

INFORMATION RISK MANAGEMENT

A practitioner's guide
Second edition

David Sutton



<https://t.me/learningbcs>

INFORMATION RISK MANAGEMENT

BCS, THE CHARTERED INSTITUTE FOR IT

BCS, The Chartered Institute for IT, is committed to making IT good for society. We use the power of our network to bring about positive, tangible change. We champion the global IT profession and the interests of individuals, engaged in that profession, for the benefit of all.

Exchanging IT expertise and knowledge

The Institute fosters links between experts from industry, academia and business to promote new thinking, education and knowledge sharing.

Supporting practitioners

Through continuing professional development and a series of respected IT qualifications, the Institute seeks to promote professional practice tuned to the demands of business. It provides practical support and information services to its members and volunteer communities around the world.

Setting standards and frameworks

The Institute collaborates with government, industry and relevant bodies to establish good working practices, codes of conduct, skills frameworks and common standards. It also offers a range of consultancy services to employers to help them adopt best practice.

Become a member

Over 70,000 people including students, teachers, professionals and practitioners enjoy the benefits of BCS membership. These include access to an international community, invitations to a roster of local and national events, career development tools and a quarterly thought-leadership magazine. Visit www.bcs.org/membership to find out more.

Further information

BCS, The Chartered Institute for IT,
3 Newbridge Square,
Swindon, SN1 1BY, United Kingdom.
T +44 (0) 1793 417 417
(Monday to Friday, 09:00 to 17:00 UK time)
www.bcs.org/contact
<http://shop.bcs.org/>



INFORMATION RISK MANAGEMENT

A practitioner's guide
Second edition

David Sutton



© BCS Learning and Development Ltd 2021

The right of David Sutton to be identified as author of this work has been asserted by him in accordance with sections 77 and 78 of the Copyright, Designs and Patents Act 1988.

All rights reserved. Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted by the Copyright Designs and Patents Act 1988, no part of this publication may be reproduced, stored or transmitted in any form or by any means, except with the prior permission in writing of the publisher, or in the case of reprographic reproduction, in accordance with the terms of the licences issued by the Copyright Licensing Agency. Enquiries for permission to reproduce material outside those terms should be directed to the publisher.


All trade marks, registered names etc. acknowledged in this publication are the property of their respective owners. BCS and the BCS logo are the registered trade marks of the British Computer Society charity number 292786 (BCS).

Published by BCS Learning and Development Ltd, a wholly owned subsidiary of BCS, The Chartered Institute for IT, 3 Newbridge Square, Swindon, SN1 1BY, UK.
www.bcs.org

Paperback ISBN: 978-1-78017-5720

PDF ISBN: 978-1-78017-5744

ePUB ISBN: 978-1-78017-5751



British Cataloguing in Publication Data.

A CIP catalogue record for this book is available at the British Library.

Disclaimer:

The views expressed in this book are of the authors and do not necessarily reflect the views of the Institute or BCS Learning and Development Ltd except where explicitly stated as such. Although every care has been taken by the authors and BCS Learning and Development Ltd in the preparation of the publication, no warranty is given by the authors or BCS Learning and Development Ltd as publisher as to the accuracy or completeness of the information contained within it and neither the authors nor BCS Learning and Development Ltd shall be responsible or liable for any loss or damage whatsoever arising by virtue of such information or any instructions or advice contained within this publication or by any of the aforementioned.

All URLs were correct at the time of publication.

Publisher's acknowledgements

Reviewers: Andrea Simmons

Publisher: Ian Borthwick

Commissioning editor: Rebecca Youé

Production manager: Florence Leroy

Project manager: Sunrise Setting Ltd

Copy-editor: The Business Blend Ltd

Proofreader: Barbara Eastman

Indexer: Matthew Gale

Cover design: Alex Wright

Cover image: Shutterstock/Pat-s-pictures

Typeset by Lapiz Digital Services, Chennai, India

DEDICATION

While updating this book, the UK was locked down due to the coronavirus SARS-CoV-2, and we were unable to leave home except for food shopping, essential exercise or medical needs. This gave me the opportunity to concentrate fully on the book instead of the usual procrastination and finding other things to do, even though there was a long list.

What struck me above all else was the dedication and sheer determination shown by many people. In particular, doctors, nurses, carers, hospital staff, police and ambulance drivers, who as front-line responders put their lives on the line to save others, and some of whom unfortunately lost their lives in doing so. But many others did not receive the same recognition and deserve a mention – shop workers, who made sure that we could buy essential items (even if some were in short supply for a while); the people producing our food and essential needs; delivery drivers, who made sure the shops and supermarkets were stocked; refuse and recycling collectors; transport workers, who kept the country moving – all of whom who carried on their daily work despite the risks and often without proper thanks.

Friends, neighbours and frequently total strangers rallied round to make sure that the elderly and the less able continued to receive food and essential medication or just to have a telephone call with another human being while in isolation.

Groups of individuals began making personal protective equipment for front-line staff who lacked it, often paying for the materials out of their own pockets or crowdfunding money for the resources they needed.

All these unselfish people did what they did without being asked to do so, and demonstrated just how much a crisis can bring communities together, and bring out the best in the human race.

Many of these people are underpaid and undervalued, and I hope that if nothing else comes of this, they will receive the recognition they so rightly deserve, and it is to all of the above that I would like to dedicate this book.

CONTENTS

List of figures and tables	x
Author	xi
Other works by the author	xi
Acknowledgements	xii
Abbreviations	xiii
Preface	xvi
1. THE NEED FOR INFORMATION RISK MANAGEMENT	1
What is information?	1
Who should use information risk management?	4
The legal framework	6
The context of risk in the organisation	7
Hot topics to consider in information risk management	9
The benefits of taking account of information risk	12
Overview of the information risk management process	14
Summary	16
2. REVIEW OF INFORMATION SECURITY FUNDAMENTALS	20
Information classification	22
Plan-Do-Check-Act	26
Summary	27
3. THE INFORMATION RISK MANAGEMENT PROGRAMME	28
Goals, scope and objectives	29
Roles and responsibilities	30
Governance of the risk management programme	30
Information risk management criteria	31
Summary	36
4. RISK IDENTIFICATION	37
The risk identification process	37
The approach to risk identification	39
Impact assessment	42
Summary	53
5. THREAT AND VULNERABILITY ASSESSMENT	54
Conducting threat assessments	54
Conducting vulnerability assessments	60

	Identification of existing controls	66
	Summary	69
6.	RISK ANALYSIS AND RISK EVALUATION	71
	Assessment of likelihood	71
	Risk analysis	74
	Risk evaluation	76
	Summary	80
7.	RISK TREATMENT	81
	Strategic risk options	82
	Tactical risk management controls	85
	Operational risk management controls	86
	Examples of critical controls and control categories	87
	Summary	90
8.	RISK REPORTING AND PRESENTATION	91
	Business cases	92
	Risk treatment decision-making	93
	Risk treatment planning and implementation	94
	Business continuity and disaster recovery	95
	Disaster recovery failover testing	104
	Summary	104
9.	COMMUNICATION, CONSULTATION, MONITORING AND REVIEW	105
	Skills required for an information risk programme manager	105
	Communication	107
	Consultation	109
	Risk reviews and monitoring	110
	Summary	112
10.	THE NCSC CERTIFIED PROFESSIONAL SCHEME	113
	SFIA	115
	The CII Sec skills framework	116
	Summary	119
11.	HMG SECURITY-RELATED DOCUMENTS	120
	HMG Security Policy Framework	120
	The National Security Strategy	120
	CONTEST, the United Kingdom's Strategy for Countering Terrorism	121
	The Minimum Cyber Security Standard	121
	The UK Cyber Security Strategy 2016–2021	121
	UK government security classifications	122
	Summary	123
	APPENDIX A – TAXONOMIES AND DESCRIPTIONS	124
	Information risk	124
	Typical impacts or consequences	126

APPENDIX B – TYPICAL THREATS AND HAZARDS	130
Malicious intrusion (hacking)	130
Environmental threats	133
Errors and failures	135
Social engineering	137
Misuse and abuse	138
Physical threats	139
Malware	140
APPENDIX C – TYPICAL VULNERABILITIES	143
Access control	143
Poor procedures	146
Physical and environmental security	147
Communications and operations management	149
People-related security failures	151
APPENDIX D – INFORMATION RISK CONTROLS	154
Strategic controls	154
Tactical controls	155
Operational controls	155
The Centre for Internet Security Controls Version 8	156
ISO/IEC 27001:2017 controls	158
NIST Special Publication 800-53 Revision 5	163
APPENDIX E – METHODOLOGIES, GUIDELINES AND TOOLS	171
Methodologies	171
Other guidelines and tools	176
APPENDIX F – TEMPLATES	181
APPENDIX G – HMG CYBERSECURITY GUIDELINES	187
HMG Cyber Essentials Scheme	187
10 Steps to Cyber Security	189
APPENDIX H – REFERENCES AND FURTHER READING	192
Primary UK legislation	192
Good Practice Guidelines	193
Other reference material	193
NCSC Certified Professional Scheme	194
Other UK government publications	195
Risk management methodologies	196
UK and international standards	196
APPENDIX I – DEFINITIONS, STANDARDS AND GLOSSARY OF TERMS	204
Definitions and glossary of terms	205
Information risk management standards	213
Index	216

LIST OF FIGURES AND TABLES

Figure 1.1	The information life cycle	4
Figure 1.2	The overall risk management process	15
Figure 2.1	The Plan-Do-Check-Act cycle	26
Figure 4.1	A general view of the risk environment	38
Figure 4.2	Typical types of information asset	39
Figure 4.3	Generic sequence of situation management	41
Figure 4.4	A simple threat, vulnerability and impact	42
Figure 4.5	Multiple threats can exploit the same vulnerability	42
Figure 4.6	A single threat can exploit multiple vulnerabilities	43
Figure 4.7	A typical chain of consequence	43
Figure 4.8	Impact types	44
Figure 4.9	Potential losses over time following a disruptive event	51
Figure 4.10	Typical impact assessment form	52
Figure 5.1	Typical threats and hazards	56
Figure 5.2	Typical threat assessment form	61
Figure 5.3	Typical vulnerabilities	63
Figure 5.4	Typical vulnerability assessment form	67
Figure 5.5	The overall scheme of risk treatment options	69
Figure 5.6	Typical existing controls identification form	70
Figure 6.1	A typical risk matrix	75
Figure 6.2	An enhanced risk matrix	76
Figure 6.3	A typical risk register spreadsheet	78
Figure 7.1	The overall scheme of risk treatment options	82
Figure 7.2	The strategic risk management process	83
Figure 8.1	The BCI life cycle	97
Figure 8.2	The generic business continuity incident timeline	98
Figure 8.3	Overall structure for disaster recovery	99
Figure 8.4	Cost versus availability	101
Figure A.1	An overall taxonomy of information risk	124
Figure A.2	Typical impacts or consequences	127
Figure B.1	Typical threats and hazards	131
Figure C.1	Typical vulnerabilities	144
Figure D.1	Information risk controls	154
Figure I.1	Concepts and relationships	205
Table 4.1	The general properties of detrimental situations	40
Table 4.2	Typical impact scales	48
Table 6.1	Typical likelihood scales	74

AUTHOR

David Sutton's career spans more than 55 years and includes radio transmission, international telephone switching, computing, voice and data networking, structured cabling systems, information security and critical information infrastructure protection.

He joined Cellnet (now Telefónica UK) in 1993, where he was responsible for ensuring the continuity and restoration of the core cellular and broadband networks, and represented the company in the electronic communications industry's national resilience forum. In December 2005 he gave evidence to the Greater London Authority enquiry into the mobile telecoms impact of the London bombings.

David has been a member of the BCS Professional Certification Information Security Panel since 2005 and delivered lectures on information risk management and business continuity at the Royal Holloway University of London, from which he holds an MSc in Information Security.

He is a Chartered Fellow of BCS, the Chartered Institute for IT, a member of the Chartered Institute for Information Security (CIIISec), a Freeman of the Worshipful Company of Information Technologists and a Freeman of the City of London.

OTHER WORKS BY THE AUTHOR

Cyber Security: A Practitioner's Guide. BCS, 2017. ISBN 978-1-78017-340-5

Business Continuity in a Cyber World: Surviving Cyberattacks. Business Expert Press, 2018. ISBN 978-1-94744-146-0

Information Security Management Principles, Third edition (co-author). BCS, 2020. ISBN 978-1-78017-518-8

Data Governance: Governing Data for Sustainable Business (contributor). BCS, 2021. ISBN 978-1-78017-375-7. Pages 87–96

ACKNOWLEDGEMENTS

I would like to thank Ian Borthwick and Rebecca Youé of BCS for kindly agreeing to publish this book; my wife Sharon for her unceasing encouragement; my children Bella, Matt and James, and their respective partners for their support; and my wonderful grandchildren for regularly reminding me that there's much more to life than work.

Finally, I would like to thank Mr Evans, my English teacher at Thomas Adams School in Wem, for reasons that I hope will be obvious.

ABBREVIATIONS

AI	Artificial Intelligence
APM	Association for Project Management
BC	Business Continuity
BCI	Business Continuity Institute
BCM	Business Continuity Management
BCP	Business Continuity Plan
BCS	BCS, The Chartered Institute for IT
BIA	Business Impact Analysis
BR	Business Resumption
BS	British Standard
BSI	British Standards Institution
BYOD	Bring Your Own Device
CCP	Certified Cyber Professional
CCTV	Closed-Circuit Television
CD	Compact Disc
CDPA	Copyright, Designs and Patents Act 1988
CEO	Chief Executive Officer
CIA	Confidentiality, Integrity and Availability
CIISec	Chartered Institute of Information Security
CMA	Computer Misuse Act 1990
CMM	Capability Maturity Model
CNSS	Committee on National Security Systems
COMAH	Control of Major Accident Hazards
DAS	Direct Attached Storage
DCMS	Department for Digital, Culture, Media and Sport
DDoS	Distributed Denial of Service
DoS	Denial of Service
DPA	Data Protection Act 1998, 2018
DR	Disaster Recovery
DVD	Digital Versatile Disc

ABBREVIATIONS

ENISA	European Network and Information Security Agency
ERM	Enterprise Risk Management
EU	European Union
FAIR	Factor Analysis of Information Risk
GCHQ	Government Communications Headquarters
GDPR	General Data Protection Regulation
GPG	Good Practice Guidelines
HMG	Her Majesty's Government
HR	Human Resources
HTML	Hypertext Markup Language
IA	Information Assurance
IASME	Information Assurance for Small and Medium Sized Enterprises
ICT	Information Communications and Technology
IEC	International Electrotechnical Commission
IISP	Institute of Information Security Professionals
IM	Incident Management
IoT	Internet of Things
IP	Intellectual Property
IP	Internet Protocol
IRM	Institute of Risk Management
ISF	Information Security Forum
ISMS	Information Security Management System
ISO	International Organization for Standardization
ISP	Internet Service Provider
IT	Information Technology
ITU	International Telecommunication Union
LAN	Local Area Network
MAO	Maximum Acceptable Outage
MBCO	Minimum Business Continuity Objective
MRI	Magnetic Resonance Imaging
MTDL	Maximum Tolerable Data Loss
MTPD	Maximum Tolerable Period of Disruption
NAS	Network Attached Storage
NCSC	National Cyber Security Centre
NIST	National Institute for Standards and Technology
NSA	National Security Agency
OCTAVE	Operationally Critical Threat, Asset, and Vulnerability Evaluation
PAS	Publicly Available Specification
PCI DSS	Payment Card Industry Data Security Standard

PDA	Personal Digital Assistant
PDCA	Plan-Do-Check-Act (aka the Deming Cycle)
PDSA	Plan-Do-Study-Act
PIN	Personal Identification Number
RAID	Redundant Array of Inexpensive Disks
RIPA	Regulation of Investigatory Powers Act 2000
RPO	Recovery point objective
RTO	Recovery time objective
SABSA	Sherwood Applied Business Security Architecture
SAN	Storage Area Networks
SFIA	Skills Framework for the Information Age
SQL	Structured Query Language
TLP	Traffic Light Protocol
UPS	Uninterruptible Power Supply
VLAN	Virtual Local Area Network
VPN	Virtual Private Network
WAP	Wireless Access Point
Wi-Fi	Wireless Fidelity

PREFACE

In the six years since I wrote the original *Information Risk Management* book, much has changed in terms of technology and the threats to information. Little, however, has changed in terms of vulnerabilities. Chief among these is that many organisations (and often the most senior executives within them) believe that information risk is purely a technology problem, and ignore the fact that processes, procedures and people are often not only at the root of information risk issues, but also one of the principal means of resolving or avoiding them.

Technology is frequently the tool we use to secure information as well as to generate and store it, and these activities are easily interchanged in people's minds, resulting in confusion and misinterpretation. After all, if you leave your car unlocked and your mobile phone, wallet or laptop are stolen, it is not the car's fault is it?

It is time we stopped blaming technology for all our woes, and concentrated instead in understanding not only what is happening, but also and more importantly, why it is happening. Then and only then we can do something positive about it; prevent it from happening in the first place, and also prevent it from recurring.

It does not actually matter whether the information is in physical or electronic form; what matters is that it is important to someone and therefore warrants protection from theft or abuse.

It is an unfortunate fact of life that we do not always value things until they are lost. This is especially true of information. Were the last digits of someone's telephone number 674 or 647? Does a colleague live at number 24 or number 42? While these are trivial examples of the loss or misunderstanding of information, they serve to illustrate how dependent we are on information of all kinds, but they fall short of recognising the effects of information either being permanently lost or (possibly worse) falling into the wrong hands.

In recent years, there have been numerous reports in the media about how the security services, particularly in the UK and the USA, are intercepting our private communications, and while this in itself is laudable in the fight against organised crime and international terrorism – it is, after all, their primary role – it is clear that some governments, and indeed organisations and people, may have different objectives and are seeking to mine our information in order to use it either for their financial gain at our expense or to take advantage of us in some way.

The general principles we use to protect our information can be found in *Information Security Management Principles* Third edition, published by BCS, [Chapter 2](#) of which deals with information risk. However, this is only a 20-page summary account of the subject, and therefore only scratches the surface.

