliability of the information provided.

In the previous section we discussed the basic limitations of audit work and the conclusions that can be drawn from that work. Addressing these issues will mean a change in philosophy, in scope, in procedures and in reporting on the part of the auditors.

The change in philosophy is fundamental and would bring about significant changes to the whole auditing approach. Typically the approach is very defensive: auditors are required to reach a conclusion or series of conclusions and issue an opinion or opinions. The work performed is designed to support the opinion and to show that sufficient consideration has been given and effort made to identifying errors, exceptions and causes for concern. The objectives of corporate management are similar: typically the desire is for a clean audit report obtained with the minimum of fuss and hopefully at the lowest possible cost. There is therefore very little obvious motivation on either side to go further into analyzing and evaluating wider issues of business design, efficiency and reliability.

The question of scope is fundamental to the audit process. The auditors need a clear definition of which aspects of the business they are looking at and in what detail. As a consequence of persistent cost pressures and strict reliance on the reporting and statutory requirements for financial auditing, the scope will always be reduced to the bare minimum with which the corporation (which is paying the fees) and the auditor (which has to uphold professional standards) will mutually be satisfied. Thus the focus will typically be placed on the areas of high financial significance, of perceived high risk, and of specific interest to regulators. Unless specifically requested otherwise, there will be no detailed evaluation of, or formal opinion on, the effectiveness of compliance with standards apart from those used for financial reporting. [11]

The choice of auditing procedures is in reality fairly limited, with set ways of performing standard tasks being widespread. There is a not entirely undeserved tendency to view the auditing process as essentially a box-ticking exercise performed under time constraints and with little scope or opportunity for creativity.

The reporting process is also very stylized, with standard formats and wordings developed over the years and carefully monitored to ensure that no undue risk is being taken in making any individual statement in the reporting package. This is entirely legitimate given the stakes in play for the auditors and corporate management, but one consequence is that reporting has become a sterile exercise in which only standard phrases are employed. It is hard for the users of reports to gain any specific insight into the real workings of the business under review.

The result of these factors is that it is not necessarily appropriate to derive much confidence from the reports that are presented. There is a comforting reassurance presented by the standardised texts and familiar layouts and in the overall opinion presented by the signatories, but users cannot go further and derive any real confidence over questions of security.

This could be addressed by addressing the issues discussed above. A change of philosophy, a widening of scope, a broadening of testing approaches and a greater flexibility in reporting could combine to provide users of the reports with a reliable, solid basis for assessing security and compliance and thus deriving confidence in the business. This would require a sea-change in the whole context for auditing, however: larger budgets, different timeframes for completing work and reporting, perhaps changes in professional standards and in the protection for auditors. It could very well be that such changes could only be forced through by legislation, which would then have to be coordinated internationally to avoid organisations shopping for the most compliant reporting jurisdictions and relocating to convenient territories, at least in name.

VII. Conclusion

Technology is not neutral and neither is the law; and so we should work towards ensuring that their development goes alongside that of society and becomes a driver for the economy. This is of particular importance in areas such as auditing, as discussed above, where the structures and standards in place continue to reflect the concerns of previous generations and do not necessarily correspond to the requirements of the modern society.

Every society is based on a vision of the world that it has created from a system of beliefs and values that its members share and that gives them their identity, their motivation and their sense.

Perhaps we should take time to respond, in a perspective of durable development and of knowledge. What are these shared values? Could it be confidence and fair practices?

Perhaps we need to dare to change our perspective on security while there is still time. We also need the wisdom to see beyond current limits, a thought process that properly considers the realities of information and security technologies for everyone. Because information security is about secrecy and because the human being who can no longer keep secrets loses some of his substance and a little of his humanity.

REFERENCES

- [1] E Morin, *Pour une politique de civilisation*, Editions Arléa, 2002.
- [2] R Berger, S Ghernaouti-Hélie, *Technocivilisation, Pour une philosophie du numérique*, Collection Focus Science, PPUR 2010.
- [3] U Beck, *Risikogesellschaft. Auf dem Weg in eine andere Moderne*, Suhrkamp, Frankfurt a.M. 1986. ISBN 3-518-13326-8.
- [4] S. Ghernaouti-Hélie, D. Simms, I. Tashi, "Optimizing security efficiency through effective risk management", presented at 25th IEEE International Conference on Advanced Information Networking and Applications (AINA 2011), Singapore, 2011.
- [5] S. Ghernaouti-Hélie, D. Simms, I. Tashi, "Reasonable Security by Effective Risk Management Practices: From Theory to Practice", presented at International Conference on Network-Based Information Systems, 2009 (NBIS '09), Indianapolis.
- [6] I Tashi, S Ghernaouti-Hélie, *Information Security Evaluation: A Holistic Approach*, EPFL Press, 2011.
- [7] <http://www.gpo.gov/fdsys/pkg/PLAW-107publ204/content-detail.html>
- [8] V Ainapure, M Ainapure, *Auditing and Assurance (2nd Edition)*, PHI Learning Private Ltd, New Delhi, 2009, pp48-49.
- [9] I Gray, S Manson, *The Audit Process (Fourth Edition)*, Thomson Learning, 2008, Chapter 16.
- [10] P Holgate, *Accounting Principles for Lawyers*, Cambridge University Press, 2006, p61.
- [11] A Chambers, G Rand, *Operational Auditing Handbook: Auditing Business and IT Processes*, John Wiley & Sons Ltd 2010, p167.

(4) Optimizing security efficiency through effective risk management

Ghernaouti-Hélie S., Tashi I., Simms D, *Faculty of Business and Economics, University of Lausanne, Switzerland*

Abstract—Security measures taken in isolation and without reference to a concrete and relevant assessment and evaluation of actual risks are doomed to be inefficient. At best they do not address the real issues facing an organization and simply waste resources; at worst they provide management with inappropriate comfort over the level of security management that is in place. This paper reviews the key points of some relevant international standards, discusses the links between effective risk management and optimized security measures, and provides a case study illustrating the benefits to be obtained from a structured and integrated approach.

Index Terms— information security governance, risk analysis, security management, security evaluation

I. Effective Risk Management in the Domain of Corporate Information Security

Information security, and security in general, includes all the actions to be undertaken preventatively in order to allow organisations to minimise risk. The outputs, however, are not necessarily identified immediately or even, in some cases, at all. Very often it is a significant problem for corporate managers to justify security spending because of the absence of visible results.