The lesson – as many a security professional will tell you – is that if a well-resourced opponent really wants to read your information, remove it or change it, then they will find a way of doing so. It may not be cheap or easy, it may involve using a mix of technology and human agents, but if they think it is worth it, you will find it very, very hard to stop them.

The intention of this book is therefore to help you to make life as difficult as possible for them to be successful.

The technology, tools, standards, regulations and methods incorporated in information systems all change at a considerably faster rate than the updates to books such as this. Although all the detail included has been verified at the time of writing, and again during the publication process, there will always be discrepancies between the book and the real world. Hopefully, there will be sufficient information in the book to allow readers to identify these, and to confirm the most up-to-date information.

1 THE NEED FOR INFORMATION RISK MANAGEMENT

In this first chapter of the book, we shall set the scene for the later chapters by focusing on what information actually is and how it is produced or obtained, why we should manage the risks to information, the legal framework surrounding information, and the context of risk within organisations.

We shall take a brief look at some of the hot topics in information risk management, including the Internet of Things and remote working, before discussing the benefits of information risk management and some of the processes by which it can be achieved.

WHAT IS INFORMATION?

Before we begin to examine the need for information risk management, it is important to understand what the difference is between information and data.

Superficially, this appears to be quite straightforward – data are merely unstructured facts and figures, whereas information consists of data that are organised into a meaningful context. For example, the temperature, wind speed and direction, rainfall and atmospheric pressure readings taken twice daily in towns and cities around the country are just data. It is only when they are recorded together, and along with those readings of previous days, that the data are placed in context and begin to have meaning, allowing meteorologists to examine trends and develop a weather forecast. It is at this point that the data have become organised and structured and can now be seen as information.

Although I have drawn the distinction between the two, for the purposes of this book I shall deal with them both under the heading of 'information', since both data and information will have value to their owners and must be equally protected, although the owner of the original data and the owner of the resulting information may be entirely different entities.

Information can exist in two different states: physical, with information recorded on paper, film, paper tape, canvas, pieces of clay with cuneiform indentations and notches in tally sticks; and with virtual binary ones and zeros stored on magnetic media or other types of electronic memory device.

Information also comes in two distinct forms. Firstly, there is information that describes or lists other information, such as a catalogue or index, and is often referred to as 'metadata'. Secondly, there is information that is something in its own right, such as a

novel, a software application or the formula for a new medicinal drug. All have value to their owner or originator, and indeed may either be of a personal nature, in which case might be subject to data protection legislation, or may be IP, in which case copyright or trademark legislation will apply.

It is not my intention to deal in any depth with either of these two aspects of legislation since each could easily be the subject of a book in its own right, but you should be aware not only of their existence and general content, but also that they need to be taken into account when developing an information risk management programme.

Recent revelations regarding the organised interception and mining of information by various security agencies have raised awareness at all levels of society of the need to take greater care of our information, but we should not be at all surprised by the extent to which this so-called 'snooping' takes place, or by the fact that these agencies are able to carry it out.

This problem lies in the distinction between the need to maintain national security and the need to gather sufficient information to be able to do so. Security agencies such as the National Security Agency (NSA) in America and Government Communications Headquarters (GCHQ) in the UK were set up precisely to carry out this kind of work, so it should not come as a shock to anybody that they are doing it, nor that they are very successful at doing so albeit subject to strict legal undertakings, at least in theory. What should be more worrying is that other nations' security agencies may be able to undertake similar surveillance and interception and may use the resulting information gathered for nefarious purposes.

Then there is the question of so-called 'Big Data', in which organisations – both commercial and governmental – collect vast amounts of information on us as individuals. Every time we use a credit card to purchase goods, the credit card agency gathers a little more information about us. This has positive benefits as well as negative connotations; for example, if a transaction falls outside your 'normal' spending profile, the credit card agency can contact you to verify that your card is still in your possession and has not been used fraudulently.

On the other hand, of course, supermarkets may target us with advertising and promotions as a result of aggregating information gained from our loyalty cards, which may or may not be something to be happy about, since they now know more about our spending habits than we do!

A recent investigation¹ into how much Amazon knows about us unearthed some interesting and somewhat alarming results – not only about how much use they make of our browsing and spending habits, what films we watch and what music we listen to, but also about how many data their 'Ring' doorbell/video camera records, and what they are able to infer from our commands to the 'Alexa' devices.

Similar concerns revolve around Google's ability to monitor our habits and movements when we use their search engine or ask the Google 'Home' devices for information. Both

¹ See 'Amazon is watching, listening and tracking you. Here's how to stop it' (phys.org).

Alexa and Google Home additionally allow us to control aspects of our homes – lights, sockets, closed-circuit television (CCTV), baby monitors and central heating, all from one application on a smartphone.

All this may be extremely useful to us as users, but continues to raise questions over whether others are learning more about us than we might care for them to know, and whether they could ultimately take at least partial control over certain aspects of our lives.

In the UK, there is an ongoing and often heated debate about the use of network infrastructure from the Chinese company Huawei. On the one hand, there is the fear that their possible links with the Chinese government might enable it to have unwanted influence on our lives, including unfettered access to more sensitive information. On the other hand, its cost to the network operators may be significantly lower than that of other suppliers, allowing them to keep call charges to users at a lower rate. The view at the time of writing is that Huawei will be allowed to provide some of the fifth-generation mobile network infrastructure, while the more sensitive 'core' of the networks will be closed to them.

Whatever the situation, we sometimes do not treat our own or other people's information with sufficient care, and the consequences of this can be severe. When scaled up from a personal to an organisational level, the consequences can be catastrophic, and it is hoped that this book will enable you to take a proactive position in preventing this from happening.

Finally, we should make the distinction between information that is about what we do, and information about who we are. Information about what we do could cover such things as where we spend our money, what our audio and visual entertainment preferences are, what we view on the internet, what we say online and anything that can be recorded about actions we have undertaken.

Information about who we are will include those so-called immutable attributes. These are absolute facts and can never be altered. They include such things as our biological parents, our biometrics (for example, iris scan, fingerprints or DNA) and where and when we were born.

Next there are so-called assigned attributes such as our nationality, names, national insurance number or title. These are generally the attributes that people and organisations rely upon to identify and communicate with us, and rarely change.

Finally, there are other related attributes, which, while being a part of our personae, are more easily changed, but still allow people and organisations to identify and communicate with us, and which may be used in identity verification, such as usernames and passwords, email addresses, memberships, qualifications and entitlements.


Many of these types of information are almost impossible to conceal since they are a matter of public record and generally speaking we are happy to make them available – indeed, it is often in our interests to do so, although there are some that we would naturally not make publicly available. For example, we are usually happy to give

someone our email address, but at the same time we would not let them know the password to the email account.

The information life cycle

It is easy to imagine that information is 'just there', but it must be created in the first place, and then generally follows a set path, as shown in [Figure 1.1](#).

Figure 1.1 The information life cycle



The creative process begins with some form of research, design or discovery, which allows the creator to record the information in some form, whether in hard copy or electronic form, and then to store it in some way. In some situations, the information may be processed somehow, either to manipulate it in a way that others can easily access it, or to make it more useful by enriching it in some way, perhaps by amalgamating it with other information.

The process continues with use, either by the information's creator alone, or more frequently by others, whether individually or collaboratively, at which point it can be widely shared within a contained environment or publicly.

At some stage, the information may become out of date but still be required as a time-based reference, in which case it will be archived. Eventually, the information will become completely redundant, at which point it can safely be disposed of or destroyed, or may be updated and recycled as new information.

At each stage of this life cycle process there will be the need to ensure that the information is adequately protected from accidental or deliberate loss, change or destruction, hence the need for information risk management.

WHO SHOULD USE INFORMATION RISK MANAGEMENT?

Quite simply, any part of an organisation can and should make use of information risk management, since all parts of an organisation are likely to have information that has value to it.

The human resources department keep records of personnel, much of which will be considered to be personal information under general data protection legislation; sales and marketing departments will hold information on past and projected sales as well as pricing schedules; finance will hold records of the organisation's income and expenditure; development will have plans and designs for both current and future products and services; and the IT department, although perhaps not owning any of this information, will be responsible for keeping it secure and making it available to authorised staff.

Non-commercial organisations too will have valuable information that must be protected. Hospitals, GP practices and health trusts hold sensitive personal information on patients; local authorities hold lists of vulnerable people; the Driver and Licensing Agency holds details of every driver and vehicle registered in the UK; and Her Majesty's Revenue and Customs hold huge amounts of financial information about every taxpayer in the country and beyond.

All these different types of information must be protected – they must be kept confidential, so that only authorised people may have access to them; their integrity must be protected, so that only authorised people may change them; and they must be available when required by those who have a need to access them. These three main tenets of information security – confidentiality, integrity and availability – underpin everything in this book, and will be dealt with in greater detail in [Chapter 2](#).

However, in order to protect our own or our organisation's information, we first need to understand exactly what it is and why it is important to the organisation.

An excellent example of the need for protecting information goes back to the 1940s, when, during the Second World War, the British government put up posters declaring 'Careless talk costs lives'. The meaning was clear. People who were aware of military plans might innocently reveal them by indiscreet conversation, and the consequences could be extreme for the military and civilian personnel who were taking part in those actions or whose environment might be affected as a result of them. The implication of this was that any information revealed could unwittingly lead to a compromise of security, but it gave no indication of how sensitive the information might be, the consequences of revealing it or how it should be protected.

This brings us to the issue of information classification, in which each piece of information can be classified for its sensitivity, handling, storage, access or distribution and ultimately its disposal. The only problem with information classification is that it does not usually reveal the potential value (monetary or otherwise) of the information either to the organisation itself, or to an adversary who might be able to benefit from obtaining it.

Yet another aspect to be considered is information aggregation, in which small pieces of information (some of which might appear completely trivial) are gathered, often from a variety of sources, and pieced together to provide a clearer picture of the whole.

All these elements are brought together in the techniques we use for information risk management, which allows us to clearly identify those information assets that have value to our organisation; determine the impact on the organisation of their unauthorised

distribution, alteration or destruction; assess the vulnerabilities exhibited by them; and assess the events that might bring these about, and the likelihood of these occurring.

All this provides us with a measure of the level of risk associated with each type or piece of information, from which we can determine the most appropriate response while balancing the possible consequences against the cost of treatment.

Some people believe that risk assessments are only necessary in a health and safety situation, but where personal information is concerned, there is also a legal obligation to ensure its proper protection, and the General Data Protection Regulation (GDPR) states that:

Personal data shall be processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures ('integrity and confidentiality').

As we have already mentioned, information is not restricted to the IT department, so every part of an organisation can therefore benefit from the use of information risk management.

If we are going to protect our information assets, we need to understand what they are, what might threaten their confidentiality, integrity or availability, how they might be vulnerable to such threats or hazards, and how likely these are to occur. This, in short, is the key role of information risk management, the essential components of which are discussed later in this chapter.

THE LEGAL FRAMEWORK

Safeguarding information did not present too many problems until computers, especially personal computers, became widespread. It was only with the introduction of the Computer Misuse Act (CMA) in 1990 that people outside government really began to take unauthorised access to information seriously. Since that time, earlier legislation has been updated to reflect the changes in the accessibility of information, and other legislation designed to better protect information has been developed.

The principal instruments of law, in the UK, regarding information risk management are:

- The Data Protection Act (DPA) 2018 and the UK General Data Protection Regulation (UK GDPR), which deal with maintaining the confidentiality and integrity of information but not its availability.
- The Computer Misuse Act 1990, which deals with the criminal offence of unauthorised access to computer systems and the information contained within them.
- The Police and Criminal Evidence Act 1984, together with subsequent addenda, which deals in part with the proper securing of information-based evidence such as computer files.

- The Official Secrets Act 1989, which deals with the disclosure of nationally sensitive information.
- The Freedom of Information Act 2000, which allows requests to be made regarding rights of access to information held by government organisations.
- The Regulation of Investigatory Powers Act (RIPA) 2000, which deals with information that may be collected by governmental organisations in the pursuit of criminal investigations.
- The Copyright, Designs and Patents Act (CDPA) 1988, which defines copyright as a property right that subsists in the following areas: original literary, dramatic, musical or artistic works; sound recordings, films or broadcasts; and the typographical arrangement of published editions.
- The European Union (EU) General Data Protection Regulation, upon which the UK's DPA 2018 is based, is slightly different, but will not be dealt with in this book since the differences do not radically affect it.

All of these relate in one way or another to information risk, and many require that the organisations collecting and holding information must take all reasonable steps to ensure its safety.

Regulation in the information risk management space is less prevalent, although the financial sector does have some regulation regarding risk generally; it is more connected with the management of business risk than information.

Standards and guidelines, however, are available in abundance, and [Appendix H](#) lists the principal publications in this area.

THE CONTEXT OF RISK IN THE ORGANISATION

Any work on information risk in the organisation must begin with an understanding of the organisation's wider view of business risk, which must necessarily examine the impacts or consequences of unexpected events. These can result in any of the following:

- Financial loss, which can include loss of business or IP.
- Legal and regulatory penalties, which can arise from either a breach of regulatory practice or failing to meet regulatory deadlines.
- Reputational damage, which generally begins with adverse reports in the media.
- Damage to the organisation's operations, which may result in subsequent financial or reputational damage.
- Harm to the organisation's staff or the public-at-large, which again can result from damage to the organisation's operations and reputation.

Although many of these are not based on pure finance, the bottom line is that it is mostly about money, since many of the other types of impact will ultimately result in some form of financial loss, whether directly or indirectly.

Information risk is a subset of business risk and relates to the confidentiality, integrity and availability of business information assets and the information management infrastructure, and, although we shall deal with the specific minutiae of impacts to the organisation's information assets, the general principles of business risk management still apply.

Some of the damage to the organisation will be as a result of failures of technology, while other damage will be due to failures to follow policies, processes or procedures, and some will be due to events that simply happen.

A wide range of factors affects the organisation's business risk environment, beginning with generic operational disruptions which affect all organisations, public and private, regardless of sector or size, such as dramatic changes in the economic or political environment; the failure of business transactions that might result from poor management decisions or the failure of parts of the organisation's infrastructure.

Other disruptions that are totally outside the control of any organisation, but which can affect a wide range of organisations, include natural disasters, such as pandemics, flooding and severe weather, terrorism and civil unrest, any of which will disrupt not only normal business operations, but also those of the public-at-large.

Other types of disruption will come within the control of the organisation, and are often sector-specific, especially in the area of hazardous operational environments such as petrochemicals and energy production. Disruptions from failures in business processes and systems will normally come within the remit of business continuity management, which, although linked to information risk management, is a subject area in its own right.

Organisations in certain sectors will also be subject to the dictates of legal and regulatory bodies, where both generic and sector-specific statutory regulations place additional responsibilities on the organisation. In those organisations where products and services fall within the range of hazardous products, the organisation will be subject to additional societal responsibilities under regulations such as the Control of Major Accident Hazards (COMAH), and may also be required to cooperate with emergency responders under the Civil Contingencies Act 2004 in order to provide protection not only for their staff within the working environment, but also for the general public.

The culture of the organisation itself will have a dramatic effect on business risk. The most visible of these in a business risk context is that of the organisation's risk appetite and the internal awareness the organisation has regarding planning for risk. While they sound similar in nature, risk awareness and risk appetite are quite different – awareness meaning that the organisation recognises risk in all its forms, whereas risk appetite means the level of risk that the organisation will accept in any given situation.

Some organisations maintain an extremely low risk appetite; for example, pharmaceutical research organisations take almost no risks at all when it comes to developing a new drug, although it could be said that the potential development costs are a business risk in themselves.

Other organisations thrive on risk – insurance companies and investment brokers being classic examples. This is where risk can be seen as opportunity as opposed to danger.

The reach – local, regional, national, continental or global – of the organisation, together with its business structure and the operational demands it places on its staff, will also be a major contributory factor, and the organisation's hierarchy and reporting channels will define to a great extent the roles and responsibilities of key staff and their accountability for risk.

Very often, those organisations whose operations have a greater degree of urgency will have an increased risk appetite, and may actively encourage staff to take risks within defined limits.

When it comes to information risk, some organisations will maintain extreme secrecy over their entire operations, while others will focus more on information that is either sensitive or confidential.

HOT TOPICS TO CONSIDER IN INFORMATION RISK MANAGEMENT

Since the original version of this book was written, a number of topics have become increasingly relevant and should therefore be considered in their own right in relation to information risk management.

The Internet of Things (IoT)

The IoT is a concept that allows connectivity between multiple physical devices, permitting communication between the digital and physical worlds. A 'Thing' comprises some form of power source; a transceiver to intercommunicate with other 'Things'; sensors, which allow the 'Thing' to take in information; a central processor of some form to make decisions; a storage capability; and, finally, actuators that allow the 'Thing' to control something to which it is connected.

While we generally think of these 'Things' as objects, they do in fact hold or deliver information, and for this reason, we should be mindful that there will always be some form of information risk associated with them.

'Things' vary considerably in shape, size and functionality, but a good example is the Google Protect smoke and carbon dioxide detectors we have on our hall and landing ceilings, together with the Google Nest controller for the hot water and central heating. We can adjust the hot water schedule or heating temperature either by using the Google Nest app on our smartphones, or by speaking to the Google display in the kitchen, and can silence the smoke detector when we overcook the toast in the same way.

The functionality these 'Things' provide is extremely beneficial, but, like most 'Things', they come with default security settings that should be changed once they were installed. However, not everybody has the means or motivation to do this, leaving devices open to attack from digital intruders.

At a personal level, the outcome of an attack could range from embarrassing to costly, while at a medical level, for example, an intrusion could possibly have extremely serious consequences.

Artificial intelligence (AI)

Much of the discussion around AI is regarding its capabilities (both for good and for evil); there has been rather less discussion regarding the information security aspects of AI.

What exactly is AI? Investopedia website,² which publishes financial information, defines it as referring to

the simulation of human intelligence in machines that are programmed to think like humans and mimic their actions. The term may also be applied to any machine that exhibits traits associated with a human mind such as learning and problem-solving.