In practice risk management, largely as a result of pressures in the financial sector, is taking on an increasingly important role in business, as inherent risks need to be addressed immediately in order to mitigate any negative impacts. Moreover, international and national regulations are increasingly specifying risk management and related activities as mandatory tasks.

The ultimate aim of risk management, together with security management processes, should be to define what constitutes “reasonable protection” for the organization’s assets, characterized by a number of best practices that aim to secure those assets in the context of specific identified risks.

A. A Definition of Risk Management

One concept of management is that it is the process of the definition and achievement of goals while optimizing the use of resources. Risk management is the process that allows business managers to control the operational and economic costs of protective measures and to achieve improvements in performance by protecting the business processes that support the business objectives or mission of the enterprise. A risk management process can also be seen as a framework for determining and implementing acceptable security controls.

Risk management is a process whereby organizations methodically address the risks attached to their activities with the goal of achieving sustained benefit both within each activity and across the portfolio of activities. This means that the identification of the activities that provide an added value for the enterprise and their prioritization is fundamental. At a second stage, after having identified the important processes and activities, the risks concerning these activities are identified and prioritized.

In some models both processes, security and risk management, are combined into the larger concept of Security Risk Management, which is viewed as being the practice of controlling and mitigating the amount of loss an organization will suffer because of an adverse action or situation, initiated intentionally or unintentionally.

B. Presentation of the Risk Management Process

1) Threats and Vulnerabilities

Regardless of the application domain, the two main components relating to risk management are threats and vulnerabilities. These are the cornerstones of a risk management approach as risk is defined as the potential that a given threat will exploit vulnerabilities of an asset or group of assets to cause loss or damage.

A threat is a potential cause of an unwanted incident that may result in harm for an organization. A vulnerability is a weakness that is susceptible to being used by a threat. Vulnerabilities can be human failings, weaknesses or flaws in technology, or, by extension, anything that does not conform to the expected state of operations. The pairs of threats and vulnerabilities lead to unwanted events, the likelihood of which needs to be estimated or measured. This likelihood is the probability that a vulnerability will be exploited by a threat that leads to harm.

2) The Processes within Risk Management

ISO /IEC TR 13335-1:1996 [1] defines risk management as the entire process of identifying, controlling, and eliminating or minimizing uncertain events that may affect information systems.

According to ISO /IEC 27005 [2], Risk Management is composed of six processes:

- 1- Risk Communication;
- 2 - System Characterization or Context Establishment;
- 3 - Risk Assessment, split into Analysis and Evaluation;
- 4 - Risk Handling;
- 5 - Risk Acceptance;
- 6 - Risk Monitoring and Review.

Risk Communication concerns the achievement of agreement on how to manage risks by exchanging and sharing information about risk between the decision makers and other stakeholders.

The Context Establishment is the definition of the boundaries and scope of the risk management process.

The Risk Assessment process concerns the identification, description and prioritization of risks harming organizational assets. The process consists of two main categories, risk analysis and risk evaluation. The risk analysis itself is made up of risk identification and the risk estimation. This phase requires a deep knowledge of the organization and its environment, as well as a detailed understanding of strategic and operational objectives. The identification and categorization of business objectives, and of the information assets supporting those objectives, are thus required.

Risk identification consists of the identification of critical assets to be risk managed and of relevant threats and vulnerabilities, in order to manage the consequences of the exploitation of a risk.

To follow the framework set out in the ISF Standard of Good Practice for Information Security (ISF) [3], the risk analysis process should include the following steps:

- Calculation of the potential level of business impact;

- Identification of threats, intentional and non-intentional;
- Identification of vulnerabilities resulting from both control weakness and circumstances.

Performance of risk analysis is considered evidence of appropriate diligence in overall risk management, which is frequently an important objective for corporate management.

During the previous stages the most critical assets will have been identified and the risks related to these assets will have been identified and estimated. The risk evaluation phase consists of prioritizing the most probable and severe risks according to the evaluation criteria decided during the establishment of context. The risks identified as important during this phase will be subject to specific measures.

Risk handling, or risk mitigation, is the activity that aims to implement the different options to address risks. There are four (not mutually exclusive) options for handling risk:

1. Reduce through the implementation of controls;
2. Retain or Accept without further actions;
3. Avoid;
4. Transfer to, or share with, a third party.

During this phase different options for addressing the risks are considered and the degree of residual risk is anticipated. Once the assessed risk is considered as acceptable, no additional actions need to be undertaken. If the risk based on the current assessment is considered to be unacceptable, some security controls will need to be chosen in order to reduce or mitigate the risk level. Of course, despite the implementation of security controls, a small element of the risk will always remain; this is the residual risk.

The risk environment is very dynamic and its attributes change continuously. As a consequence, a successful risk management process includes continuous review and development.

II. Security Management

A. Security as a management process

A good risk analysis process should provide a good panoramic view of the implemented security controls. ISO 27001 Information Security Management Systems – Requirements [4] specifies that the controls implemented in an Information Security Management System (ISMS) scope should be risk based. From that point of view, the risk management process could be considered as an input into the Information Security Management Process.

According to ISO/IEC 27001, Information Security Management is a system, part of the overall management system, that is based on a business risk approach and that seeks to establish, implement, operate, monitor, review, maintain, and improve information security. This includes the establishment of an organizational structure and procedures, the definition of the information security policies and responsibilities with respect to information security in general and the policy itself, the development of planning activities, and the identification of the processes and resources needed to operate the information security programme.

ISO/IEC TR 13335-1 defines security management as a process to achieve and maintain appropriate levels of confidentiality, integrity, availability, accountability, authenticity and reliability. At the same time, the maturity model ISM3 defines information security management as a system to prevent and mitigate the attacks, errors and accidents that can jeopardize the security of information systems and the organizational processes supported by these systems.

Generally, information security is considered as a process aiming to implement appropriate measures in order to eliminate the impact that various security-related threats and vulnerabilities might have on

an organization.

The ISF defines information security as a process for keeping the risk associated with information systems under control by means of clear direction from the top levels of management, the allocation of resources, security awareness, and the establishment of a secure environment. The involvement of top management's should provide evidence that security controls have been implemented across the enterprise.

Information security concerns the protection of information assets. This suggests that there are two main categories of actions to be undertaken:

1. The identification of the subjects for protection, for example against risks and losses;
2. The definition of protection strategies.

Risk management as a process must be included in an information security programme. From a pragmatic point of view information security management allows management to ensure business continuity, minimize damage, and organize security activities in a cost effective manner. Information security management ensures that appropriate policies, procedures, standards, and guidelines are implemented to provide a proper balance of security controls and business objectives.

Security management is the framework that ensures that risks are identified and an adequate control environment is established to mitigate those risks. Information security issues cover security policy, risk analysis, risk management, contingency planning and disaster recovery.

B. Security management activities

Managing security means planning, organizing, commanding, coordinating and controlling. Information security management is thus a system used to establish and maintain a secure environment by defining, achieving and maintaining the appropriate safeguards to counter the exploitation of risks. It should be a collaborative effort that utilizes a broad array of organizational capabilities in order to be successful. There is a consensus about the drivers for structured, formalized and managed information security management processes: the most important are technical and environmental complexity, as well as the increasing dependence on technology in doing business. Information security management is a process that is focused on the organization, while risk management contains variables that are outside the organization's boundaries.

Security control activities are designed during the risk handling phase of the risk management process as a means of mitigating impacts. Risk assessment, another activity within the risk management process, is a source of security requirements beyond those arising from legal, regulatory, statutory, operational and business related requirements that will also need to be address by appropriate measures.

According to ISO 13335-1, "Guidelines for the Management of IT Security", eight processes contribute to Information Technology security management.

The first is configuration management, requiring knowledge of changes that occurred within the system in order to evaluate if the safeguards in place are appropriate.

The second is change management, the logical consequence of the first, requiring the identification of new security requirements as a result of changes to the system.

The third and the fourth concern risk-related processes, those of risk management and risk analysis. ISO 13335 considers risk management as the process for implementing the security strategy and security policy, where safeguards are selected and thus the level of residual risk is defined. By way of contrast, risk analysis concerns the identification of risks, that is to say the identification and the analysis of threats, vulnerabilities, and impacts.

The fifth is accountability, dealing with the assignment and knowledge of responsibilities. These responsibilities include both the identification and valuation of assets and the management of the protection process.

The sixth process is security awareness. This is a process involving personnel and their adherence to the information security programme, practices and operations, since individuals are widely considered to be the weakest link in relation to security. This process consists of a number of distinct activities such as the provision of awareness materials, training, and educational activities.

The seventh process is information security monitoring, focused on the appropriate functioning of control activities. This activity is similar in concept to risk monitoring but is focused on the control activities rather than any changes in risk factors. The aim is to be sure that the controls are in place and are being operated in the expected manner.

The last processes considered in ISO 13335-2 [5] are contingency planning and disaster recovery. These processes document how business operations will continue in the event of an adverse occurrence and the unavailability of IT systems, as well as how unavoidable losses due to the event will be minimized.

Having an information security management system means that the organization has a formal management framework in which the security controls are implemented and documented and where records are maintained in order to allow the evaluation of the controls and of compliance with procedures and standards.

III. The Links and Differences between Risk Management and Security Management

These activities make up the phase where the risk management process coincides with the information security domain. These two domains are closely related and indeed information security can, in a sense, be considered as an extension of the risk management process in that both processes share a common objective: a secure information infrastructure for the organization.

Information security management can be considered as a business process focused on the protection of information technology assets, with an underlying objective that is the continuity of the organization as a whole. While the risk management process involved in risk identification and treatment is rather a technical task, information security is driven by policies and procedures. Unlike risk management, information security is based on a top-down approach. Consequently, in order to be effective, information security management has to be performed in an integrated manner.

The aim of a risk management process is generally to identify, evaluate and mitigate risks by choosing appropriate protective measures. The information security process goes beyond this by making sure that the information security measures in hand are correctly implemented, and they are both effective and efficient. In fact an effective information security process ensures that the protection measures are implemented in an appropriate and applicable way and that they are consistently applied. It also falls into the scope of the information security process to ensure that the impact of every organizational, technical and human change is considered and the process is updated regularly.

One specific process is "Performance Measurement". During this process all information security fundamentals, such as objectives, strategies and policies, are reviewed to be up to date and provide an overall picture of the security situation. From this perspective, information security management activities allow management to optimize the core actions undertaken during the risk management phase.

Risk management focuses on the risks, while security management focuses on the organization. During the risk management phase risks are identified and during the security management phase the identified risks are placed in the context of the organization's environment. The risk treatment within a

given organization changes over time according to many variables such as time, current security level, and current exposure of the important assets.

The National Institute of Standards and Technology (NIST) [6] considers that risk management is an important component of a successful Information Security programme. Within the Information Technology security lifecycle, as modelled in the Plan-Do-Check-Act (PDCA) approach, the “Plan” phase matches exactly the risk management process. Activities such as the selection of the risk management method, risk classification, risk assessment and selection of control activities correspond to this “Plan” phase.

To summarize, risk management and information security management should not be considered as two distinct processes. It has been shown that they respond to the same objective, protecting valuable assets. In order to do that, both processes are necessary. Risk management will outline the appropriate security controls to be operated; security management ensures that the security controls are operated in the most effective and efficient way. Different and complementary competencies are clearly required to perform each process.

The information security process is driven by requirements, while risk management is driven by the importance of the impact and by acceptance criteria. This is because information security is policy-based and responds to efficiency-related objectives. Risk management identifies a large range of events and chooses the appropriate controls to be implemented. Information security management, via policies, strategies, and requirements, prioritizes risk and related controls.

Another distinction could be made when considering the result of processes. For risk management the result is the selection of countermeasures to mitigate the impacts of risks. In the case of information security, the result is a security status based on, but not restricted to, security controls.

IV. Case Study

A. Background

As a result of a takeover, a small manufacturing business in Switzerland found itself part of a large, publicly owned US-based group. Although not automatically subject to the requirements of the Sarbanes-Oxley (SOX) legislation as the activities of the entity did not meet any of the materiality criteria used to define the scope of SOX compliance within the group, senior group management insisted that the entity undertake a project to bring its security management into compliance with corporate standards.

During the technical discussions held before the acquisition, it was decided that in the short and medium terms at least the business would continue to operate its existing systems rather than migrate onto centralized corporate systems. This was because its core activity was in a highly-specialised sector of the industry and its current systems for resource planning, production control and financial reporting, which had been in place for several years, were highly customised and thus met the business’s needs far better than the acquirer’s own global ERP system.

Before its purchase, the entity had followed the traditional philosophy of family-run businesses, with very few formal policies and procedures and very little systematic and formal risk assessment. As a consequence, the security measures in place were ad hoc and did not consistently provide comfort to management over the integrity, confidentiality and availability of data.