In theory, this suggests that AI machines may make better judgements, recommendations and decisions than humans, and will invariably make them in a more time-efficient manner.

While this section does not go into the more detailed capabilities of AI, we do need to examine some of the issues raised by its use. Let us consider, then, just one of the many uses to which AI can be put – that of autonomous vehicles.

This is a subject close to my heart, since my wife and I now drive an all-electric car, which (the manufacturer claims) is 'Car 2 Car'-ready,³ and therefore capable of driving itself autonomously as well as communicating with other similarly equipped cars. In fact, this use of AI is yet another example of the IoT.

I should say straight away that given the car's current state of development, I would not dare take both my hands off the steering wheel, having experienced some of the AI functions built into it. However, leaving aside its drawbacks, let's look at some of the information-related security aspects in the Car 2 Car context.

Firstly, there is the confidentiality aspect, in which the so-called 'Infotainment' system is connected by Bluetooth to my smartphone, and therefore has access to all my contacts.

From the integrity point of view, if an attacker was able to gain access to the car's software, as has been reported in *Wired* magazine,⁴ they can – at least in theory – control almost any aspect of the car's performance, the results of which could in fact be life-threatening. Now that the awareness level of this kind of threat has been raised, automotive manufacturers are beginning to take the security aspects very seriously indeed.

From an availability perspective, if the information contained or gathered by the car itself, such as performance, speed, direction and what obstacles are ahead, is missing for some reason (perhaps some form of instrumentation failure), then the car might

² <https://investopedia.com/>.

³ <https://car-2-car.org>.

⁴ 'Hackers Remotely Kill a Jeep on the Highway—With Me in It', WIREd.

not respond as quickly in the way it should, might even respond in a totally incorrect manner or might not respond at all. A sobering thought when a cyclist suddenly pulls in front of it!

The other aspect of the car's information is that of correctness. If we have engaged the automatic cruise control facility, when our car leaves a speed restricted zone (say 40 mph) and moves into a de-restricted zone, it will accelerate up to what it has been programmed is an acceptable speed – say 60 mph. This is fine on a straight road, but very much the opposite when leaving our village, where the road has several sharp bends. Conversely, on entering a de-restricted zone, the car may 'decide' to drop its speed to 30 mph for a few seconds, which comes as something of a surprise to the driver of the vehicle immediately behind!

The manufacturer argues that these software bugs will be ironed out in the next software release (no date given), but our experience of the previous release is that they do not tell us what has changed or what improvements have been made, so we have no idea whether or not the problems have been fixed.

Another example of information security and AI is that of medical diagnosis and treatment. The use of AI in this area is widely hailed as having the potential to eliminate many diseases, and has already contributed greatly to scientists having a better understanding of viruses and how they replicate.

Consider for a moment the use of AI in identifying the size, shape and location of a prostate cancer tumour. The systems that analyse the magnetic resonance imaging (MRI) scans can be used by an AI system to inform the radiotherapy team exactly how to target the tumour while not impacting other vital organs nearby. If the integrity of this information were to be compromised, at best the treatment would be a failure but at worst it could result in further medical issues or worse.

A compromise of availability might mean that tests had to be rerun, causing a delay in treatment, which could well result in a worsening condition for the patient.

It is clear that the information used by AI systems may differ in some respects from routine personal and business information, but its safeguarding must be no less secure, and in some cases there is still a long journey ahead before the algorithms used are truly fit for purpose.

Remote working

Working from a location other than an organisation's office or from an individual's normal workplace – occasionally while mobile – began to be popular some years ago, but since the coronavirus pandemic struck in early 2020, the opportunity to permit staff to work remotely (usually from their home) became a major factor in organisations being able to continue their operations, albeit often at a reduced level.

Many organisations found themselves having to gear up for remote working with little experience of the technology needed to achieve it, and with even less of the information risks that remote working can bring about.

The issues they faced ranged from securing individuals' personal computers (as opposed to company ones), provision of a secure communications infrastructure, storage of non-digital documents and access by other household members whether family or not, to educating the users about how to work remotely and securely. This also included ensuring that both the organisation's central communications hub and the individuals' homes had sufficient capacity and resilience to operate at the required level of performance.

There have been numerous examples in the media of people on videoconferencing calls whose partners appear on camera while inappropriately dressed, participants being unable to be heard or seen through lack of bandwidth, and other distracting noises or images. While some of this requires changes to business processes, there are strong links between them and information risk management.

THE BENEFITS OF TAKING ACCOUNT OF INFORMATION RISK

As we have seen, risk is inherent in any organisation or business, and failure to take account of risk in any context can be disastrous. This is also true of information risk, in which information that is critical to the survival of the organisation must be protected, or the consequences could be severe and the organisation could be subject to the same types of impact or consequence.

Information risk management – a subset of business risk management – addresses these issues in order to prevent them, and after understanding the business context, organisations will identify risks, analyse, evaluate and treat them.

There are two basic actions that can be used – firstly to reduce the likelihood and secondly to reduce the impact or consequence of adverse events. In either case, it is also necessary to limit the possible escalation of events so that matters do not deteriorate once they have begun.

Within the context of information risk management, organisations will need to budget for the prevention of disruptive incidents that would otherwise result in some form of impact, and, in those cases where prevention is either not possible or too costly, to budget for the costs of recovery from them.

The potential benefits to organisations of taking serious account of information risk are manifold:

- There will be an improved view within the organisation of the information assets, their value and the degree to which they are protected.
- There will be a noticeable decrease in the overall level of risk borne by the organisation.
- There may well be a reduction in premiums for those information assets that the organisation insures when transferring or sharing the risk.
- There will be an enhanced view of the organisation in the eyes of its various stakeholders and the media.

- The organisation will be able to respond to and recover from disruptive events more quickly and more effectively.
- There will be reduced levels of impact and loss when unexpected events occur.
- The organisation will be able to claim commercial advantage over those of its competitors that do not follow an information risk management strategy.
- Information risk management is a 'must' for organisations that are seeking to gain accreditation against ISO/IEC 27001:2017 (Information security management systems) and/or ISO 22301:2019 (Business continuity management systems).

Capability Maturity Model

For those organisations that decide to invest seriously in information risk management, there is also the option of gaining additional benefit from following the so-called Capability Maturity Model (CMM),⁵ which can be used for almost any business function including information risk management.

The CMM consists of five levels of capability maturity:

Level 1 – Initial

Processes at this level are typically undocumented and in a state of dynamic change, tending to be driven in an ad hoc, uncontrolled and reactive manner by users or events.

Level 2 – Repeatable

Processes at this level are repeatable, possibly with consistent results. Process discipline is unlikely to be rigorous, but where it exists it may help to ensure that existing processes are maintained during times of stress.

Level 3 – Defined

Processes at this level are defined and documented standard processes established and subject to some degree of improvement over time. These standard processes are used to establish consistency of process performance across the organisation.

Level 4 – Managed

Processes at this level use process metrics. Management can effectively control the process and, in particular, can identify ways to adjust and adapt the process to particular projects without measurable losses of quality or deviations from specifications.

Level 5 – Optimising

It is a characteristic of processes at this level that the focus is on continually improving process performance through both incremental and innovative technological changes/improvements.

⁵ Details of the Capability Maturity Model may be found at: <https://cmminstitute.com>.

ISO/IEC 15504-2:2003 and COBIT 5®

Along similar lines, there is also reference to a measurement framework for process capability in ISO/IEC 15504-2:2003 – Software engineering – Process assessment – Part 2: Performing an assessment. The identical generic process capability attributes appear in COBIT 5:⁶

Level 0 – Incomplete

The process is not implemented, or fails to achieve its purpose. At this level there is little or no evidence of any systematic achievement of the process purpose.

Level 1 – Performed process

The implemented process achieves its process purpose.

Level 2 – Managed process

The previously described Performed process is now implemented in a managed fashion (planned, monitored and adjusted) and its work products are appropriately established, controlled and maintained.

Level 3 – Established process

The previously described Managed process is now implemented using a defined process that is capable of achieving its process outcomes.

Level 4 – Predictable process

The previously described Established process now operates within defined limits to achieve its process outcomes.

Level 5 – Optimising process

The previously described Predictable process is continuously improved to meet relevant current and projected business goals.

OVERVIEW OF THE INFORMATION RISK MANAGEMENT PROCESS

Figure 1.2 illustrates the generic information risk management process, found in a number of standards, including ISO/IEC 27005:2018, ISO/IEC 31000:2018 and ISO/IEC 31010:2019. While being a useful aide-memoire, it does suffer from being rather high level, and fails to show the more detailed steps involved. In later chapters of this book, we shall expand this diagram to explain the steps more fully.


At a very high level, the information risk management process consists of four key steps:

1. The identification and qualification of inherent risk – that is, the risk that an activity would pose if no controls or other mitigating factors were in place.

⁶ Details of the COBIT 5 attributes can be found at: <https://cobitonline.isaca.org/about>.

2. Decision-making regarding the most appropriate form of risk treatment for the risks identified in step 1.
3. The application of suitable controls to achieve the objectives determined in step 2.
4. The acceptance of any residual risk following the implementation of the controls applied in step 3.

Figure 1.2 The overall risk management process



Naturally, this process only scratches the surface of information risk management, and each of these steps is covered in much greater detail in the remaining chapters of this book.

The process itself begins with gaining an understanding of the context in which the organisation finds itself, and includes both the internal context – that is, strategies and policies from within the organisation itself – and the external context, which includes areas such as legal and regulatory constraints and so on. This is dealt with in greater detail in [Chapter 3](#).

Once the organisational context has been established, the process can continue with the risk assessment, which is broken down into three distinct phases: firstly the

identification of the risk, dealt with in [Chapters 4 and 5](#); secondly risk analysis; and third risk evaluation, which are both dealt with in [Chapter 6](#).

Following this, the process takes us into the realm of risk treatment, which is discussed in [Chapter 7](#), with risk reporting and presentation covered in [Chapter 8](#).

At each stage, there are links between the various steps and those of communication and consultation, in which a dialogue is conducted with major stakeholders, and finally also with the process of monitoring and review, both of which are covered in [Chapter 9](#).

SUMMARY

Having now examined the more general aspects of information, hot topics, the capability maturity model and the information risk management process, we can move on to the remaining chapters in this book, which are organised as follows:

Chapter 2 – Review of information security fundamentals

This chapter includes a review of the basic concepts of information security fundamentals, the process of information classification and the Plan-Do-Check-Act model.

Chapter 3 – The information risk management programme

This chapter deals with the goals, scope and objectives, roles and responsibilities and governance of an information risk management programme, and information risk management criteria.

Chapter 4 – Risk identification

In this chapter, we deal with the approach to risk identification, how information assets and their owners are identified, how a business impact assessment (BIA) is conducted and the types of impact we might encounter, and discuss the pros and cons of qualitative and quantitative assessments.

Chapter 5 – Threat and vulnerability assessment

In this chapter, we describe how threat and vulnerability assessments are carried out and also examine the view of existing controls.

Chapter 6 – Risk analysis and evaluation

In this chapter, we cover the process of assessing the likelihood of risks arising and, by combining the likelihood with the impacts or consequences of a threat, calculate the relative levels of risk for each threat type. We then examine how the risk matrix is developed and evaluate the risks in terms of priority.

Chapter 7 – Risk treatment

This chapter discusses the approach to making risk treatment plans, and describes the four strategic, four tactical and three operational risk treatment options.

Chapter 8 – Risk reporting and presentation

In this chapter, we describe how to report and present the findings of the risk assessment process and explain the need for robust business cases.

Chapter 9 – Communication, consultation, monitoring and review

This chapter includes details of the importance of consulting with stakeholders throughout the entire risk management process, and with the process of monitoring and reviewing the work undertaken and how the risk management programme should continue.

Chapter 10 – The NCSC Certified Professional Scheme

In this chapter, we describe the National Cyber Security Centre Certified Cyber Professional (CCP) scheme, the Skills Framework for the Information Age (SFIA) levels and the Institute of Information Security Professionals (IISP) skills framework upon which the CCP scheme is largely based.

Chapter 11 – HMG security-related documents

This final chapter provides a detailed summary of the UK government approach to information risk management, and includes descriptions of:

- Her Majesty's Government (HMG) Security Policy Framework – the security of information;
- the UK government security classifications.

Appendix A – Taxonomies and descriptions

In this appendix, we provide two useful taxonomies that can be used in information risk management:

- information risk;
- typical impacts or consequences.

Appendix B – Typical threats and hazards

Here, we discuss various types of threat and hazard, including malware threats, physical threats, misuse and abuse threats, social engineering threats, hacking threats, environmental hazards and, finally, threats caused by errors and failures.

Appendix C – Typical vulnerabilities

In this appendix, we examine various types of vulnerability, including those of communications and operations, people-related vulnerabilities, access control vulnerabilities, systems acquisition, development and maintenance vulnerabilities and physical and environmental vulnerabilities.

Appendix D – Information risk controls

In this appendix, we look at the three levels of controls, beginning with strategic controls – avoid or terminate, reduce or modify, transfer or share and accept or tolerate. We then move to the tactical level, which includes detective, directive, corrective and preventative controls. Finally we examine the operational level controls – procedural controls, physical controls and technical controls.

Appendix E – Methodologies, guidelines and tools

In this appendix, we provide a brief description of some of the more popular information risk management methodologies:

- CORAS;
- FAIR;
- OCTAVE;
- SABSA.

Appendix F – Templates

In this appendix, we provide a number of useful templates and guidance information that can be used in the information risk management programme:

- impact assessment template;
- threat/hazard assessment template;
- vulnerability assessment template;
- existing controls assessment template;
- risk register template.

Appendix G – HMG cybersecurity guidelines

This appendix examines the main UK government guidelines for cybersecurity, including the:

- HMG Cyber Essentials Scheme;
- 10 Steps to Cyber Security.

Appendix H – References and further reading

This appendix provides the reader with a large number of information sources:

- primary UK legislation;
- good practice guidelines;
- other reference material;
- NCSC Certified Cyber Professional Scheme;
- other UK government publications;
- risk management methodologies;
- UK and international standards.

Appendix I – Definitions, standards and glossary of terms

The final appendix provides the reader with a summary of the definitions, standards and glossary used throughout the book.

2 REVIEW OF INFORMATION SECURITY FUNDAMENTALS

Let us now take a brief look back at the fundamental concepts of information security, as it is these that will form the basis of the risk assessment process itself.

We shall then examine the means by which information is classified and labelled, and how the Plan-Do-Check-Act methodology may be used as a high-level process for information risk management.

It is a widely held belief that the three main pillars of information security are confidentiality, integrity and availability, often referred to simply as 'CIA'. While this is essentially true, other factors also contribute to the overall scheme of things. Accountability, authenticity, non-repudiation and reliability are all contributing factors, and need to be considered along with the 'main' three.

Let's take a look at some definitions and explanations of these, together with those for information assurance, information governance and data governance.

Confidentiality – 'the property that information is not made available or disclosed to unauthorised individuals, entities or processes' (ISO/IEC 27000:2018). Confidentiality is concerned with ensuring that information is available to authorised entities and is not allowed to become available to unauthorised entities, whether they are able to obtain this deliberately or by accident. It follows, therefore, that users should only have as much access as they require in order to carry out their task and that a formal process is required in order to administer access rights.

PRIVACY AND SECRECY

Both the terms 'private' and 'secret' have the same basic meaning, but whereas privacy generally indicates the need to protect an individual's information, secrecy can be seen to have a darker side and can indicate a more sinister motive.

Integrity – 'the property of accuracy and completeness' (ISO/IEC 27000:2018). While this definition is fine as far as it goes, the term 'integrity' also suggests a high degree of reliability and assurance, and can apply equally to people as well as to information. Integrity considers both the completeness and accuracy of the information, and as with confidentiality, users should only have as much access as

they require in order to carry out their task and that a formal process is required in order to administer access rights.

At best, integrity failures can lead to misinterpretation or poor decision-making; at worst they can lead to serious financial impact and embarrassment to the organisation.

Availability – ‘the property of being accessible and usable upon demand by an authorised entity’ (ISO/IEC 27000:2018). Availability is often considered the poor relation of CIA, and, while the other two are very important, if information is not available then it becomes frustrating to those who require access to it at the time they require it, and under certain circumstances this can have extremely severe consequences.

Availability is now a critical element in the delivery or provision of information, not only to customers who shop online at any hour of the day or night, but also to multinational organisations operating across multiple time zone boundaries.

Also – a business continuity (BC) issue – the tolerable length of time for which any information asset is unavailable may well vary from one organisation to another, and indeed from one service to another.

Non-repudiation – ‘the ability to prove the occurrence of a claimed event or action and its originating entities’ (ISO/IEC 27000:2018). Non-repudiation can be used both to prove not only that an entity has carried out a certain action but also equally that an entity has not carried out an action, whether this be carrying out a commercial transaction, editing a document or sending an email. An example of non-repudiation is the use of digital signatures and certificates, which establish the identity of an individual beyond all reasonable doubt.

Authentication – ‘the provision of assurance that a claimed characteristic of an entity is correct’ (ISO/IEC 27000:2018). In order to ensure both confidentiality and integrity, authentication mechanisms are used to validate an entity’s credentials – this can be either an individual or an application requiring access to information or applications. Authentication mechanisms include such things as passwords, fingerprint and iris scanning and token generators.

Identification – this is a mechanism by which an entity begins the process of authentication. It may refer to systems, peripherals, people or processes. For example, a user may submit his or her identification in the form of a user ID when logging on to a system or application.

Accountability – ‘the assignment of actions and decisions to an entity’ (ISO/IEC 27000:2012 – for some reason, the term ‘accountability’ has been omitted from the ISO/IEC 27000:2018 version). Accountability is often confused with responsibility. The two are very different – an entity may be made responsible for carrying out an action, for example an engineer may be responsible for configuring firewall rules, whereas a more senior manager is likely to be accountable for the firewall and/or its rule set, and may be held to account if things go wrong.

Accountability is also linked to non-repudiation, in that it may be desirable to correlate transactions with individuals or processes.

Reliability – ‘property of consistent intended behaviour and results’ (ISO/IEC 27000:2018). Reliability has similar connotations to integrity, but whereas integrity

refers mainly to ensuring accuracy and completeness, reliability leans more towards something that can be repeated with accuracy, for example a process that works in a consistent manner every time.