Realizing that the SOX project, if performed properly, would require the re-evaluation of the internal controls procedures in place over all parts of the business, management decided to invest resources in re-engineering business processes as a whole and take the opportunity to implement a new, risk based security framework.

Local business management assumed entire responsibility for the quality of the project, but liaised very closely with group management in order to ensure that corporate standards in respect to controls design and scope were adhered to.

B. Project Activities

1) Planning and Risk Identification

Senior management first convened a steering committee that had the responsibility of driving the project. Among the first tasks was the definition of the key activities, which included risk assessment as a fundamental basis for further steps. Accordingly, the committee organized workshops to which financial, operational and IT management and key staff were invited and during which the company's overall control objectives were identified and the most significant risks that would prevent these objectives from being met were identified and discussed. The output from these sessions was a list of risks, operational, financial, technical and environmental, that were classified and prioritized and which the contributors felt presented the most significant areas of concern to the business.

2) Defining control objectives and designing control activities

The next stage required key staff in the business departments and IT to design individual control activities that would address the risks identified during the workshops. The input of IT into this process was crucial, as security underlay many of the objectives (particularly in relation to segregation of duties and restricting access to sensitive transactions and data) and business users did not necessarily possess sufficiently detailed knowledge of the configuration and structure of the business information systems to be able to design appropriate controls.

One specific area in which business management, IT and the internal audit function worked closely together was in redesigning the whole framework of access rights to the key business applications. The existing rights framework had evolved over a number of years based on a vague sentiment that access rights needed to be restricted; this has been mitigated by the demands of users for additional access rights in order to perform specific functions. As time passed, users who changed functions or moved between departments had tended to accumulate access rights and these accumulated profiles had often been copied for new joiners, leading to a situation in which management had no realistic idea of which users had access to which functionality or whether their access rights were appropriate.

Several person-weeks were accordingly spent in redesigning access rights from the ground up, considering which transactions were necessary for each position in the organization, which rights corresponded to the workflows being documented and formalized, and which rights would be necessary to ensure the smooth continuation of business during staff absences.

At the same time, procedures were defined for the request, attribution, removal and monitoring of access rights. For the first time, these procedures related to a formalized workflow involving human resources, line managers, IT and internal audit: a long-standing weakness in the control environment had been the lack of coordination and information passing between departments to ensure that access rights were granted, modified and revoked in an efficient and accurate manner.

3) Reviewing the operational effectiveness of the control activities

Understanding that a fundamental element of operating effective internal controls consists of measuring how well they are performed and how well they meet the control objectives, management implemented processes for monitoring the quality of their control environment.

The first measure was the implementation of a formal process for the recording and storage of documentation as evidence of the operation of controls. This fed into a management dashboard in which statistics on the quality and quantity of evidence were prepared and presented to senior management on a regular basis, allowing exceptions to be identified and investigated.

Management then planned two audit reviews of the internal control environment. The first, performed five months after the implementation of the new control framework and thus after the controls had had a reasonable time to be bedded in, was performed by the corporate internal audit

function. Their mandate was to review both the design effectiveness and operational effectiveness of the controls and conclude whether the controls corresponded to the group's guidelines and requirements.

The second review was performed at the end of the financial year by the group's external auditors as a part of their financial audit procedures. Their objectives were to conclude on the compliance of the internal control system with legal requirements in Switzerland and in the USA, and to determine whether the control framework was sufficiently robust and reliable to allow a move from largely substantive audit procedures to a more controls-based approach in future years.

C. Results

1) Administrative perspective

From a purely administrative perspective, the results of the project were highly positive for the company. The audit reviews found very few weaknesses in the internal control framework related to security, and those that were detected were easily remediated. In addition, the external auditors determined that the security measures in place, taken into consideration alongside with wider controls over IT management (such as controls over change management and operational activities), were sufficient to allow the proposed migration towards a controls-based audit approach. This had an impact both on the cost of the audit, as a number of detailed substantive audit procedures were replaced by fewer tests of the operation of controls, and on the focus, as more attention could be paid to issues of quality and consistency rather than traditional areas of recalculation and validation.

2) Business efficiency perspective

From a business efficiency perspective, one reported by staff and monitored by management, the results were equally positive. The focus on genuine risks to the business as a basis for the design and implementation of security measures meant that a number of processes had been streamlined or reengineered in order to be better managed, and this allowed staff to concentrate on the quality aspects of their tasks. The redesign of access rights allowed staff to perform their tasks more fluidly as they had access to the functionality they required, while the reinforced procedures for granting and modifying access rights meant that cases where changes were necessary were handled rapidly, systematically, and in a documented way.

3) Security perspective

Management rapidly saw the benefits of the risk-based security implementation from the confidence that it gave them over the use of the company's systems and resources. From low-level transactions such as the input and approval for payment of invoices to such high-level transactions as the input and approval of manual journal entries, an area of concern for many companies, management could be reasonably confident that only authorized staff were performing transactions. Supported by monitoring controls over the nature and timing of transactions, management felt confident for the first time over the way key business systems were being used and hence over the confidentiality and integrity of key business data.

An unanticipated consequence of the increased confidence of users in the quality and integrity of the data held in the key corporate systems was that increasing demands started being made for extractions of data to be used for detailed analysis, decision support and audit purposes. As a result of this, management was obliged to consider the risks related to the extraction, manipulation and storage of data and thus begin an additional cycle of the risk management and security management process.

Interviewed more than a year after the implementation of the revised security structures, management were reluctant to vaunt the quality of their security measures. The process of assessment and evaluation of risks and the requirement to adopt solutions based on compromises between access restrictions and usability of the systems, together with an understanding that security is an ongoing

process because risks and threats are constantly evolving, led management to conclude that their current level of security was reasonable based on their evaluation of their requirements and the risks they faced.

V. Conclusion

In the authors' experience, there are two significant barriers to the implementation of effective and efficient information security in organisations.

One is the popular belief that security is a technical issue, to be handled by IT alone and that given an adequate budget, a reasonable level of security can be attained.

The other is that very often security is bolted onto systems and integrated into operational processes as an afterthought, sometimes as a gesture towards auditors and sometimes as a result of a general understanding that some security measures are necessary, even if they are of little interest to users and managers.

As a result of this, the security around and within information systems in many organisations is inappropriate and ineffective. Organisational and cultural resistance to internal control measures often mean that measures are not implemented fully, are ignored, or are only operated to the letter rather than to the spirit. At the same time, the configuration of systems or the nature of business processes might not allow the security functionality in the information system to be properly utilised. Finally, as has been mentioned above, the controls in place might not correspond to the actual risks and thus serve little or no purpose to the organisation.

If an organisation takes the time, however, to identify, classify and prioritise its risks and then design a framework of internal controls including and based around information security, it will discover that numerous advantages accrue. Monitoring within the business becomes easier as there is greater clarity over the purpose and importance of activities; adherence to internal controls improves as staff understand the importance and value of the tasks they are performing; and users, management and third-parties can begin to have confidence in the quality and reliability of the information contained in and extracted from the information systems.

VI. References

- [1] ISO/IEC TR 13335-1: 2004 - Information technology—Security techniques—Management of information and communications technology security—Part 1: Concepts and models for information and communications technology security management
- [2] ISO/IEC 27005:2008 - Information technology -- Security techniques -- Information security risk management
- [3] Information Security Forum, *The Standard of Good Practice for Information Security*, 2007
- [4] ISO/IEC 27001:2005 - Information technology—Security techniques—Information security management systems
- [5] ISO/IEC 13335-2:2003 - Guidelines for the management of IT security Part 2: Managing and planning IT security
- [6] NIST Special Publication 800-30: Risk Management Guide for Information Technology Systems, National Institute of Standards and Technology, 2002

(5) Structured and unstructured data in the Cloud:

A Swiss perspective on readiness and internal controls

David Simms and Solange Ghernaouti
 Faculty of Business and Economics University of Lausanne
 Switzerland
 david.simms@unil.ch

Abstract— The concept of cloud computing has become one of the hot topics in the world of corporate information systems in recent years. Organisations are always interested in initiatives that allow them to pay less attention to the more mundane areas of information system management, such as maintenance, capacity management and storage management, and free up time and resources to concentrate on more strategic and tactical issues that are commonly perceived as being of higher value. A moment's consideration reveals, however, that the use of cloud computing services, like the use of outsourcing facilities, is not necessarily a panacea. Management will always retain responsibility for the confidentiality, integrity and availability of its applications and data, and being able to develop the confidence that these issues have been addressed. This paper outlines the background to cloud computing and presents the results of a survey of Swiss organisations and their attitudes towards data management and the use of the cloud.

Keywords - outsourcing, cloud computing, data management, unstructured data, risk management, security management, internal control, compliance

I. Introduction

Cloud computing is the subject of a great deal of discussion in information systems circles at the moment. The technology, described by NIST as “a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction” [1], is being taken up by individuals, in the form of Google Apps and Flickr and other on-line services, and by organisations and corporations that use contracted cloud services such as those offered by Google, Amazon and IBM, for example.

This article sets out to discuss the basic principles of cloud computing and the motivations of organisations who are already, or are considering, using cloud services specifically for data storage. It contains definitions and discussions of non-centralised and unstructured data and presents the background to, results of, and conclusions to be drawn from, a small-scale survey of attitudes towards data security and the use of the cloud carried out among organisations based in Switzerland. In conclusion the key issues arising from the research are identified and ongoing concerns are highlighted.

II. An Analytical Approach to Cloud Computing

A. Definition

Cloud computing is a paradigm in computing that has emerged from the spread of high-speed networks, the steady increase in computing power, and the growth of the Internet. The interconnection of resources that may be separated by significant distances is allowing users access to what appears to be a single resource. As a result, there is an opportunity to outsource resource-intensive or resource-specific tasks to service providers who will deliver the service for a fee, often based on consumption, rather than developing and maintaining the infrastructure and competences in-house.

Neither of the central ideas in this model, distributed computing or outsourcing, is new. Distributing resources within and across networks has been performed since the days of mainframe computers, where data were entered on remote terminals and processed centrally, while outsourcing of computing activities and services has taken place for many years for strategic, operational and financial reasons.

In technical literature reference is often made to cloud computing, grid computing and utility computing; there is no real consensus over whether there is a distinction to be made between the three and, if so, whether the distinctions are clear or are rather subtle. In general, though, the models are identified by features such as ubiquitous access, reliability, scalability, virtualisation, the exchangeability of and independence from location considerations, and cost-effectiveness [2].

A key area of growth is that of the types of service that can be offered by providers. Historically providers typically offered remote data processing, storage and systems management as major services, but the newer paradigm is to group offerings together as types of services: Software as a Service (SaaS), Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Service-Oriented Architecture (SOA). The underlying principle is that every facet of computing activities can be offered as services to consumers to be paid for on a usage basis.

This paradigm has been likened to the electrical power grid, both in the impact it has on social and industrial development [3] and because of its nature: elements of the grid can be owned and operated by different entities in different locations and might not share any physical characteristics, but the users, those who pay for the service, will only see a single interface or point of contact and will consume a consistent and homogenous service [4].

B. Advantages

The advantages of such services for certain users are, at least in principle, clear. The users need not worry about maintaining the application infrastructure, with all the operating system, database and application patches and upgrades that are necessary. Nor need the users worry about data management practices such as backups or transfer between machines and environments. Both individuals using office-style applications and remote storage services and large corporations using large-scale applications remotely will see the benefits of avoiding questions of ongoing software compatibility and upgrade paths. Essentially these services are seen as an opportunity to avoid having to deal with certain aspects of the complexity involved in managing a technological infrastructure.

For corporate users there are also the advantages related to the use of scalable architectures. Instead of making potentially significant capital investments in hardware and software that might never be fully exploited and thus represent an unnecessary cost to the business, management will be tempted to subscribe to the model of being able to request and exploit additional resources when required, as for as long as required, on something like a Pay-As-You-Go basis. This has the potential to allow a much more accurate and dynamic management of costs while still permitting absolute flexibility in the access to and use of resources.

The advantages for service providers of operating in this sector also seem to be well established. Once a data centre has been established and the reasonably fixed costs of operation have been established and incorporated into the business model, the variable costs of service provision and marginal costs related to the acquisition of additional clients or providing additional services or additional capacity to existing clients should be straightforward to manage and reasonably simple to recover (and exceed, of course) through appropriate pricing. Thus the provision of such services can be viewed as a potentially lucrative business, with important initial investment but steady and permanent future revenue streams.