Information assurance – information assurance is the practice of assuring information and managing risks related to the use, processing, storage and transmission of information or data and the systems and processes used for those purposes.

Information assurance includes protection of the confidentiality, integrity, availability, authenticity and non-repudiation of information. It uses physical, technical and procedural controls to accomplish these tasks.

While focused predominantly on information in digital form, the full range of information assurance encompasses not only digital information but also analogue or physical information. Protection applies to information in transit, both in physical and electronic forms as well as information at rest in various types of physical and electronic storage facilities. Information assurance as a subject area has grown from the practice of information security.

Information governance – information governance is the set of multi-disciplinary structures, policies, procedures, processes and controls implemented to manage information at an enterprise level, supporting an organisation's immediate and future regulatory, legal, risk, environmental and operational requirements.

Data governance – data governance refers to the general management of key data resources in a company or organisation. This broad term encompasses elements of data use, storage and maintenance, including security issues and the way data flow from one point to another in an overall information technology architecture.

Because data or raw information are a key resource for most businesses and organisations, data governance is a logical area of overall information technology strategy focus for many large enterprises.

INFORMATION CLASSIFICATION

All information assets have some degree of value to the organisation. Unless users of these understand their sensitivity and how to deal with them, they could unwittingly – or even deliberately – make them available to unauthorised or unsuitable recipients to the detriment of the organisation or themselves. So when dealing with either raw data or processed information, whether it is our own or someone else's, it is vital to ensure that users of these understand fully how to access, process, store, transmit, transport and (if necessary) ultimately destroy them. This is otherwise referred to as data or information handling.

To provide these handling specifications for each individual item of data or information would be an enormous task, so in order to simplify matters we first classify each data or information item according to a set of rules, which will then allow us to specify the handling procedures for each type.

Within government circles this has been undertaken for many years and is a well-established process. In the private sector, however, although organisational data or information handling guidelines may exist, they are not always rigorously enforced, and in some sectors organisations that do not adequately classify and protect certain types

of information may face regulatory penalties. Further, any organisation wishing to attain an accreditation relating to information risk will have to satisfy the accreditor that due diligence has been undertaken and that information has been classified appropriately.

The term 'privacy marking' is also used in connection with this topic, but differs from information classification in one critical respect – privacy marking deals solely with the labels applied to the information, whereas information classification deals with the privacy marking and the handling of information.

Information classification includes all forms of media, whether in storage (at rest) or in transit from one location to another, such as:

- magnetic media, including hard disks, USB sticks and magnetic tape, locally or in the cloud;
- PDAs, tablet computers, mobile phones and digital cameras;
- optical media, including CDs, DVDs and microfiche;
- paper, including handwritten notes, printed files, diagrams and plans;
- information passing across both wireless and wired networks, including telephone calls, video calls and facsimile transmissions;
- email, text messages and related social media such as Facebook, Instagram, LinkedIn and Twitter.

The value of information assets to the organisation or individual is not necessarily limited to their commercial value, but also the impact they could have on the organisation or individuals were they to become known to an attacker, a competitor, a hostile state or the public-at-large.

All information assets must be identified and rated in value against an agreed impact system – a form of risk assessment in itself. The UK government updated its information classification scheme in 2018, which greatly simplifies the previous system:

- TOP SECRET – HMG's most sensitive information requiring the highest levels of protection from the most serious threats; for example, where compromise could cause widespread loss of life or else threaten the security or economic wellbeing of the country or friendly nations.
- SECRET – very sensitive information that justifies heightened protective measures to defend against determined and highly capable threat actors; for example, where compromise could seriously damage military capabilities, international relations or the investigation of serious organised crime.
- OFFICIAL – the majority of information that is created or processed by the public sector. This includes routine business operations and services, some of which could have damaging consequences if lost, stolen or published in the media, but are not subject to a heightened threat profile.

Commercial organisations, on the other hand, may have a system such as:

- Strictly Confidential – information that the loss or damage of which could cause extremely serious financial impact or embarrassment to the organisation. This

might include future business plans, future product development information or information that might have an adverse effect on the organisation's share value.

- Confidential – the loss or damage of which could cause some financial impact or embarrassment to the organisation.
- Personal – the loss or damage of which could cause some financial impact or embarrassment to one or more individuals within the organisation, and could have regulatory repercussions on the organisation.
- Internal use only – the impact of which might be low, but could be aggregated with other information for use by a competitor.
- Public – which can be made available to any person or organisation.

In order to assign privacy markings, the concepts of confidentiality, integrity and availability must be taken into account. For example, any information asset labelled as Strictly Confidential would almost certainly have a very high degree of all three, whereas Public information need only have a certain degree of integrity and availability.

In terms of confidentiality, the most frequently used guideline is referred to as the 'need-to-know' principle – information should not be made available to people who do not need to know it. Integrity is often addressed by segregation or separation of duties, so that one person might generate information, but in order for it to be made available it may need to be verified by a second person. Availability is most frequently addressed by the use of backups, disaster recovery (DR) and BC plans, processes and procedures.

In addition to these privacy markings, information can also be assigned caveats, known in government circles as descriptors. These are additional attributes that ensure a finer layer of granularity. Some examples are:

- Human resources (HR) only – referring to personnel files containing sensitive personal data.
- XXX Project Team only – not to be shown to anyone who is not a member of a particular project team.
- Not for general release until xxxx – not to be further distributed until a certain date.

Another information classification scheme has become popular in recent years, which emanates from the European Network and Information Security Agency (ENISA), known as the 'Traffic Light Protocol'.¹ See <https://enisa.europa.eu/topics/csirts-in-europe/glossary/considerations-on-the-traffic-light-protocol>.

- RED – Personal for Named Recipients Only – in the context of a meeting, for example, distribution of RED information is limited to those present at the meeting, and in most circumstances will be passed verbally or in person.
- AMBER – Limited Distribution – recipients may share AMBER information with others within their organisation, but only on a 'need-to-know' basis. The originator may be expected to specify the intended limits of that sharing.

¹ See <https://enisa.europa.eu/topics/csirts-in-europe/glossary/considerations-on-the-traffic-light-protocol>.

- GREEN – Community-Wide – information in this category can be circulated widely within a particular community or organisation. However, the information may not be published or posted on the internet, nor released outside the community.
- WHITE – Unlimited – subject to standard copyright rules, WHITE information may be distributed freely and without restriction.

Once information assets have been identified, it will be necessary to match each of them (or groups of similar information assets) against an information owner. There must then be a process in which the security classification of the information assets are verified through interviews with the information owners.

Handling of information assets

Once the security classification scheme has been established, thought must be given to how the information asset is handled.

Creation and storage of an information asset

The originator or creator of any information asset should consider assigning its security classification immediately, especially if the information is of a sensitive nature. Even if the information asset is in draft form – for example an early version of a project plan, design or simple document – it should be stored in the most appropriate manner.

Not only should the item be stored in a secure location, but it might also be necessary to password protect the item as an additional means of securing it, or by encrypting the item when stored.

Sharing and review of an information asset

Once an information asset has been created, it is possible that other people will review it; for example, a draft project plan might require input from a number of team members, each of whom may need to view and update the plan. This brings in another level of protection – that of the item's security attributes, and the ability of individuals to read from and write to the item.

If multiple people are able to access the item simultaneously, there needs to be a 'lockout' mechanism to prevent more than one person trying to edit the item at the same time.

For this reason, and in order to minimise the possibility of an item going astray, sharing is better achieved by allowing controlled, shared access to it rather than by sending it by email, for example.

Transmission of an information asset

At times, it will be necessary to transmit the information to another person, and the security of the information during transit must be considered. Depending upon the sensitivity of the information, it may be possible to transmit it 'in clear' over a public network such as the internet; a virtual private network (VPN); or a heavily secured

private network. Some information may be required to be encrypted when transmitted, or it may have to be hand carried by courier in a secure container.

Disposal of an information asset

Most information assets will have some kind of life expectancy, and once this point in time has been reached, it may be necessary or desirable to dispose of the asset rather than storing it indefinitely. Suitable methods of destruction will depend, as always, on the sensitivity of the information, and may range from simple file deletion for an unclassified asset to physical destruction of the platters of a hard disk drive or from shredding to burning of paper documents.


The main point for any system of information classification is that once an information asset has been given a security classification, it automatically imposes constraints on the methods that can be used to process, store, transmit and ultimately dispose of it. These conditions must inevitably be imposed on anyone who may come into contact with that information asset.

Because the nature and sensitivity of information assets may change over time, a periodic review of information assets and their classifications is essential.

PLAN-DO-CHECK-ACT

Strictly speaking, although it is not an information risk topic, for many years, and for a variety of purposes, organisations have made use of a system known as the Plan-Do-Check-Act or the Plan-Do-Study-Act (PDCA or PDSA) cycle, otherwise known as the Deming Cycle,² illustrated in Figure 2.1.

Figure 2.1 The Plan-Do-Check-Act cycle



² See <https://deming.org/explore/pdsa>.

The PDCA cycle has been widely adopted as a basic reference framework in information security, information risk management and business continuity management disciplines as well as many others.

The four stages are described as follows:

PLAN

In this stage, we establish the objectives and the processes necessary to deliver the required results. In the information risk management context, this equates to understanding the organisation and its context.

DO

The next stage of the process implements the plan, initially as a means of testing that the plan has been successful. In the information risk management context, this equates to implementation of the information risk management framework.

CHECK/STUDY

In this stage, we examine the results we have achieved either by measurement or observation. In the information risk management context, this equates to monitoring and review of the framework.

ACT

In the final stage, we put the validated plans into action when an incident occurs and bring lessons learnt from incidents into revision of the plan. In the information risk management context, this equates to continual improvement of the framework.

Earlier British and international standards made considerable use of this cycle in their introductory sections, but since 2013 their use of the PDCA cycle has diminished inexplicably.

SUMMARY

In this chapter, we have examined information security fundamentals, information classification and the Plan-Do-Check-Act cycle. The next chapter will cover the overall information risk management programme.

3 THE INFORMATION RISK MANAGEMENT PROGRAMME

Due to its possible scale, for many organisations risk management will involve a number of areas of work rather than simply a project, and, while the mechanics of managing information risk are relatively straightforward, there needs to be an overall framework around the activity if there is to be any real chance of success.

This chapter discusses the goals, scope and objectives of such a programme, together with the various roles and responsibilities and governance of the programme.

The organisation should ideally establish an information risk management programme, which will have oversight of all the work. Such a programme might contain the following elements:

- The goals, scope and objectives of the programme and the organisation's overall information risk management policy.
- The overall roles and responsibilities of the programme leaders and key players, including their areas of authority, ownership and accountability.
- The governance processes for the programme, and, if necessary, those for the individual components or projects within the programme.
- The internal standards that must be observed, including risk criteria, reporting, documentation and management processes.
- The financial arrangements, including budgetary constraints if these are known.
- Training of those involved in the programme and awareness for staff generally.
- Monitoring and review of progress and results.

Many organisations make use of the PDCA model, which was described in greater detail in the previous chapter. PDCA is a useful method in the management of any project or programme, but although it has featured in many UK and international standards, more recently published standards have omitted it.

The PDCA model takes the view that, although an information risk management programme may have defined start and finish points, it is in fact a continuous process, and that organisations should revisit all risks on a regular basis or when any facet changes. This highlights the need for the integration of information risk management into business-as-usual operations.

GOALS, SCOPE AND OBJECTIVES

The organisation's ultimate goal might be to obtain ISO/IEC 27001 accreditation, and an effective information risk management programme will be an essential component of this. Alternatively, it may just be the case that the organisation wishes to establish as risk-free an information environment as possible. If the former, the accreditor will be seeking evidence not only of the outcomes of such a programme, but also the means by which it has been executed and its ongoing monitoring. If the latter, it may still be worthwhile for the organisation to employ the services of an auditor or accreditor simply to verify that the programme has been conducted thoroughly and completely.

In developing the strategic approach to the information risk management programme, the team will be required to establish both the internal and external contexts in which the organisation operates and how the information risk management process fits in to the overall business environment.

A key aspect of this is the requirement to define, document and agree with senior management the organisation's risk appetite and its criteria for accepting risks that cannot be treated by other means as well as for accepting residual risk.

Many of the drivers for the information risk management programme will originate from the organisation's existing information security policies, if they exist. If not, they must be developed as an integral part of the programme. Many of the controls applied in order to treat the risks identified will involve IT and the security team.

It is crucial, therefore, that the information risk management programme is not viewed as being a stand-alone or separate programme from that of the information security community, and that constant communication and consultation takes place between the two disciplines wherever their management structure differs.

Some of the requirements identified by the information risk management programme will originate from the organisation's legal and regulatory department, who will therefore be heavily involved; and there will likewise be the need to ensure communications and consultation between stakeholders at all levels, both within the organisation and outside it where necessary.

Setting the programme scope

While the strategy of the information risk management programme sets out the goals and objectives of the programme, it is also essential to go down a level and set the scope:

- Those elements of the organisation's information assets that are to be within the scope of the programme.
- Equally importantly, those elements of the organisation's information assets that are to remain outside the scope of the programme.

ROLES AND RESPONSIBILITIES

No programme can be successful unless it has strong leadership, and in setting the overall roles and responsibilities for the programme, the organisation's senior management team must bear in mind that, in addition to leaders, the programme requires dedicated assignment of resources to undertake the detailed work.

Managers of these members of staff must be made aware of the programme and understand that people will need to take time from other duties in order to contribute to the programme, in which they might report to another manager, and that the staff themselves may require additional training.

GOVERNANCE OF THE RISK MANAGEMENT PROGRAMME

Within the overall framework of governance of risk management, there will be three distinct layers of involvement. At the strategic layer, there will exist the overall responsibility, accountability and authority for the programme, some if not all of which will lie at board level. The designated board member or members should ensure that the tactical and operational work is understood, and that the organisation's business and cultural contexts are taken into account.

Legal and regulatory compliance issues can become a complex subject in their own right, especially in cases where organisations are spread across multiple legal and regulatory jurisdictions. The organisation's legal and regulatory department must identify all necessary obligations to the programme, and ensure that these are complied with. However, it will also be necessary to maintain oversight of the legal and regulatory liabilities that exist, since the costs of achieving these may have an impact on the costs of the programme.

In parallel with the identification of any work packages, the nomination of key individuals having specific roles within the programme will have a profound effect on its outcome, and the board will need to ensure that these individuals' performance is monitored and reviewed, and that their terms of reference are verified at intervals to ensure that the programme's objectives are being achieved.

The board should ideally have the information risk management programme as a running agenda item, as it may well be a component part of the organisation's annual reporting, especially if it is in a highly regulated sector.

At the tactical layer, a slightly lower-level view is required, and the adoption of risk intelligence procedures enable the organisation to discover the existence of risks that have either not yet been taken into account – these might be obtained from both business and information security sources – or those that have occurred previously within the organisation. In some contexts, and especially in terms of BC, this is known as horizon scanning.

Also within the tactical layer of governance is risk policy management, in which a strategic steer from board level is translated into the day-to-day policies that must be

followed by the organisation. This includes an overall policy framework and the general format of the policies, together with their interdependencies where these exist.

Finally, at the operational layer, there will be the key activities of the information risk management programme: those of risk assessment, including the identification of threats, vulnerabilities and impacts or consequences, the formulation of the likelihood and subsequent analysis of the risks and, finally, the evaluation of risks and the proposals for risk treatment.

A final element of the overall programme governance will be the need for regular communications and reporting both upwards and downwards through the chain of command, especially in the reporting and logging of new risks and in progress in the treatment of existing risks.

INFORMATION RISK MANAGEMENT CRITERIA

In support of the information risk management programme, and to guide its internal standards, a number of business risk management requirements will be needed.

The legal and regulatory framework in which the organisation operates, both within its host country and other jurisdictions, will have a major impact on the standards and criteria adopted. These will include so-called 'primary legislation' – the laws of the country concerned, such as the Computer Misuse Act, and the data protection legislation and secondary regulation that is generally sector-specific.

The nature of business within the sector itself will have some influence on standards and criteria, as well as the way in which the organisation is structured, both organisationally and geographically.

High-level business risk estimation

In order to provide a starting point for the later risk assessment work, the organisation may benefit from producing a series of high-level estimates of business risk. These need not be very specific or accurate, since they are intended only as a 'starter for 10', but might include such areas as the possible losses that might be incurred by the business through being unable to answer customer calls or by being unable to reach a minimum regulatory threshold for some reason.

Important at this stage is the 'what', rather than the 'why', since it will lead the later stages of the work into a more detailed impact analysis and provide the analysts with a number of high-level headings with which to begin their work.

Risk appetite

In order to proceed with the process of information risk management, the organisation must commence by setting its risk appetite. Unfortunately, this is not a one-off exercise, as each type or class of information asset may have a different risk appetite associated with it.

The following factors will determine the risk appetite for each type or class of information:

- the information's classification;
- the information's confidentiality, integrity and availability requirements;
- the organisation's sector type;
- the organisation's culture;
- the organisation's legal and regulatory obligations.

We'll deal with impact and likelihood scales in greater detail in the next chapter, but for the meantime, it is worthwhile understanding that the terms 'low', 'medium' and 'high' are qualitative as they stand and are therefore not completely meaningful, and where possible should be placed in what is often referred to as a semi-quantitative context. So, for example, 'low' might refer to a range of values up to £100,000, 'medium' to a range between £100,000 and £1 million and 'high' to a range of values greater than £1 million.

This still does not dictate the exact risk appetite, but it does provide the assessor with reasonably objective guidelines as opposed to less meaningful subjective ones. Naturally, the level of granularity of the ranges can be increased if desired, but as the level of granularity increases this brings about a more complex assessment process that takes longer to achieve.

Risk treatment criteria

At the strategic level of risk treatment, there are four basic options:

- risk avoidance or termination;
- risk reduction or modification;
- risk transfer or sharing;
- risk acceptance or tolerance.

Below this are the tactical and operational levels of risk treatment, which we will deal with in greater detail in [Chapter 7](#), but, for the moment, we will examine these four in a little greater depth.

There are also several key factors that influence the decision as to which course of action is most appropriate:

- Whether the choice is actually achievable. For example, it may not be possible to take out an insurance policy against the prospect of a fine for violating data protection law.
- Whether the choice brings about additional risk. For example, if the organisation decides not to enter into a new development programme, there may be consequential losses incurred by not doing so.

- Whether multiple choices are appropriate. For example, treating a particular risk might involve halting one part of a business operation, insuring against a capital loss and introducing additional procedures to reduce the likelihood.

Whatever the choice made, it should be understood that there may always be some residual risk, regardless of the effectiveness of the actions taken, and this residual risk will have to be accepted by the organisation, recorded as such and subjected to ongoing monitoring and review. Further, the process for recommending the choice or choices should be according to defined criteria, and should follow a consultative process to ensure both consistency and fairness.

Risk avoidance or termination criteria

These can be quite difficult to define, partly due to the subjective nature of some of the inputs to the decision-making process. Trying to overcome this subjectivity can be a time consuming and expensive process in its own right. For example, if an organisation wishes to undertake a particular activity, but does not feel it possesses the skills and expertise to do so, it may be possible to outsource the work to another organisation. However, without a detailed financial analysis, it might not be clear as to whether the cost of this approach would be greater or less than the losses incurred by not undertaking the activity at all.

Alternatively, the organisation may feel that the risk is so high that the possible costs of treatment would be unacceptable. Again, without further detailed financial analysis, the decision to avoid the risk becomes subjective and subject to a large degree of uncertainty.

The only objective factor driving the decision to avoid or terminate a risk is when the organisation is fully aware without further analysis that the costs of treating the risk exceed the possible impact of not treating it.

Taking the route of risk avoidance will normally remove both the impact and the likelihood of the risk, but may result in some form of consequential risk caused by not undertaking the activity.

Risk reduction or modification criteria

This option allows us to reduce either the impact or the likelihood of the risk, and possibly even both. Risk reduction, however, does not imply that the risk is reduced to an acceptable level (as determined by the organisation's risk appetite), but merely that it has been reduced to some degree. As we mentioned earlier, it may be necessary to use several different forms of risk reduction, or use them in combination with other types of risk treatment.

The decision to reduce or modify a risk will be based on whether or not the costs of doing so are above or below the level set by the organisation's risk appetite for the particular information asset and whether the organisation possesses the skills and expertise to do so from within.

Risk transfer or sharing criteria

In contrast to risk reduction, risk transfer can only ever reduce the impact of a risk, but never the likelihood. In transferring the risk to a third party, the organisation can only

transfer the treatment of the risk – the ownership must remain with the organisation. A good example of this is the case of the BP *Deepwater Horizon* oil spill in 2010, in which almost 5 million barrels of crude oil were discharged into the Gulf of Mexico with disastrous results to the ecology of the region. Although the oil rig was operated by a third-party organisation, the US government held BP responsible for the incident.

Transferring a risk can usually mean insuring against it, but can also refer to outsourcing arrangements, especially of information technology hardware and software and also of information security management.

Transferring risk will have up-front costs (premiums in the case of insurance) and may also have downstream costs – for example, there may be an excess penalty to pay in the event of a claim, and also the policy payment may not fully cover the cost of replacement, repair or recovery if certain exclusions apply through circumstances in force when the risk event takes place.

Finally, transferring the risk depends both upon the availability or willingness of a third party prepared to take on the risk, since some risks are not insurable, and the usual constraint of whether the potential losses exceed the potential costs.

Risk acceptance or tolerance criteria

The final choice for risk treatment is that of accepting or tolerating the risk. This must always be done knowingly and objectively, and the residual risk must always be monitored in case either the impact or the likelihood changes with time. Ignoring a risk is never an option, since, although it may be very low at one point in time, either the possible impact or likelihood could increase dramatically or gradually, with the result that it becomes necessary to take an alternative (and frequently more costly) approach to treat the risk.

Accepting risks does not alter either the impact or the likelihood of the risk occurring, and will generally be the option when the costs of treating the risk are greater than the potential losses that might be incurred.

Costs of risk treatment

In a later stage of the information risk management programme, a list of recommendations will usually be presented to senior management for their consideration before risk treatment commences. Some of these recommendations will require significant financial investment in order to fully treat the risks identified. For example, the risk of a key system becoming unavailable might be so high that a decision is made to treat it by providing a high-availability standby system.

The magnitude of cost incurred by a project such as this would be very significant, and therefore the organisation's senior management might well request that a full business case be provided, and that a financial threshold is set as an additional criterion for the information risk management programme.

Training In organisations that are highly developed in terms of capability, the level of training required by staff in the process of information risk management may not be great. However, in those organisations that are less experienced in this kind of work,

training of staff at all levels – strategic, tactical and operational – may be a necessary preliminary to the programme.

Many skills are relatively easily learnt, and can be acquired on readily available industry training courses. Others, however, especially in the legal and regulatory domain, may not be so straightforward to acquire and will need time to develop, and will possibly require considerable mentoring before staff are fully proficient.

The main point, of course, is that the overall information risk management programme must necessarily include an element of training and development in order to ensure its success.

Communication and consultation From the very beginning, the information risk management programme will require a high degree of communication up and down the organisation, and the need for consultation, particularly in the early stages of the programme, cannot be overstated.

It is a common mistake for inexperienced information risk managers to make broad assumptions regarding the value of assets, the impacts on the organisation of the loss or damage to those assets, the threats and hazards faced by the assets and the vulnerabilities they exhibit.

The information risk manager should strive to consult at every stage of the programme, and resist the temptation to make rash assumptions, the consequences of which can be highly detrimental.

At an early stage in the programme, the information risk manager should take great care to identify all those who are directly responsible for the information assets in question – or who are able to take an objective view of impacts and consequences, threats or hazards and vulnerabilities – to make contact with them as soon as possible and to maintain that dialogue throughout the programme.

Another common error is to assume that these 'subject matter experts' will either be aware of the programme, its importance to the organisation or the likely involvement they may have in it. While it might be acceptable to fire off a quick 'heads up' email to someone whom the information risk manager knows well, this might be less appropriate for others, and it is strongly recommended that contact should be established on a personal level before making regular use of email to exchange information.

Monitoring and review The final piece in the overall information risk management programme puzzle is that of monitoring and regular review of the entire activity.

It may be useful for the organisation to define some basic metrics for the purpose of monitoring progress. However, care should be taken in being too general in this approach, since the risk management for some risks will take quite a short time, while that for others may take significantly longer. It is perhaps better to report on individual areas in terms of percentage completeness, and then to combine these individual amounts to provide an overall status.

Whichever approach is taken, it should be clearly defined and consistently applied, so that different teams report their own activities in the same way as all others, and that the senior management team do not receive a skewed view of overall progress.

Occasionally, risks will emerge that, despite the best efforts of the risk management team, appear to have reached an impasse and no clear indication can be made as to how – or whether – to treat them. Risks such as these should be flagged to the organisation's senior management team at the earliest possible moment in order to ensure that they do not become overlooked simply because they are too difficult to deal with.

The overall owner and person responsible and accountable for the information risk management programme should monitor its progress on a frequent and regular basis, should make a point of reviewing all risks addressed by the programme whether they have been treated or not, and should establish that their treatment is on track.

SUMMARY

In this chapter, we have discussed the goals and objectives of an information risk management programme, including the scope and governance as well as the criteria for conducting the programme. We then examined the overall process and the basics of the information risk management programme, so now it is time to move on to its first stage – that of risk identification.

4 RISK IDENTIFICATION

The first stage of the risk assessment process is that of risk identification, the purpose of which is to determine the threats or hazards that could cause loss or damage to an information asset, to identify any vulnerabilities exhibited by the information asset and to determine the possible impact or consequences to the information asset.

Regardless of whether or not the risks identified fall within the remit of the organisation, they must be included in the assessment, even though the root cause may remain hidden.

Just to recap:

- The **impact** or **consequence** of a threat or hazard acting on an information asset is the result of that threat or hazard taking advantage of one or more vulnerabilities that are present within the information asset.
- The **likelihood** or **probability** of the threat or hazard succeeding in this depends on the type of threat and any vulnerabilities exhibited by the information asset.
- The **risk** is the combination of the impact or consequence on the information asset combined with the likelihood or probability of the threat or hazard successfully taking place.


Looking slightly deeper into this, we should also take into account the motivation for the attack for certain types of threat – where an attacker either wishes to obtain or change information (confidentiality or integrity) or to deny access to information asset (availability).

In the diagram at [Figure 4.1](#), we can see that risk is the result of combining impact and likelihood. The impact is determined by a threat exploiting a vulnerability within an information asset, and that the presence of both threat and vulnerability give rise to the likelihood.

Now let us take a quick look at the more detailed process of risk identification.

THE RISK IDENTIFICATION PROCESS

Risk identification begins with identifying the information assets that are relevant to the organisation. This is almost certainly the most crucial part of the whole process – failing to identify an asset at this stage will mean it is never risk assessed, and it is vital that, along with the asset itself, an asset owner is identified. For some assets, this will be immediately obvious, whereas for others a senior management decision may be required in order to allocate ownership of the information asset to a suitable person or team.

Figure 4.1 A general view of the risk environment

Once the information assets have been identified, the asset owners can identify the impact or consequences of their damage, loss or destruction. This part of risk identification is commonly referred to as impact assessment.


Let us take a look firstly at the types of information we might identify. Mostly, these days, we think of information being in electronic form, but it is important to remember other information that remains in the form of paper documents, microfiche, microfilm and so on. Also, while much of the information we are concerned with will be alphanumeric in nature, there may be other types that are critical to the organisation, such as software, audio and video files, photographs or 'raw' data, such as that recorded from seismographic surveys and weather observations. [Figure 4.2](#) illustrates many types of information, but individual organisations will doubtless identify others that are peculiar to their sector or business type, and which must also be considered.

The main factor that will help to identify those information assets that must be within the scope of the programme is a list of business-critical activities. However, even though this will identify many of the organisation's information assets, the business-critical activities themselves should be subjected to some form of analysis to take into account other non-critical activities upon which the business-critical activities have a greater or lesser degree of dependence.

For most organisations, the main areas of information assets may include:

- Operational information that underpins the very nature of the organisation, such as customer order history, stock control or product designs.
- Personal information, as defined within data protection legislation such as names, postal and email addresses and telephone numbers.
- Strategic information, such as sales forecasts, development schedules or product launch plans.
- Information that is expensive to collect, store or process, such as census information.

It is unlikely that information that falls outside these categories will be critical to an organisation's performance, and it may be omitted from the information identification scope with the agreement of the relevant information owners, and should be recorded as such.

Figure 4.2 Typical types of information asset

THE APPROACH TO RISK IDENTIFICATION

There are several different skills required in order to conduct a successful impact assessment. These are described in greater detail in [Chapter 9](#) – Communication, Consultation, Monitoring and Review.

It may be the case that no individual within the organisation possesses all these skills, in which case the combined skills of two (or possibly more) people may be employed. In any event, it is always advantageous to have a second opinion when undertaking this kind of activity, and another pair of eyes and ears can pick up things that an individual might overlook.

An effective method for commencing an impact assessment is to hold an introductory workshop to inform and engage the senior management team, who will then be best positioned to brief their departmental managers and staff on what to expect and what to research. Beyond that, the approach taken for interviews will need to be tailored to the type of audience.

Some interviews are best conducted on a one-to-one basis, especially if there are sensitive issues at stake, whereas others may benefit from a group discussion or workshop in order to ensure that the views of different sections within a department are considered.

Telephone interviews should be avoided where possible unless there is no alternative, as the visual signals given in face-to-face meetings can provide additional clues; however,

video conferencing calls such as Zoom, FaceTime, Skype and Webex may well be an acceptable alternative. Follow-up telephone discussions to review the findings present less of a problem.

At all times, the interviewer should give the interviewee time to reflect on the question and, once an answer has been recorded, it is worth the interviewer summarising the discussion as a quick means of verifying the findings. It is also important to choose the order of questions carefully. Questioning should follow a logical sequence – for example, the same sequence as a development or production process in which information is used and generated – in order to permit the interviewee to follow a more logical thought process when answering.

At the beginning of the risk identification process, the most appropriate format of the impact analysis should be agreed with the programme sponsor, although flexibility is always important and it may be necessary to modify the format as the interviews progress and findings are recorded.

EVENTS OR INCIDENTS?

For the purposes of this book, we shall refer to events and incidents as having the same characteristics, since both terms are frequently used interchangeably to discuss something detrimental that has occurred. However, it is worth taking a very brief look at how the incident management (IM) – especially the BC management – communities view the terms. A general view is illustrated in [Table 4.1](#).

Table 4.1 The general properties of detrimental situations

Type	Glitch	Event	Incident	Crisis	Disaster	Catastrophe
Timeframe	Seconds	Minutes	Hours	Days	Weeks	Months
Focus	Equipment		Operations	Management	Board	Government
Recovery	Automatic		Process	Improvise	Ad hoc	Rebuild
Method	Proactive			Reactive		

Glitches

Extremely short occurrences are often referred to as glitches. They usually last a few seconds at the most and generally refer to brief interruptions in power, computer processing, television and radio transmissions. Activities usually return to normal following most glitches as equipment self-corrects automatically.

Events

Events normally last no more than a few minutes. Like glitches, the equipment they affect is frequently automatically self-correcting, but may on occasion require a degree of manual intervention.

Incidents

Incidents are viewed as lasting no more than a few hours. Unlike glitches and events, they usually require operational resolution, normally requiring manual intervention that follows some form of process.

The methods of dealing with glitches, events and incidents are all proactive in nature.

Crises

Crises generally last for at least several days. Although organisations may have plans, processes and procedures to deal with them, and although operational staff will carry out any remedial actions, some degree of improvisation may be required. Crises almost invariably require a higher layer of management to take control of the situation, make decisions and communicate with senior management and the media.

Disasters

Disasters generally last for weeks. As with crises, operational staff will carry out remedial actions, although at this stage a degree of ad hoc action may be necessary, and, even if a higher management layer will control activities, the senior management layer will take overall charge of the situation.


Catastrophes

Catastrophes are the most serious level, often lasting for months, or in some cases for years. Their scale tends to affect many communities, and so, although individual organisations may be operating their own recovery plans, it is likely that local, regional or even national government will oversee the situation and that a complete rebuilding of the infrastructure may be required.

Despite any proactive planning or activities to lessen their impact or likelihood, crises, disasters and catastrophes all require reactive activities.

Incidents, crises, disasters and catastrophes generally follow a set sequence of planning, preparation and testing, incident response, business continuity operations and resumption to normal operations, as shown in [Figure 4.3](#).

Figure 4.3 Generic sequence of situation management



IMPACT ASSESSMENT

The next real piece of work in risk assessment is that of identifying the impacts or consequences to the information assets.


Let us begin by understanding what actually comprises impact or consequence. In general terms, an impact is the adverse result of some action (a threat or hazard), taking advantage of a weakness or vulnerability in an information asset, as shown in [Figure 4.4](#).

Figure 4.4 A simple threat, vulnerability and impact



However, the situation is not always as simple as this, since there can be more than one threat against an information asset that takes advantage of the same vulnerability and results in the same impact, as shown in [Figure 4.5](#).

Figure 4.5 Multiple threats can exploit the same vulnerability



Also, of course, it is quite possible that an information asset exhibits a number of vulnerabilities and that a single threat attacks each, resulting in a number of different impacts, as depicted in [Figure 4.6](#).

Finally, there is the so-called 'knock-on effect' or 'chain of consequence', in which a threat exploits a vulnerability in one information asset, and the resulting impact becomes a threat to another information asset, and so on. These chains of consequence can be quite complicated to follow, but whenever impact assessments are carried out, it should be borne in mind that such a chain may well exist, and where the assessor can imagine links of dependencies between information assets, the possibility of such a chain should be consciously explored.



Figure 4.6 A single threat can exploit multiple vulnerabilities

Figure 4.7 illustrates such a chain of consequence. A threat exploits a vulnerability in asset 1, resulting in an impact that has a knock-on effect to become the threat facing assets 2 and 3, and in turn the impact on asset 3 becomes the threat that affects asset 4.

Figure 4.7 A typical chain of consequence

An excellent example of chains of consequence is the failure in a fully automated environment of an organisation's order processing system. Without the output from this, the organisation's production line would not be able to produce the order, the despatch

department would have nothing to deliver and the billing department would be unable to produce an invoice.

When examining chains of consequence, it is important that progress through the chain is traced back to its root cause – sometimes referred to as 'root cause analysis', since it is the root cause that triggers all the subsequent impacts.

Types of impact

Before we examine the various types of impact, we need to understand that they will affect the organisation in slightly different ways, depending on their origin. On the first level, there are two categories:

Primary or immediate impacts. These result from the event itself, when a business function is detrimentally affected or unable to continue. They come in two varieties:


- **Direct primary impacts** – for example, if a customer database is hacked and personal information is stolen, the organisation will lose control of that valuable information resource.
- **Indirect primary impacts** – these may occur as a consequence of the direct impact. In the above example, the Information Commissioner might levy a fine against the organisation for failing to adequately protect the information.

Secondary or future impacts are those that result from responding to or recovering from the event.

- **Direct secondary impacts** include such things as customers purchasing their products or services from another supplier.
- **Indirect secondary impacts** include such things as fines imposed for failing to file statutory returns on time because the necessary information is unavailable.

Figure 4.8 takes the 'simple' impact model described above, and breaks it down further.

Figure 4.8 Impact types



For the purposes of this book, we have organised impacts into five types, and we shall examine each in greater detail:

- operational impacts;
- financial impacts;
- legal and regulatory impacts;
- reputational impacts;
- wellbeing of staff and the public-at-large impacts.

Operational impacts

As the name suggests, operational impacts are those that impair the day-to-day activities of the organisation. Very often, direct operational impacts will result in subsequent indirect financial impacts, so an inability to meet a service contract may well result in lost orders or claims for contractual damages. These include (but are not limited to):

- the loss of or damage to the confidentiality, integrity or availability of information;
- the loss of or damage to premises and equipment;
- order backlogs;
- productivity losses;
- industrial action;
- reduced competitive capability;
- the organisation's inability to meet service contracts;
- the organisation's inability to progress new business or developments;
- damage to third-party relations;
- impaired management control;
- the loss of customers to competitors, known in some sectors as 'churn'.