C. Comparison with traditional outsourcing

Conceptually the techniques of cloud computing are not significantly different to those of traditional outsourcing of IT services. Many service providers offer outsourced data processing, system management, monitoring and control services and these services are very popular with many organisations who either do not wish to have internal IT services and competences for organisational reasons or prefer to outsource for financial or logistical purposes.

Key elements of any outsourcing agreement are the contract terms and the service level agreements. With these terms, it is clear to both parties which services are being provided, what resources are being made available, how these resources are being managed, and how the quality of service can be evaluated and managed.

Many structured cloud services will provide such terms and conditions but might not be able to specify every element of interest to the customer, such as the precise location where data are stored or processed.

It is in the nature of any kind of outsourcing or service provision activity that both parties will wish to maximise their revenues and benefits from the agreement while at the same time accepting the minimum level of responsibility for addressing the risks involved and for handling any issues that arise. In the context of cloud computing, customers need to pay particular attention to the clear definition of roles and responsibilities in order to try to avoid situations of blame-shifting and cost avoiding should problems subsequently arise.

D. Costs

A key driver for the use of cloud facilities is costs: organisations do not wish to tie up capital in IT infrastructure that might not be used to capacity and which may only have a short life before obsolescence when there exists the possibility of renting services as a regular P&L charge. Along with the resource and competency questions, which of course also incur ongoing costs and can be expensive to update or replace, this has long been a prime mover for outsourcing services. Experience in the domain of outsourcing, however, reveals that the cost savings may not always be as significant as hoped for. As mentioned above, the prudent management team will look to ensure that its systems and data are secure and reliable by implementing additional internal controls to generate evidence that there are no weaknesses in the service provider's controls that can be or are being exploited. Typically these internal controls will take the form of data and transaction analysis to identify exceptions or the appointment of specialist staff that can manage the relationship with the service provider and ensure that trends are identified, that service levels are maintained, and that issues are identified, reported and rectified. Very often the costs associated with implementing and operating such additional internal controls in an effective and robust manner are such that they can eat into the margins created by the whole outsourcing initiative.

E. Disadvantages

If the advantages for users of cloud services, as set out above, appear to be obvious, then the disadvantages are equally clear, particularly if viewed from a management and control perspective. The management of organisations always retains responsibility for the security and availability of their systems and data, in particular for the three key attributes of confidentiality, integrity and availability. Ensuring that these attributes are understood, managed and evaluated has traditionally been a significant challenge in systems management, even when systems and data are kept within a well-defined and efficiently policed perimeter. Once this perimeter is extended outside the sphere of direct control of relevant management, the challenge becomes increasingly difficult.

In an article published in the Gartner blogs [5], Thomas Bittman argues that the widely-used analogy of cloud computing to provision of electricity or water supplies is not particularly illustrative, for two reasons: "(1) Computing is a rapidly evolving technology, and (2) Service requirements vary widely for computing. Electricity production and distribution hasn't evolved much since the invention of AC that made distance distribution possible. How many forms of water are needed around the world? It's H₂O – maybe it can be potable, purified, or come at a special temperature, but it's still pretty basic stuff."

His analogy is that of transmitted music: radio broadcasts of music began in 1916 and phonograph cylinders were used for storage, with the idea being that people would not wish to bear the expense of storing their own copies of music when it was available over the airwaves. But technologies have advanced, both for the broadcast and transmission of music and for the local storage and consumption. Individuals continue to be prepared to pay a premium, for infrastructure, material and content, for a quality and rapidity of service.

The choice between broadcast and local access to music is the same as the cloud computing question: it will depend on a number of factors and requirements that are constantly evolving and ultimately become a question of the best balance between costs, risks, and quality. A key element in the cloud computing choice will be assessing the development of requirements and adopting a strategy that will maximise the Return on Investment, however that is calculated.

Such analogies do reach their limits, however, because they do not consider the specific nature of the service provided. When plugging in a laptop in a foreign hotel room, the consumer broadly speaking does not care who has provided the electricity and the identity of the provider has no impact on the use of the service. Using cloud services is fundamentally different, though, for as soon as a provider is storing data or performing processing for a client, there exist questions of the confidentiality and integrity of the systems and data as well as the availability of services.

A further aspect in which the analogy breaks down is in the area of the difficulty of changing approach once decisions have been taken. There is no question of not continuing to use electricity to power devices, but there will always be questions about the most efficient way to have access to and use IT infrastructure and resources. Moving from one cloud solution to another is likely to prove as, if not more, complicated than moving from one traditional outsourcing provider to another.

III. What Data are being put on the Cloud?

The opportunity to exploit the storage potential of the cloud can feasibly be viewed as an encouragement to organisations to place less priority on good practices in relation to data management. As storage space increases, so does the tendency simply to store all data rather than to classify, prioritise, archive and delete them as appropriate, and there is nothing in the principles of cloud computing to reverse this tendency.

It is therefore reasonable to envisage situations in which individuals and organisations simply do not know what data they have placed in the cloud. There may be multiple copies of the same documents and datasets. There may be inconsistent and incoherent datasets. There may be quantities of unstructured data – data extracted from central systems and databases that are not in standard, easily recognised and easily classified structures – that by their nature have not been classified and evaluated.

Of course, the question of unstructured and unmanaged data is not unique to the cloud: it exists in virtually all IT environments. Where it becomes particularly significant in the cloud environment, however, is as a consequence of the lack of visibility the data owners have over their data. Access to in-house, and properly secured outsourced data collections, can be defined, logged, and reviewed, at least in theory given sufficient resources and sufficient motivation on the part of the data owners. Once the data are in the cloud, however, the owners have very little visibility and control over the security of their data. If this weakness is exacerbated by an absence of real knowledge of what data are out there, the risk of data loss and disclosure is increased.

IV. Non-Centralised and Unstructured Data

A. Definition

The classical model for many businesses and other large organisations is to have one or several centralised key systems in which all of the organisation's financial and operational activities are recorded. Typically the data used in key applications and upon which users depend are stored in structured, centralised and at least theoretically well-controlled databases.

There is also, however, widespread use of non-centralised data repositories. Individual departments or users may have their own specific applications that do not fall into the overall organisation-wide systems landscape and thus are subject to different standards and procedures for management and control. This not to say that these systems and data are necessarily badly controlled, simply that management cannot necessarily be certain that standard policies and procedures are being applied.