Most operational impacts are felt by the organisation very quickly – their presence will usually be very obvious, and not only within the organisation itself, but also to the public-at-large and the media. Because of this, operational impacts are frequently dealt with in a reactive way by invoking some form of IM process to bring the situation under control, to recover from the incident and eventually to return to normal. IM is more usually associated with BC, but very much has a place in information risk management as well.

Financial impacts

Unsurprisingly, financial impacts or consequences are normally those that gain the greatest attention within the organisation. It is frequently against a backdrop of possible financial loss that the costs of remedial actions will be compared. While this is certainly correct, it is also true for other types of impact as well:

- loss of current and future business opportunities;
- increased cost of borrowing;

- cancellation of contracts;
- contractual penalties;
- loss of cash flow;
- replacement and redevelopment costs;
- loss of share price;
- increased insurance premiums;
- loss of tangible assets.

Many of these impacts – for example, lost sales immediately following the event – will be felt very quickly, while others – for example, increased insurance premiums – may not manifest themselves until a later date, possibly some considerable time after the costs of the event have been counted.

Financial impacts may also not be as noticeable to the whole organisation – for example, staff may not be aware of the financial implications of an event at all, and have no appreciation of the position in which the organisation finds itself until they read about it in the media or find that pay increases and bonuses are reduced.

Legal and regulatory impacts

As with reputational impacts, legal and regulatory impacts can have serious repercussions on an organisation, and the handling of these is best dealt with by a specialist team within the organisation – who may communicate information regarding an event through the corporate communication department. Legal and regulatory impacts, which can also be referred to as consequential losses, include:

- warnings or penalties from sector regulators;
- fines for late submission of company accounts;
- fines for late payment of taxes;
- breach of contract damages;
- penalties for fraud and other criminal acts.

Reputational impacts

Reputational impacts are almost always highly detrimental to the organisation. For this reason, many organisations employ communication specialists who are skilled in countering negative publicity and putting a positive spin on any bad news. In such organisations, most staff are advised not to talk directly to the media, but to pass enquiries through to the corporate communication department.

Reputational impacts, which can also be referred to as consequential losses, include:

- stock market confidence;
- competitors taking advantage;
- customer perception;

- public perception;
- industry and institutional image;
- confidential information made public.

The reputation of an organisation can be destroyed overnight. Take the case of the Gerald Ratner chain of cut-price jewellery shops in the 1990s. Ratner made a speech at the Institute of Directors during which he made several derogatory remarks about the products he sold. Despite the fact that he had thought the remarks were 'off the cuff', they were reported in the press and customers exacted their revenge by staying away. The value of the organisation plummeted by around £500 million, and the company very nearly ceased trading.

Wellbeing of staff and the public-at-large impacts

Although more rare, safety incidents are generally highly visible outside the organisation, and occasionally have an impact on the public. More common, however, are any events that may have an adverse effect on the organisation's staff, and these can also cascade into financial and operational secondary impacts. Wellbeing impacts include:

- safety, health risks and injuries;
- stress and trauma;
- low morale of staff.

Qualitative and quantitative assessments

Qualitative assessments are almost always subjective. The terms 'low', 'medium' and 'high' give a general indication of the level of impact, but do not tell us how much. Quantitative assessments, on the other hand, can be very specific, and rather more accurate. The main problem with quantitative assessments is that, if they are to be entirely useful, they can take a great deal of time to undertake, and mostly total accuracy is not a requirement.

In the case of impact assessment, it may be worthwhile taking an initial qualitative pass in order to give an idea of the levels of risk and to identify those risks that are likely to be severe, with the objective of a more qualitative assessment later on.

Alternatively, a compromise solution often works well, since it can combine objective detail with subjective description. Instead of spending significant amounts of time in establishing the exact losses that might be incurred by a particular activity, the compromise solution takes a range of impact values and assigns descriptive terms to them. This is also known as semi-quantitative assessment.

So, for example, we might state that in a particular scenario, the term 'very low' approximates to values up to £25,000, 'low' to values between £25,000 and £250,000, 'medium' to values between £250,000 and £1 million; 'high' to values between £1 million

and £25 million; and 'very high' to values in excess of £25 million. Although we have provided boundaries for the levels, there will be a degree of uncertainty about the upper and lower limits of each, but in general the ranges should be sufficient to provide a reasonable assessment, placing it in terms that the board will understand quickly. Clearly, these ranges will differ from one scenario to another, but set a common frame of reference when there are a substantial number of assessments to be carried out. These ranges should be agreed when setting the criteria for assessment.

For those situations in which there are no applicable financial values, such as reputational or operational impacts, a similar method of quantification can be used. [Table 4.2](#) illustrates some possible options.

Table 4.2 Typical impact scales

Level of impact	Operational	Financial	Legal and regulatory	Reputational	Wellbeing
Very low	Partial loss of a single service	Loss of less than £25K	Warning from regulatory body	Minor negative publicity	Inconvenience to several people
Low	Total loss of a single service	Loss between £25K and £250K	Penalties up to £10K	Local negative publicity	Injury or harm to one person
Medium	Partial loss of multiple services	Loss between £250K and £1M	Penalties between £10K and £50K	National negative publicity	Injury or harm to several people
High	Total loss of multiple services	Loss between £1M and £25M	Penalties between £50K and £500K	EU-wide negative publicity	Loss of single life
Very high	Total loss of all services	Loss exceeds £25M	Penalties exceed £500K	Worldwide negative publicity	Multiple loss of life

Impact assessment questions

Let's take a look at the questions we need to ask when conducting an impact assessment. Note that a suggested template form for this is included in [Appendix F](#).

To begin with, apart from basic information such as name of department, name of contact(s) and contact details, for each information asset we will need to understand:

- The capital value of the asset if it has one.
- The current and projected revenues the asset generates or contributes to.

- The value of any resources and activities required to maintain and support this asset, including:
 - staff;
 - systems hardware;
 - systems operating system and application software;
 - information such as databases, designs, etc.;
 - services, including third-party services;
 - premises;
 - other infrastructure.
- The required operational levels for each of these resources, in the short, medium and long terms.
- How long the organisation can survive without these resources, expressed in minutes, hours, days or weeks (note that in BC terms, this is often referred to as the maximum tolerable period of disruption, or MTPD).
- The impact on the information asset in terms of operational, financial, legal and regulatory, reputational and people impact.
- How long any relevant activities take to complete.
- What effect disruptions have on these activities and other activities within the business.
- What dependencies there are for activities to be completed.

Who should be involved in an impact assessment?

Clearly, the information asset owner should be the first port of call for this activity. However, although the information asset owner may be familiar with the technical details of the asset itself, he or she may not be able to express its value in business terms, so it is always worth double-checking with another person – or group of people in the case of substantial information assets such as a customer database, which may cut across several different areas of the organisation.

In larger organisations, it is quite possible that impact assessments may be carried out at three different levels, with the senior management team defining what information assets are critical to the organisation at a strategic level, departmental managers refining this assessment and identifying the individual components, and operational managers making the final impact assessments themselves.

Strategic impact assessment

The strategic impact assessment begins with input from senior management who have a detailed high-level view of the organisation. They will be able to interpret the organisation's position in terms of both the internal and external contexts, and will appreciate the value of the customer base and the organisation's overall information assets.

The strategic impact assessment will determine which information assets are deemed to be essential to the survival of the organisation, and will make an initial estimate of the importance of each. This may well be modified in the tactical impact assessment when details of the information assets become clearer. If the organisation decides to omit any information assets from the impact assessment, the reasoning behind this decision should be clearly stated.

The strategic impact assessment should take into account the business-specific factors or general impact types that might affect the organisation, such as the interests of stakeholders, any statutory, legal or regulatory obligations, the reputation of the organisation and its financial future. The organisation must begin by clearly defining at what point the level of impact constitutes a change from 'tolerable' to 'intolerable'.

When conducting a strategic impact assessment, each information asset should be separately investigated against each general impact type, and estimation made of the likely time that might elapse before it becomes business-affecting, since loss of information assets may not be felt immediately. It is also useful at this stage to provide any reasoning for this, so a more objective analysis can be conducted, especially in cases where two or more information assets are considered, and the organisation feels it is desirable to prioritise one over the other.

Tactical impact assessment

Senior managers will undoubtedly have a general idea of what information assets are critical to the organisation, but may not have a detailed knowledge of their makeup, or who is responsible for them, so, having agreed which information assets are being covered by the strategic impact assessment, we now turn to the tactical impact assessment and set the scope of this work.

In this case, we are less interested in the external context as this will already have been addressed by the strategic business impact assessment, but we are now interested in those of the organisation's activities that contribute to the delivery of the organisation's products and services and any interdependencies between them. It may be extremely useful to piece together a block diagram or process flow diagram illustrating these interdependencies in order to simplify later work.

Operational impact assessment

Finally, we turn to the operational level of impact assessment, in which those managers having a detailed knowledge of and responsibility for the information asset provide an assessment of the actual potential impact of the asset's damage, loss or destruction.

Over- and underestimation

When carrying out impact and likelihood assessments, a mistake frequently made is either to over- or underestimate the realities of the situation. Often this is done with the best of intentions, but either of these can cause considerable difficulties later in the risk management process.

If over-estimation takes place, more time, effort and money might be spent in treating risks than is necessary, or indeed the balance between the potential losses and the costs of treatment may change to make the correct form of risk treatment uneconomical.


Where underestimation occurs, risks that should be viewed as being more serious might well not be treated to a sufficient extent, or may be unwittingly accepted when avoidance, transfer or reduction might have been a more appropriate option.

Time dependency

Impacts can also change with and over time. Those threats that affect capital items – buildings, equipment and so on – will almost certainly result in an instantaneous financial impact, while those that affect sales revenue streams, for example, will possibly show little impact at first, but over the course of days or weeks the impact would increase in a more or less linear fashion. The other factor affecting this form of impact will be the hours during which the information must be available, since a systems failure that occurs ‘out of hours’ may have no effect on revenue at all, but will invariably cause considerable disruption when business resumes on the next working day.

Figure 4.9 illustrates how the organisation might view the impact of financial losses over time.

Figure 4.9 Potential losses over time following a disruptive event



There are, of course, some times that are the worst possible for an incident to affect information assets. These include:

- regular events, such as end-of-month accounting dates;
- seasonal occasions, such as Easter, Christmas and summer holidays;
- heavy load occasions, such as the beginning of the school or university year.

Validating the results

Once the interviews or workshops have been completed, and the results of the questionnaires have been recorded, they should be passed back to the originators for review and comment. The originators should be allowed time to consider their input and invited to verify the accuracy of the information before it is finalised for presentation. A typical impact assessment form might look something like that shown in [Figure 4.10](#).

Figure 4.10 Typical impact assessment form

Impact/Asset			Date	
Asset owner			Analyst	
Asset location			Reference	
	Primary impacts		Secondary impacts	
	Direct impacts	Indirect impacts	Direct impacts	Indirect impacts
Operational				
Score				
Legal and regulatory				
Score				
Reputational				
Score				
People-related				
Score				
Financial				
Score				
Total financial impact				
Total impact rating				

VH = Very high

H = High

M = Medium

L = Low

VL = Very low

SUMMARY

In this chapter, we have examined the approach to identifying the risks to our information assets, seen how an impact assessment can be conducted and looked at a number of different types of impact, together with understanding the differences between qualitative and quantitative assessments.

The next stage of the information risk management programme is to understand the threats the organisation faces to its information assets, and the vulnerabilities that these may exhibit.

5 THREAT AND VULNERABILITY ASSESSMENT

In 2002, US Secretary of State Donald Rumsfeld said the following during a briefing:

... there are known knowns; there are things that we know that we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns, the ones we don't know we don't know.

This is partly true of threats, but very true of vulnerabilities.

Threats, as we see in the glossary in [Appendix I](#), are the 'potential cause of an unwanted incident, which can result in harm to a system or organisation'. Threats are slightly different from hazards, which may still cause an unwanted incident resulting in harm to an information asset, but whereas threats are generally man-made and deliberate, hazards are more usually accidental or naturally occurring.

Vulnerabilities are defined as 'the intrinsic properties of something resulting in susceptibility to a risk source that can lead to an event with a consequence'. Vulnerabilities are weaknesses either in information assets themselves, or in the infrastructure that supports the information assets, including the people who have ownership or responsibility for them.

CONDUCTING THREAT ASSESSMENTS

Some experts believe that the threat and vulnerability assessments should be carried out ahead of the impact assessments; others disagree and opt for the reverse arrangement.

I believe that, in practice, either method will suffice as long as the information assets have been clearly identified, but that it can be extremely helpful if the threat and vulnerability assessments can be performed at the same time as impact assessments, since many of the threats and vulnerabilities will be apparent to the information asset owners. Further threat and vulnerability assessments can be conducted at a later time with other knowledgeable staff, especially with information security specialists.

For every threat identified, there may well also be some data on the frequency of historical events where the threat has either been known to have been used or to have succeeded.

It is also worthwhile remembering that a threat can only cause an impact on an information asset if the asset contains a vulnerability for the threat to exploit.

To begin with, it may well be worth running a brainstorming session to identify possible threats. In the first pass through, as with normal brainstorming rules, no suggestion should be discounted, no matter how bizarre it might seem, since sometimes those ideas that seem crazy at first glance turn out to be viable threats. The weeding out of the very unlikely threats can come at a later stage when the likelihood assessments are carried out.

It is worthwhile remembering, however, that even the most thorough threat assessment might not identify all the threats and hazards that the organisation faces, and also that new ones may emerge with time, so, as with all other aspects of an information risk management programme, this should be an ongoing activity.

Following the first pass through a brainstorming session, it may also be beneficial when conducting both threat and vulnerability assessments to use a mind map, so that all threats, hazards and vulnerabilities can be grouped, as shown in [Figure 5.1](#).

The output of the threat assessment will include threats and hazards from a number of different sources including, but not limited to:

- malicious intrusion or hacking;
- environmental threats and hazards;
- errors and failures;
- social engineering;
- misuse and abuse;
- physical threats;
- malware.


These threats and hazards are described briefly below, and in greater detail in [Appendix B](#).

Malicious intrusion or hacking

Hacking is a generic term applied to many forms of unpleasant behaviour, although it began as a description of what people did in order to find out how computers worked and how to improve their performance. Hacking almost invariably results in a breach of confidentiality, integrity or availability as hackers use software tools to intercept and decrypt legitimate information, and either steal it or change it. Occasionally, hacking is used to deliver so-called 'denial of service' attacks, designed to prevent legitimate access to systems, often to make a political point.

Since the introduction of the Computer Misuse Act in 1990, hacking is now treated as a crime, since it invariably involves accessing a computer without the owner's permission to do so.

Figure 5.1 Typical threats and hazards



It is becoming more and more common to hear news stories about hackers who steal large quantities of information, such as user identifiers and passwords, and sell this information on – usually to criminal gangs – for use in wider fraud.

On 22 September 2016, Yahoo! Inc. announced that a massive data breach of its service had taken place – with roughly 500 million user accounts' passwords, along with other sensitive information, stolen – claiming that the breach was perpetrated by a 'state actor'. Interestingly, details of the attack were revealed just as Yahoo! was trying to negotiate a buy-out by Verizon – the news would certainly have affected the company's share value.

Hacking includes:

- denial of service (DoS) attacks;
- unauthorised access;
- unauthorised network scanning;
- interception;
- session hijacking;
- website modification;
- software modification;
- data modification;
- decryption;
- credential theft.

Environmental threats and hazards

These types of threat are almost always concerned with availability, since they affect the environment in which a system resides. Those threats that occur as a result of natural events – for example, severe weather – are often referred to as hazards in order to distinguish their origin from those of malicious threats. Many of these hazards affect a wide geographic area, and can cause serious disruption to multiple organisations rather than to a specific organisation or system.

Examples of environmental threats include:

- natural hazards such as severe storms and flooding;
- accidental and malicious physical damage;
- fire;
- communications jamming or interference;
- communications failures;

- power failures;
- pandemics.

Errors and failures

Errors fall neatly into two categories: those made by users and technical staff, and those things that simply fail. Neither form is generally regarded as being malevolent, even though some user and technical errors are caused by lack of attention or poor training. Despite the view of many technicians that both hardware and software are designed to cause them grief, there is no evidence to suggest that this is actually the case.

Examples of errors and failures include:

- software failures;
- software interdependencies;
- system overloads;
- hardware failures;
- user errors;
- technical staff errors;
- internal and external software errors;
- change failures.

Social engineering

Social engineering is a technique used by hackers and other ne'er-do-wells to acquire information, generally about access to systems so that their hacking activities are simplified. Social engineering comes in several forms – not only the traditional approach where a hacker attempts to engage with a user by conversation (usually over the telephone or by email), but also by disguising malware as legitimate software and web links and by copying the style-naming conventions and language of a target organisation. For example, they may send a user an email that appears to originate from their bank, but in which embedded web links take the user to the hacker's own website.

Examples of social engineering threats include:

- spoofing, masquerading or impersonation;
- phishing;
- spam;
- disclosure.

Misuse and abuse

Whereas hacking is usually deemed to originate from outside an organisation, misuse normally originates from within. The net result may well be the same for either approach,

but in the case of misuse, the internal user or technician has the added advantage of already being on the inside of the organisation's firewall and security systems, may have access to the required passwords and, critically, may also have elevated access privileges. For this reason, the threat from internal attackers potentially presents a significantly greater level of likelihood of success than that of an external attacker.

Examples of misuse threats include:

- modification (invariably elevation) of privileges;
- unauthorised system activity;
- software and business information theft.

Physical threats

Physical threats may also be easily undertaken by employees – many will have access to systems and equipment that they can readily remove from the organisation's premises without the fear of discovery, whereas an external attacker would have to pass through the organisation's layers of physical security in order to do so.