Unstructured data is a term that has been widely employed in technical fields over recent years but which has suffered from a lack of precision in meaning. Definitions along the lines of “non-tabular” data suffer from being a description of what the data lack rather than a description of how and what they are [6]. For the purposes of this paper, we will use the term to mean data that do not reside in a structured, managed datastore, and which are not subject to requirements over structure, format or contents.

B. Use and frequency

In the modern business environment great use is made of unstructured data in a variety of contexts. Typically users will extract data from central databases, often those underlying ERP systems, and then manipulate these data in a variety of ways as part of their business activities. Such data are thus often found in spreadsheets, user-built databases and desktop or departmental server-based applications. These data can also be found in text files, pdfs, and even multimedia formats, depending on the use to which they have been put.

From an internal controls perspective the presence and use of unstructured data can pose numerous problems in respect of the confidentiality and integrity of data, two thirds of the famous “CIA” triad of information security objectives (the third being availability) [7]. When data are located in a structured central database, they can, at least in theory, be controlled, managed, verified and secured. Once extracted from the database, though, they can easily escape the internal control environment [8]. They can be used for decision taking without the assurance that they are still current, complete or valid. They can also, depending on the security measures in place and the efficiency with which these are enforced, leave the organisation easily, typically on USB media that have become ubiquitous, on laptop hard drives, or even as attachments to emails.

v. Activities in the Swiss Environment

A. Background to the survey

50 questionnaires were sent out to contacts in 43 organisations across Switzerland, using the lead author’s business experience to identify correspondents across a range of industries and sectors of activity who would be likely to respond. Completed questionnaires were received from 34 organisations, with three sending two responses and two sending three. The organisations that declined to respond did so on the grounds of confidentiality, not wishing to divulge details of their IT strategy or approach to security to a third party.

The organisations were selected in order to provide a wide cross-section of the range of IT environments and attitudes to the management of data and the use of new technologies. Of the organisations that did respond, five were publicly owned, eighteen privately, and the remaining eleven were public administrations or NGOs. The breakdown by organisation size also shows variety: eight with less than fifty employees; ten with between fifty-one and one hundred; five with between one hundred and five hundred; and eleven with more than five hundred.

The rationale behind sending multiple questionnaires to the same organisation was to attempt to discover whether there would be cases of poor communication of strategy or developments within those organisations. It is the lead author’s experience of large corporations in particular that there can be differences in the understanding of the overall strategy, the firm objectives, the initiatives undertaken and the impact on users between top management, middle management and IT management, for example.

B. The Survey

The questions in the survey were split into two sections, dealing with data security within the organisation and with data storage in the cloud, as follows:

Section A

Q1 Is there an overall policy concerning data management, security and retention?

Q2 Is there awareness at the level of senior management and/or security management of the concept of “unstructured data”?

Q3 Has there been any assessment of whether the organization should be concerned, from security or efficiency perspectives, about the existence of unstructured data?

Q4 Has the organization established any guidelines on the classification of data according to their sensitivity, age and relevance?

Q5 Has there been any structured attempt to identify and quantify the data stored around the organization outside centralized databases?

Q5A If so, was this a manual process?

Q6 Does the organization have policies on the extraction, use and storage of data by end-users?

Q6A If so, is compliance with these policies monitored and enforced, and how?

Q7 Are there policies and/or restrictions in place of the use of removable storage media for file transfer or storage?

Q7A If so, is compliance with these policies monitored and enforced and how?

Q8 Does the organization have any mechanisms for determining whether there have been breaches of security or confidentiality in respect of its sensitive data?

Section B

Q9 Is the organization using, or planning to use, cloud computing services for the storage of data?

Q10 If yes, have policies and guidelines been drawn up for the nature of data that can be stored in this way?

Q11 If cloud computing is being used, is the organization’s approach based on the centralized management, monitoring and retrieval of data, or do departments and/or individuals retain responsibility for their data?

C. Survey results

Section A

	Yes	No	Yes %	No %
Q1	32	9	78%	22%
Q2	22	19	54%	46%
Q3	17	24	41%	59%
Q4	9	32	22%	78%
Q5	6	35	15%	85%
Q5A	6	0	100%	0%
Q6	14	27	34%	66%
Q6A	5	9	36%	64%
Q7	12	29	29%	71%
Q7A	8	4	67%	33%
Q8	4	37	10%	90%

Section B

Q9	39	2	95%	5%
Q10	6	33	15%	85%
Q11	7	32	18%	82%

D. Interpretation and analysis of the results

The results demonstrated two major trends in respect of unstructured data. The first was that although 78% of organisations reported having designed and implemented an overall policy concerning data management security and retention, the exceptions being overwhelmingly small organisations with informal internal control structures, there was little systematic follow-up in terms of the management of unstructured data or in terms of managing and monitoring the use of data by users. 54% reported that senior management were aware of the issue of unstructured data; but only 41% reported that a risk analysis had been carried out to evaluate their exposure; 22% reported that guidelines for the classification of data had been developed and published; and only 15% reported having carried out a structured attempt to identify and quantify the data stored outside centralised databases. In each case it was a large multinational company that had undertaken such an initiative; interestingly, each one reported that the process had been largely manual.

In respect of the management of user activities related to data handling, 34% reported having policies on the extraction, use and storage of data by end users, and only 36% of these were able to report that compliance with these policies was monitored and enforced. Only 29% of organisations reported having policies or procedures in place concerning the use of removable storage media for file transfer or storage, and 67% of these were able to describe compliance mechanisms. Finally, only 10% of organisations were able to describe the existence of mechanisms for determining whether breaches of security or confidentiality in respect of sensitive data had occurred.

In the cases where more than one response was received from an organisation, it was noted that end users were less aware of the existence of strategies and policies than senior management, a situation that is consistent with the lead author's long experience of auditing large organisations.

In respect of the use of cloud computing facilities for the storage of data, 95% of the responses reported that the organisation was using or was planning to use such facilities. The reasons most frequently given were cost management, flexibility, and a desire to streamline IT activities to concentrate on more value-added activities in-house. Of these organisations, however, only 15% had drawn up policies and guidelines concerning the nature of data that could be stored in this way, and 18% were adopting an approach based on the centralised management, monitoring and retrieval of data rather than leaving it to departments or individuals to manage.

Overall it was possible to draw a clear distinction between large and small organisations and publicly and privately owned ones in respect of their approach to structured and formalised control environments. It was also possible to identify with high accuracy the responses received from publicly-owned corporations and those from organisations subject to other strict and demanding controls requirements such as Sarbanes-Oxley or local industry or environment-specific regimes. It was also possible to identify organisations that had significant internal audit or internal controls functions, or that had been alerted to the risks involved in going through a process of discovery in the context of a legal dispute.