There is a salutatory (possibly apocryphal) anecdote from the building trade that tells of the employee who pushed a wheelbarrow covered with a tarpaulin home at the end of every day's work. Every evening the site foreman checked beneath the tarpaulin to find there was nothing there and let the employee go on his way. Eventually it was discovered that the man was stealing wheelbarrows and tarpaulins.

Physical threats include:

- unauthorised access;
- theft of computers and portable devices;
- theft of authentication devices.

Malware

The term 'malware' is used to refer to malicious software that can be used to attack an information system. Examples of malware include software entities that result in the collection of, damage to or removal of information. Such software is almost always concealed from the user, often self-replicating (attaching itself to an executable program) and can spread to other systems when the user unwittingly activates it.

Some malware goes to great lengths to conceal its existence, appearing to the user as legitimate software. Its purpose, however, is usually sinister in that it may collect, damage or remove information when the user activates what they believe is a legitimate program.

Examples of malware include:

- viruses;
- worms;
- Trojan horses;
- rootkits;
- spyware;
- active content;
- botnet clients;
- ransomware.

Who should be involved in a threat assessment?

As with the impact assessments covered in the previous chapter, the information asset owners should be the first port of call for this activity, since they may well already be aware of many of the threats their information assets face. However, other parts of the organisation will be able to provide input on this, such as the IT department, human resources and the organisation's information security team. A typical threat assessment form might look something like that shown in [Figure 5.2](#).

Additionally, there are comprehensive examples of threat types to be found in [Appendix C](#) of ISO/IEC 27005:2018, including suggestions as to the origin of threats – accidental, deliberate and environmental – as well as the possible motivations and consequences of threats resulting from various types of threat source.

CONDUCTING VULNERABILITY ASSESSMENTS

Vulnerabilities are weaknesses either in information assets or in the infrastructure that underpins them, while threats exploit vulnerabilities in order to achieve an impact.

The output of the vulnerability assessment will include such vulnerabilities as:

- access control failures;
- systems acquisition, development and maintenance procedures;
- physical and environmental failures;
- communications and operations management failures;
- people- and process-related security failures.

However, as already stated, vulnerabilities alone cannot cause an impact on an information asset, as an impact requires the presence of a threat to exploit the vulnerability.

Figure 5.2 Typical threat assessment form

Date		Reference	
Threat description			
Hacking			
Environmental threats and hazards			
Errors and failures			
Social engineering			
Misuse and abuse			
Physical threats			
Malware			
Operating systems affected			
Applications affected			
Information types affected			
Previous attack history (if known)			
Previous success rate (if known)			
Attack motivation (if known)			

For every vulnerability identified, there may well also be some additional data on whether the vulnerability has been known to have been exploited, whether such exploits have succeeded and whether there might be known controls that are already in place in partial mitigation.

As with threat assessments, it may well be worth running a brainstorming session to identify possible vulnerabilities, and in the first pass through, as with normal brainstorming rules, no suggestion should be discounted no matter how bizarre it might seem, since sometimes those ideas that seem crazy at first glance turn out to be viable vulnerabilities. The weeding out of the very unlikely vulnerabilities can come at a later stage when the likelihood assessments are carried out.

It is worthwhile remembering, however, that even the most thorough vulnerability assessment might not identify all the vulnerabilities that affect the organisation's information assets, and also that new ones may emerge with time, so as with all other aspects of an information risk management programme, this should be an ongoing activity.

Following the first pass through a brainstorming session, it may also be beneficial when conducting vulnerability assessments to use a mind map, so that all vulnerabilities can be grouped, as shown in [Figure 5.3](#).

Typical vulnerabilities are described briefly below, and in greater detail in [Appendix C](#):

- access control failures;
- poor systems acquisition, development and maintenance procedures;
- physical and environmental failures;
- communications and operations management failures;
- people-related security failures.

Access control failures


Access control has two complementary uses – firstly to permit access to resources for authorised persons, and secondly to deny access to those resources to unauthorised persons. Failures in access control are very likely to increase the likelihood of successful attacks against information assets.

If you are able to obtain a copy, Appendix C4 of BS 7799-3:2017 provides a highly comprehensive list of vulnerabilities, together with a brief description of how they might be exploited.

Access control failures include:

- The lack of, or poorly written access control policies.
- Failure to change access rights of users changing roles within the organisation.

Figure 5.3 Typical vulnerabilities



- Failure to revoke access rights of users leaving the organisation.
- Inadequate user password management.
- The continued use of default system accounts and passwords.
- The use of passwords embedded in software applications.
- The lack of security of mobile devices.
- The lack of network segregation.
- Failure to impose a clear desk and clear screen policy.
- The use of untested software.
- Failure to restrict the use of system utilities.

Poor systems acquisition, development and maintenance procedures

When acquiring systems hardware and software, developing software and maintaining both, it is vital to ensure that selection or specification is carried out according to a formal set of criteria that include appropriate security features. Unlike access control failures, this type of vulnerability is rarely noticed immediately but can result in serious consequences at a later time.

The root cause of this is often a failure to correctly specify appropriate criteria prior to acquisition or development, and may result either from a lack of forethought or a desire to achieve cost savings.

This type of vulnerability includes:

- The lack of clear functional purchasing specifications.
- The lack of clear functional development specifications.
- Failure to validate data entry.
- The use of untested software.
- The use of unauthorised software.
- The lack of BC and DR planning.

Physical and environmental failures

Physical security is normally highly visible, both to staff and to potential intruders. Very often, the mere presence of robust security is sufficient to deter an intruder, but even so it is important that physical security measures are appropriate and well maintained, and failure to do this increases the organisation's likelihood of experiencing intrusion of some description.

Environmental vulnerabilities tend to be rather more difficult to address, but are generally relatively easy to identify and can either relate to the location or construction of premises (for example, a data centre built on a flood plain) or to the environmental subsystems that underpin major premises such as large office buildings, factories, warehouses and data centres.

Physical and environmental vulnerabilities include:

- Poor management of access to premises and to areas within them.
- Inadequate physical protection for premises, doors and windows.
- The lack of suitable protection for stored equipment and supplies, and especially waste.
- The use of unsuitable environmental systems, including cooling and humidity control.
- The location of premises in areas prone to flooding.
- The uncontrolled storage of flammable or hazardous materials.
- The location of premises in proximity to flammable or hazardous materials or facilities that process them.
- The lack of standby power supplies.

Communications and operations management failures

Along with access control failures, failures of operations management and communications systems rank high among the vulnerabilities that can be successfully exploited, whether deliberately or accidentally.

Many of these exploits are possible due to process failures – again, either through failure to observe them or failure to have processes in the first place.

Communications and operations management failures include:

- The failure to ensure the appropriate segregation of duties where necessary.
- Inadequate network monitoring and management including intrusion detection.
- The use of unprotected public networks.
- The uncontrolled use of users' own wireless access points (WAPs).
- Poor protection against malware and failure to keep virus protection software up to date.
- Failure to maintain patching of software.
- Inadequate and untested backup and restore procedures.
- Improper disposal of 'end of life' storage media.
- The lack of robust 'bring your own device' (BYOD) policies.
- Inadequate change management procedures.
- The lack of audit trails, and non-repudiation of transactions and email messages.
- The lack of segregation of test and production systems.
- The uncontrolled copying of business information.

People-related security failures

The vulnerabilities that are caused by the failures of users and operational staff are almost all related to policies and processes:

- The insufficient or inappropriate security training of technical staff.
- The lack of appropriate security awareness training for users.
- The lack of monitoring mechanisms, including intrusion detection systems.
- The lack of robust policies for the correct and appropriate use of systems, communications, media, social networking and messaging.
- The failure to review users' access rights whenever they change roles or leave the organisation.
- The lack of a procedure to ensure the return of assets when leaving the organisation.
- Demotivated or disgruntled staff.
- Unsupervised work by third-party organisations or by staff who work outside normal business hours.

Who should be involved in a vulnerability assessment?

As with the threat assessments covered earlier in this chapter, the information asset owners should be the first port of call for this activity, since they may well already be aware of many of the vulnerabilities their information assets possess. However, other parts of the organisation will be able to provide input on this, such as the IT department, human resources and the organisation's information security team.

Since many of the vulnerabilities identified will be of a technical nature, it is probable that various forms of security testing will identify them, including penetration testing and vulnerability scanning tools.

A typical vulnerability assessment form might look something like that shown in [Figure 5.4](#).

IDENTIFICATION OF EXISTING CONTROLS

Controls are generally designed to treat risks – either by removing or reducing the impact on an asset, or by removing or reducing the likelihood of the threat occurring. However, controls may also be used to monitor processes, ensuring predictability without actually modifying them.

Before we can move on to the next stage of risk assessment, we must identify all existing controls that are in place and also verify that they are working as expected. There are three reasons for doing this:

- Firstly, to ensure that when we conduct the likelihood assessment, we have all the necessary information to do so accurately.

Figure 5.4 Typical vulnerability assessment form

Date		Reference	
Vulnerability description			
Access control			
Systems acquisition, development and maintenance			
Physical and environmental			
People-related			
Communications and operations			
Operating systems affected			
Applications affected			
Information types affected			
Previous attack history (if known)			
Previous success rate (if known)			
Attack motivation (if known)			

- Secondly, to ensure that when we begin the process of determining the appropriate controls to mitigate the risks we have encountered, we do not attempt to duplicate controls that are already in place, and that they are effective and functioning as expected.
- Thirdly, so that any controls that are not effective or functioning as expected can be reviewed and removed or replaced as necessary.

If the organisation has already been through one complete information risk management programme, then the information to complete this activity will be in the risk register. If, however, this is the first time it has been conducted, the identification of existing controls will have to be carried out from scratch.

We deal with controls in much greater detail in [Chapter 7](#), but in order to assist the identification process for this stage of the work, the following should prove useful.

Controls are divided into four strategic, four tactical and three operational types, but not all combinations of these exist, as illustrated in [Figure 5.5](#).

Strategic controls are:

- Avoid or terminate – that is, either cease an activity that incurs risk or do not begin it.
- Transfer or share – that is, spread the risk between the organisation and one or more third parties.
- Reduce or modify – that is, change either the impact or the likelihood in some way.
- Accept or tolerate – when the other options cannot treat the risk, or when there is residual risk remaining after any of the above methods have been used, this is the one choice remaining.


Tactical controls are:

- Detective – that is, being alerted to something happening.
- Preventative – that is, stopping something from happening.
- Directive – that is, putting in place some form of instruction.
- Corrective – that is, altering the state or condition of something.

Operational controls are:

- Physical or environmental – these are usually concerned with the infrastructure that underpins the information assets. Typical examples of physical controls are the use of CCTV to monitor areas within a site, and electronic door locking mechanisms to control access into restricted areas.
- Procedural or people – these are concerned with ensuring that processes and procedures are followed. Typical examples of procedural controls are change control mechanisms to manage additional systems or services, and the segregation of duties that could otherwise result in fraud.
- Technical or logical – these are usually concerned with hardware and software in some form. Typical examples of technical controls include the use of antivirus software to quarantine or delete malware, and firewalls to block unauthorised network intrusion.

Figure 5.5 The overall scheme of risk treatment options



Who should be involved in the identification of existing controls?

As with the threat and vulnerability assessments covered earlier in this chapter, the information asset owners should be the first port of call for this activity, since they may well be aware of many of the controls already present for their information assets. However, other parts of the organisation will be able to provide input on this, such as the IT department, human resources and the organisation's information security team.

A typical controls identification form might look something like that shown in [Figure 5.6](#).

SUMMARY

In this chapter, we have examined different types of threat and vulnerability, the process by which we assess them, and briefly looked at the types of mitigating controls that may already be in place to protect the assets.

We now need to move on to the next stage of the process, which is to assess the likelihood of a threat taking advantage of a vulnerability, how we analyse the resultant risk, and how we evaluate it before treatment can be undertaken.

Figure 5.6 Typical existing controls identification form

Date		Reference	
Description of controls			
Asset name			
Asset location			
Preventative controls	Physical or environmental		
	Technical or logical		
Detective controls	Physical or environmental		
	Technical or logical		
	Procedural or people		
Directive controls	Procedural or people		
Corrective controls	Physical or environmental		
	Technical or logical		
	Procedural or people		

6 RISK ANALYSIS AND RISK EVALUATION

The process of risk assessment continues with risk analysis, in which we develop an understanding of the risks. We begin by identifying the likelihood or probability of a threat or hazard having an impact on an information asset, and using that likelihood assessment together with the impact we established earlier, we calculate the overall level of risk.

In risk identification, we examined the general impacts or consequences faced by an information asset, then the threats that might cause them, followed by any vulnerabilities they might possess. These three assessments were carried out in isolation, since at that stage of the risk management process the relationship between them did not matter.

In risk analysis, however, we bring the three assessments together, along with any controls already implemented, and examine the impacts that might occur to information assets as a result of specific threat events. This may also require an understanding of any motivations that might exist for deliberate incidents.

Firstly, we assess how likely it is for any given threat or hazard to exploit a vulnerability and cause harm to an information asset.

ASSESSMENT OF LIKELIHOOD

As we mentioned earlier, the likelihood of a threat taking advantage of a vulnerability to cause harm to an information asset will depend on a number of factors, including:

- The value of the asset to the attacker, which is usually, but not always, financial.
- The history of previously successful attacks.
- The risk of an attacker being detected either during or following an attack, whether successful or not.
- The complexity of the attack.
- The motivation of the attacker, which is sometimes financial or sometimes the result of a grievance.
- The skills and tools required to carry out the attack.
- The level of any vulnerabilities within the information asset, including how well they are known.
- The presence or otherwise of existing controls.

Clearly, for environmental hazards, errors and failures, the concept of an attacker has no meaning, but for the other forms of threat described – physical threats, malicious intrusion or hacking, misuse, abuse and malware – an attacker will be present in some form, and this will have a profound effect on the likelihood of the attack succeeding.

An attacker may simply wish to steal the information asset or make a copy of it; this is straightforward theft, and hence has a confidentiality impact. Alternatively, the attacker may wish to alter the information asset to gain some benefit – a higher examination grade, for example – or to reduce the benefit to the information owner, perhaps because of some perceived grievance – to cause customer dissatisfaction, for example. These have an integrity impact on the information asset. Finally, the attacker may wish to deny the information owner and/or the owner's customers rightful access to the information asset, again possibly because of some perceived grievance, and this is an availability impact.

Clearly, the greater the motivation an attacker has, the greater will be the likelihood of the attack being successful, and possibly also the greater the impact on the information asset.

However, likelihood or probability can still be extremely difficult things to rate. Many of the events we will consider occur quite randomly, rather than at predefined or likely intervals, and so we are unable to determine exactly when a particular incident is likely to happen. Estimation is therefore the order of the day, but we need to take care with our approach so that it is both meaningful and consistent.

As with the assessment of impact, we turn to the two possible methods of likelihood assessment – qualitative and quantitative.

Qualitative and quantitative assessments – which should we use?

Qualitative assessments are almost always subjective. The terms 'low', 'medium' and 'high' give a general indication of either the degree of impact or the likelihood of an event occurring, but little more. Quantitative assessments, on the other hand, can be very specific, and are generally rather more accurate. The main problem with quantitative assessments is that the greater the level of accuracy that is required, the greater will be the amount of time the assessment will take to complete. Only in a very few cases is total accuracy a requirement.

It may be worthwhile taking an initial qualitative pass in order to gain an idea of the levels of risk and to identify those risks that are likely to be severe, with the objective of a more quantitative assessment later on.

Alternatively, a compromise solution often works well, since it can combine objective detail with subjective description. Instead of spending significant amounts of time in establishing the exact likelihood of a particular event, the compromise solution takes a range of likelihood values and assigns descriptive terms to them. This is also known as semi-quantitative assessment. So, for example, we might decide that in a particular scenario, the term 'very unlikely' approximates to an event that occurs once in a decade; 'unlikely' to one that occurs once a year; 'possible' to one that occurs once a month;

'likely' to one that occurs once a week; and 'very likely' to one that might occur at any time.

The criteria upon which these decisions are made may have some inherent business origin, and may well be based on input from an panel of 'experts', who will reach some form of consensus on the most appropriate values. This approach is sometimes referred to as the Delphi method, and, even if the decisions are incorrect in any way, the organisation can be confident that they have been reached by a formal method.

Although this provides boundaries for the levels, there will be a degree of uncertainty about the upper and lower limits of each, but, with forethought, the ranges should be sufficient to provide a fairly reasonable assessment while placing it in terms that the board will understand quickly. Clearly, these ranges will differ from one scenario to another, but will set a common frame of reference when there are a substantial number of assessments to be carried out. These ranges should be agreed when setting the criteria for assessment.

Fully quantitative assessment relies on our being able to predict the probability of an event occurring based on the likely frequency of it doing so. This requires a knowledge or experience of statistics, which is outside the scope of this book, but unless probability information is available and deemed to be trustworthy, frequency is unlikely to yield reliable results.

Historical data may be useful in providing an initial assessment and in developing appropriate likelihood scales, but it is important to base these on sufficient data so that the results are meaningful, since too little data will almost certainly provide a skewed view of likelihood.

Another problem that increases the uncertainty of likelihood assessments is that of proximity. Just because an event occurs statistically once every 10 years does not mean that it will not occur in two successive years or, conversely, that it will occur at all in the next 20 years.

Likelihood may also be influenced by other factors; for example, when a new vulnerability is discovered, or when some newsworthy event provides attackers with a new-found opportunity, such as the development of a new form of treatment for a disease, allowing them to send out spam emails inviting potential customers to purchase medicines at a low cost.

If quantitative assessment proves to be too challenging, we may opt instead for simple qualitative assessment. However, as with impact assessment, qualitative likelihood assessment tends to be very subjective, and the terms 'highly unlikely' and 'almost certain' mean little except to the person who sets them.

In order to provide a more objective likelihood assessment, we might combine a qualitative scale with quantitative values, as we suggested with impact assessment, so that we can place each assessment on a meaningful basis. For example, severe cold weather-related events tend to be more common in the winter months, while others, such as extreme flooding, may only occur only once in every few years and at any time of year. The Environment Agency in the UK provides estimates of the possible depth of

floodwater, but, as always, these are only estimates and neither the extent of flooding nor the frequency at which it might occur can be relied upon for accurate likelihood determination. The timeframe scale for this kind of event might therefore range from months to decades.