E. Initial conclusions from the survey

The tentative conclusions that can be drawn from this very specific and targeted survey are the following.

Firstly, the importance of having a structured and documented approach to data management and security is widely understood among the organisations surveyed, with a particularly positive attitude towards risk management and compliance from larger organisations and those subject to definite compliance regimes because of their ownership or industry. How this understanding actually translates into positive measures designed to ensure compliance is another question, however, and IT security managers in particular reported seeing greater enthusiasm for establishing policies within their organisations than for implementing and complying with the necessary procedures. There was also a question of priorities and resources raised by smaller organisations that did not feel that such policies corresponded to or were a part of their core daily activities.

Secondly, the concept of unstructured data and the particular challenges posed by such data is reasonably widely understood, but there has been little activity outside large publicly-owned corporations to address the issue in any systematic way. In general IT departments had a good grasp of the nature and impact of the matter, and the subject was frequently raised by internal and external auditors and by legal advisers, but it was rarely considered to be a subject of great priority by senior management.

Thirdly, even if thought has been given within some organisations to managing employee access to and extractions of sensitive data, in the majority of cases monitoring is weak and compliance cannot be ensured. Generally speaking, users who have access to data stored in centralised databases tend to have the ability and the opportunity, and frequently the encouragement, to extract those data and use them for analytical or reporting purposes. Once the data have been extracted, access controls around them are usually weaker, often being restricted to network or workstation access controls, and even these restrictions reach their limits once data are copied onto portable devices such as USB keys.

Fourthly, very few organisations are in a position of being able to detect reliably whether their security has been breached or their data compromised, even when all their systems and data are hosted and managed internally. Indeed, in respect of both security breaches and employee misuse of data, it was reported that incidents were typically identified either by chance or on the basis of information received, rather than on the basis of regular and reliable compliance measures. Of course, in practice being able to design, implement and monitor the operation of such control activities is frequently non-trivial and requires competence, resources and careful planning. Whether organisations would be able to apply such controls to outsourced data, or be satisfied by the monitoring and reporting services provided by their cloud service provider, is the same question with an added layer of complexity.

Fifthly, the idea of cloud computing as a financially attractive option for outsourcing a number of traditional IT activities including data storage is widespread and there is a great deal of enthusiasm for it across a wide range of organisations, but this enthusiasm is not yet being widely and systematically backed up by detailed risk assessments and careful consideration of the approaches needed to identify, classify, manage and monitor the data being transferred into the cloud.

The comments provided alongside the answers also provided useful information. In particular, several respondents referred to the different types of cloud that are beginning to exist: in certain industries such as financial services it would be unthinkable to use cloud services in which confidential data might be stored outside Switzerland or for which it would be difficult to obtain adequate audit comfort over key concerns, but the use of some kind of industry-specific Swiss cloud, perhaps set up as a joint venture, with appropriate controls and safeguards in place, might be conceivable.

Several respondents also flagged up the importance of the proper management of backup media, which of course need to be subject to the same policies and procedures for data management and

security as live datasets. From the perspective of the management who will retain the responsibility for ensuring the availability and reliability of system and data backups for good practice and going concern reasons, and for the auditors who will be verifying this, it will be a challenge to identify exactly which data are backed up where, how this is managed from a security and availability perspective, and what the timelines, sequences and interdependencies would be for restoring part or all of a missing dataset.

VI. Overall Conclusion

Within all businesses there is constant pressure to reduce costs and cloud computing could be seen as an effective method of managing and reducing costs, particularly in the short-term. If there are no significant in-house IT systems, a case will always be made for reducing to a bare minimum, or even eliminating entirely, the IT function, thereby reducing staff costs alongside the operational costs of monitoring and maintaining systems.

The decision to choose cloud computing services is not one to be taken lightly. For individuals, the use of personal services in the cloud (such as Facebook, gmail and Dropbox) is, or rather should be, a matter of a calculated assessment of risks and benefits, for the potential negative impacts of breaches in security, for example, can be significant. For businesses and other organisations, the same concerns apply but on a larger scale. For all organisations that have a responsibility to keep their, and others', data secure and confidential, the decision can only be taken after a detailed analysis of how they will obtain, and continue to obtain, the necessary comfort that this is the case. If they cannot build into their own procedures, into enforceable contract terms, and into audit plans, the means of confirming the confidentiality, integrity and availability of their systems and data, they should not consider externalising it.

In addition, organisations should not forget the immediate costs and efforts involved in moving onto the cloud. Datastores need to be identified, classified, cleaned up and archived, and serious technical and operational decisions are needed to determine what data will sit where. This will frequently be a project of a significant size requiring expertise, resources and input from a number of people across the business who understand the business, the systems, the data, and their use.

Experience of traditional outsourcing suggests that it is very easy for organisations to overestimate the cost savings generated by a move towards service providers and to underestimate the amount of internal competence and dedicated management required to make a success of such initiatives. As long as the organisation relies on its data and retains responsibility for all aspects of its business from a regulatory perspective, it will need to ensure that its management of the relationship with its service providers and its access to critical operational information are both adequate and appropriate. Typically this will require retaining or recruiting skilled, experienced and reasonably senior staff to liaise with and monitor the performance of the service provider.

The long-term consequences of opting for a cloud solution also need to be examined. Once on the cloud, systems and data are likely to stay there, and feasibly with the same provider. What begins as a simple and cost-effective solution to a small problem could develop into a long-term strategic commitment with little scope for alteration.

References

- [1] "The NIST Definition of Cloud Computing", Special Publication 800-145, NIST, September 2011
- [2] "Distributed Computing: Utilities, Grids and Clouds", ITU-T Technology Watch Report 9, 2009, p2
- [3] *ibid*, p1
- [4] *ibid*, pp1-2
- [5] http://blogs.gartner.com/thomas_bittman/2009/09/22/a-better-cloud-computing-analogy/
- [6] <http://www.dataversity.net/what-are-unstructured-data/>, retrieved on 10/05/2012
- [7] <http://www.techrepublic.com/blog/security/the-cia-triad/488>, Retrieved on 10/05/2012
- [8] "Optimizing security efficiency through effective risk management", Ghernaouti-Hélie S, Tashi I, Simms D, Conference Paper, AINA 2011, p.6