Alternatively, hacking attempts can and do occur much more frequently, and so a timeframe scale for these would be rather different.

The likelihood of whichever kind of threat or hazard we face must therefore be judged against an appropriate scale devised as part of the setting of general risk management criteria, and if the worst-case scenario is used for likelihood assessments, this should ideally be the standard for all.

An example of possible likelihood scales for the different categories of threats described earlier is shown in [Table 6.1](#).

Table 6.1 Typical likelihood scales

Level of likelihood	Hacking, malware and social engineering	Environmental threats	Errors, failures, misuse and physical threats
Very unlikely	The event is likely to occur once a week	The event is likely to occur once a decade	The event is likely to occur once a month
Unlikely	The event is likely to occur once a day	The event is likely to occur once a year	The event is likely to occur once a week
Possible	The event is likely to occur several times a day	The event is likely to occur once a month	The event is likely to occur once a day
Likely	The event is likely to occur several times an hour	The event is likely to occur weekly	The event is likely to occur several times a day
Very likely	The event is likely to occur at any time	The event is likely to occur at any time	The event is likely to occur at any time


The final stage in the assessment of likelihood is to place it against each threat identified earlier in the risk identification process, so that we can begin the next stage – risk analysis.

RISK ANALYSIS

Once we have assessed the impact or consequence of an incident and the likelihood of it occurring, we are in a position to analyse the risk. Typically, this is carried out using a risk matrix of the type shown in [Figure 6.1](#).

A risk matrix is simply a tool to assist in the ranking of risks in terms of overall severity, and in order to help prioritise the assessed risks for treatment. Since it is purely a visual representation of the risks identified, it should only be used at the very end of the risk assessment process, once both the impact and the likelihood have been assessed.

Figure 6.1 A typical risk matrix



Since we will have been consistent in using scales that truly represent the impact and the likelihood regardless of the type of threat, we can safely plot every impact and likelihood assessment on the matrix. In this example, a simple scale of 1 to 25 is defined.

If there are a number of risks assessed, it may be helpful to number them in some way so as to avoid too much clutter on the risk matrix itself.

This allows us to prioritise the risks into five different rankings:

1. Those risks in squares 1, 2 and 6 are very low risk, and we can probably accept them.
2. Those risks in squares 3, 4, 5, 7, 8, 11, 16 and 21 are not urgent, but may require treatment.
3. Those risks in squares 9, 10, 12, 13, 17 and 22 are worth treatment.
4. Those risks in squares 14, 15, 18, 19 and 23 are important and require treatment.
5. Those risks in squares 20, 24 and 25 require urgent treatment.

However, this presents a rather simplistic view of risks when it comes to later prioritisation, since, if we are dealing with a very large number of risks, a significant number could be ranked as urgent and we would have no easy way of prioritising them for treatment. For that reason we might instead use an enhanced risk matrix, in which a slightly altered scale of 1 to 25 is defined, and some bias is given to those risks that are more likely to occur, as shown in [Figure 6.2](#).

Figure 6.2 An enhanced risk matrix

Likelihood or probability	Very likely	6	11	16	21	25
	Likely	5	10	15	20	24
	Possible	3	8	13	18	23
	Unlikely	2	7	12	17	22
	Very unlikely	1	4	9	14	19
		Trivial	Minor	Moderate	Major	Critical
		Impact or consequence				

We will now be in a position to record the risk ranking value and begin prioritising the risks for treatment. The next stage in this process is to transfer everything we know about the risks to a risk register.

RISK EVALUATION

Risk evaluation is the stage in which the various risks analysed are transferred onto a risk register and agreement is reached on which risks require treatment and examines how this should be achieved.

The risk register

The risk register is a method of tracking all risks identified within the information risk management programme, regardless of whether or not they are able to be treated by avoidance or termination, transfer or sharing, modification or reduction, or to be accepted or tolerated.

There is no set structure for a risk register, but it should be remembered that 'less is more', and that too much information can render the risk register unusable. It is better, therefore, to record only the very necessary information in the risk register and to provide links – possibly hyperlinks – to any background information. The following information is necessary and sufficient for completeness:

- A risk reference number, possibly prefixed by the organisation's department name.
- The date the entry was added to the risk register.

- The name of the information asset.
- The name of the asset owner¹.
- A brief description of the risk the asset faces.
- The level of impact assessed.
- The likelihood assessed.
- The resulting risk level.
- The recommended operational control(s).
- The name of the person responsible for implementing the control(s).
- The date by which the controls should be in place.
- The next review date for the risk.

It may also be desirable to indicate whether the control(s) have been successful.

Some organisations develop a Structured Query Language (SQL) database to act as a risk register, since this can be made prescriptive in terms of who may make changes, and how and when the entries are made or changed; smaller organisations often find that a simple spreadsheet is sufficient for their needs, as shown in [Figure 6.3](#).

There is, however, a major caveat when it comes to sharing the risk register. The register itself becomes a valuable information asset, and may contain information that could be detrimental to the organisation if it were to become publicly available, since it provides a comprehensive list of all the organisation's information assets and the risks they face. Therefore, the risk register should be protected with its own set of controls and only shared with those people within the organisation who have a genuine need to see it.

It will be clear at this point that some fields within the register cannot yet be completed, since we have not actually conducted the risk evaluation. This is the next and final stage in the risk assessment process.

Risk evaluation

In the process of risk evaluation, we take each risk in turn and compare it with the risk criteria, which we covered in [Chapter 3](#) – The Information Risk Management Programme.

We begin by comparing the potential impact or consequence with the organisation's risk appetite for that particular information asset. If the impact is less than the level set as the risk appetite, the organisation may decide to accept the risk, but to record the fact and ensure that the risk is monitored over time in case either its impact or likelihood changes.

¹ While one person or department may be the asset owner from the point of view of impact, it is quite conceivable that another person or department might be the owner from the point of view of likelihood. Therefore, the risk register will need to record this.

Figure 6.3 A typical risk register spreadsheet

[illegible]

* Treatment: X = Avoid; S = Share; R = Reduce; A = Accept

Next, we compare the risks with the risk treatment criteria – also established in [Chapter 3](#) – to decide on the most appropriate form of risk treatment, which includes risk avoidance or termination, risk transfer or sharing, risk reduction or modification and risk acceptance or tolerance.

The recommendation at this stage of the process will only be which of the strategic options the organisation should employ, and will not go down to the tactical or operational level, since this is part of the risk treatment process.

Some organisations assist the recommendation process by introducing three separate bands of risk:

- An upper band in which the level of risk is regarded as being intolerable, regardless of the costs or complexities of treatment.
- A middle band in which the cost/benefit of treatment is considered along with other risk appetite criteria.
- A lower band, in which the level risk is so low that it can clearly be accepted or tolerated without the need for debate.

We should also remember that, although any of these choices might reduce the risk to a level below that of the organisation's risk appetite, it may well be possible to reduce it further by employing more than one of the risk treatment options and, in addition to recommending which methods to employ, the organisation may also need to consider in which order they should be employed.

Again, at this point in the process it may become clear that either the risk appetite criteria or the risk treatment criteria are insufficiently well defined to allow the organisation to make an informed decision as to the most appropriate form of risk treatment, since at the time too little may have been known about the exact nature of the risks. These criteria may therefore require refinement, and the evaluation process must be repeated.

Likewise, there may be risks for which the recommended option still remains unclear, and in such cases the recommendation might be to conduct a further and more detailed risk assessment in order to better understand the risk and to make a more informed choice as to the method of treatment.

The recommendations should also provide some indication of the likely cost of treatment, since if the cost outweighs the benefit, the organisation may decide to accept the risk as being too expensive to treat. As with any other accepted risk, this should be documented and monitored over time.

While recommendations on risk treatment tend to be based upon the acceptable level of risk for each information asset, it is important that the evaluation process takes into account three key attributes:

- The aggregation of a number of lower-level risks may result in significantly higher overall risk.

- The importance to the organisation of any information asset when compared to others.
- The need to consider whether legal, regulatory and contractual risks might be more damaging to the organisation than others.

The output of the risk evaluation process will be:

- Whether or not a risk requires treatment and, if so, which strategic option(s) are recommended.
- A prioritised list of the risks identified for treatment.
- An updated risk register reflecting the recommendations made.

It is normal practice at this point to submit a proposal for the overall risk treatment plan to senior management for consideration, together with any business cases required for treating those risks for which the cost exceeds a set threshold, again defined as part of the criteria for the information risk management programme.

The component parts of a business case are described in greater detail in [Chapter 7 – Risk Treatment](#).

SUMMARY

In this chapter, we have examined how we can assess the likelihood of a threat occurring, how we assess the resultant risk, and how we evaluate this risk in order to treat it. The next chapter describes the process of risk treatment.

7 RISK TREATMENT

Now we have completed the risk assessment process, it is time to begin to consider how to deal with the risks we have identified. The actions we take to treat risk are referred to as controls.

A control is any measure or action that modifies risk. Controls include any policy, procedure, practice, process, technology, technique, method or device that modifies or manages risk. Risk treatments either become controls or modify existing controls once they have been implemented.

However, some controls may monitor risks without actually modifying them in order to ensure predictability of the process. Actual risk modification then only occurs if the monitoring activity detects results that deviate from those expected.

Controls are the tools we use to take a level of inherent risk and modify it to a level that falls within the organisation's risk appetite, at which point the organisation is willing to accept the residual (if any) risk.

This chapter begins by taking an overview of the principal options for risk treatment.

Firstly, at the strategic risk treatment level, we have four options:

- To avoid or terminate the risk.
- To transfer or share the risk.
- To reduce or modify the risk.
- To accept or tolerate the risk.

At the end of this process, there may remain some 'residual' risk that we cannot treat, simply because we have exhausted all other possibilities or because the cost of further treatment would be more costly than the financial losses if the risk came to fruition. In this case, we must accept this residual risk, but subject it to ongoing monitoring.

In earlier chapters, we discussed the criteria for information risk management, and noted that there is a level of risk for each type of information asset, known as the risk appetite, above which the organisation will wish to treat the risk, and so it is important to understand that there will be situations in which more than one choice of treatment may be required in order to take the level of risk to a point that is acceptable to the organisation, bringing it below the risk appetite level.

Next, for each strategic risk treatment option, there are one or more tactical options:


- detective controls;
- corrective controls;
- preventative controls;
- directive controls.

Finally, at the operational level, we have three types of control:

- physical controls;
- procedural controls;
- technical controls.

Figure 7.1 illustrates all the possible combinations of risk treatment, and we shall deal with each of these in turn, but for the moment let us look at the start of the process for deciding which strategic option or options are most suitable.

Figure 7.1 The overall scheme of risk treatment options




STRATEGIC RISK OPTIONS

Strategic risk management controls themselves do not actually treat the risk, but set the approach for treatment, rather like giving travel directions – ‘Head north/south/east or west’. Figure 7.2 illustrates the process for determining the most appropriate strategic risk option.

Risk avoidance or termination

Firstly, we will examine whether or not we can avoid or terminate the risk. This implies either not commencing an activity that would incur risk or to cease doing it if it has already commenced.

Figure 7.2 The strategic risk management process

For example, if the organisation were considering building a new factory and we had identified that there was a significant risk due to the building site being within a known flood plain, we would recommend that the organisation found a more suitable location. As long as the organisation accepted this recommendation, that particular risk would be terminated immediately.

However, it should be remembered that ceasing or terminating some risks may introduce others. In the example above, if the building the factory in the original location was halted, there would be a requirement either to find a more appropriate location, or live with the results of not building at all. Either of these options would potentially introduce additional risks.

As a slightly different example, if the organisation were in the process of moving large amounts of information to a cloud service provider and we had identified that this would place the organisation's confidential customer information in a jurisdiction that does not meet the requirements of the data protection legislation, we would recommend that the project be halted.

This would terminate that particular risk, but there would be two secondary effects resulting from a decision to stop the project. Firstly, there would be some cost for any work that had already been carried out, and possibly some contractual penalties; and secondly, the organisation would be left with the problem of finding a replacement cloud service supplier if it still wished to outsource the information, and again there would be a cost involved in doing this.

Risk transfer or sharing

If terminating the risk is not an option, the next option is to decide whether the risk can be transferred or shared with a third party.

This might seem to be a simple way of disposing of a risk, but it is not as straightforward as it may first appear. While an organisation may transfer or share the actual treatment of the risk to a third party, the responsibility and accountability for the risk remains with the organisation itself, and in some situations it may not be possible to transfer or share the entire amount of the risk.

An example of the first of these possibilities is where an organisation places its confidential customer information with a cloud service provider whose systems are attacked by hackers and the information is made public. Although the outsourcing organisation would almost certainly have a legitimate claim of negligence against the cloud service provider, the responsibility for the information leakage would still lie with the organisation, and the Information Commissioner would expect to penalise the organisation itself rather than the cloud service provider.

So, although the risk has been shared, there remain a number of issues that the organisation faces, including secondary and indirect risks. All these must be taken into account when transferring risk.

An example of the second instance might be where an organisation takes out an insurance policy against the revenue losses incurred if their information systems fail. An insurance company might be happy to take on the entire risk, but the premiums to cover all possible losses might exceed the organisation's budget. In order to reduce the premiums to an acceptable level, apart from demanding a number of guarantees such as fully tested DR systems, the insurer might well set an upper limit on the possible payout in the event of a claim, so the organisation would be faced with accepting some residual risk.

Risk reduction or modification

In the past, this was often referred to as risk treatment, but the definition has now become more precise so as to avoid confusion with other forms of treatment.

When an organisation examines the possibilities for reducing or modifying risk, it means reducing the impact, reducing the likelihood or sometimes a combination of both, since there are a number of different approaches to this at tactical and operational levels that may be used in combination to bring the level of risk down to below the organisation's level of risk appetite.

One typical example of risk reduction is to reduce the likelihood of unauthorised remote access by replacing user ID and password authentication with two-factor authentication using a smartcard, security token or biometric scanner as well as the existing user ID and password. This would strengthen the authentication process considerably and reduce the likelihood of a successful intrusion, but the organisation would then have to balance this benefit against the costs of deployment and ongoing support, together with the additional risk that a lost or stolen token could present a new threat!

Another example is to reduce the impact of theft of a company laptop by encrypting the entire hard drive. The capital value of the laptop would still have been lost – although the impact of this might have been reduced under an insurance policy – but the

information contained on the laptop would be secured and almost impossible for the thief to recover.

Again, organisations might have to accept some residual risk – the replacement cost of the laptop in this last example would be covered, but there would be some additional expenditure in configuring it and installing replacement software.

Risk acceptance or tolerance

The final strategic option is to accept or tolerate the risk. Ordinarily this only happens when a risk is too costly to treat by any other means or, following other forms of risk treatment, has reached a point where no further treatment is possible.

The most important aspect of risk acceptance is that it differs entirely from ignoring risk, which should never be an option under any circumstances. Accepted risks must always be documented as such and reviewed either at intervals, in case the level of risk has changed in the meantime, or if there has been a sudden change in either the impact or the likelihood that produced the risk in the first place.

Not all forms of strategic control involve the use of all forms of tactical control, as we shall see in the next section.

TACTICAL RISK MANAGEMENT CONTROLS

Using the earlier analogy, tactical risk management controls are similar to suggesting the route – ‘Take the motorway as far as junction 20’. They consist of four different types: detective, preventative, directive and corrective.

As with strategic controls, tactical controls themselves do not actually treat the risks, but determine a more specific course of action.

Detective controls

Detective controls are those controls that advise or warn the organisation that an incident is taking place, but that is all they do – they themselves do not change either the impact or the likelihood of any risk, but operate alongside other types of tactical controls.

If action needs to be taken as a result of the warning, then this must necessarily be as a separate control (probably corrective) initiated either automatically by the detecting system or by a member of staff who is alerted by the warning.

An example of a detective control is an alarm generated when an intruder forces open a security gate or door.

Preventative controls

Preventative controls are designed to stop an incident from taking place before it has begun. The choice of any subsequent operational control will determine the actual steps

to be taken, but preventative controls will ultimately reduce or eliminate the likelihood of an incident occurring, and may therefore reduce or terminate the risk.

An example of a preventative control is the lock on a door or window to an area to which an intruder might otherwise have free access.

Directive controls

Directive controls are totally instruction-based. They comprise policies, procedures, processes and works instructions, all of which dictate what must or must not be done.

As with detective controls, directive controls do not change either the impact or the likelihood of an incident occurring, since they only dictate what may or may not be done. If people follow the organisation's policies, then the controls will be successful, but if they do not, the controls will not be effective. For this reason, many organisations like to couple directive controls with other types, such as preventative or corrective, in order to enforce the policies.

An example of a directive control is a policy stating that passwords must consist of a minimum number of alphanumeric characters, sometimes requiring both upper and lower case and so-called 'special' characters such as hyphens. This could be reinforced by a preventative control within the computer's operating system that forced users to update their passwords according to the constraints of the policy.

Corrective controls

Corrective controls are those that, through appropriate operational controls, will make a difference either to the impact or to the likelihood of a risk. Corrective controls are invariably introduced after an event of some kind has occurred or has been detected, and their purpose is to fix the problem. An example of a corrective control is the requirement to update a firewall rule in response to a newly discovered threat.

Not all forms of tactical control involve all forms of operational control, as we shall see in the next section.

OPERATIONAL RISK MANAGEMENT CONTROLS

Using the earlier analogy, operational risk management controls are similar to providing more detailed directions to the destination, such as 'Now take the A4303 east, then the A426 south for 2½ miles'.

Operational controls come in just three varieties: procedural, physical and technical, also known sometimes as people, environmental and logical controls, respectively. As with tactical controls, operational controls may be used singly, or in conjunction with others in order to minimise a risk.

Procedural controls

Procedural controls, as the name suggests, simply put procedures in place to set out actions that users and technical staff must or must not take in any given situation. Procedural controls on their own do not change either the impact or the likelihood of a risk, but do so only where users follow them.

An example of a procedural control is that of a clear desk policy stating that users must leave no materials (for example books, papers or electronic media) on their desk when out of the office. This is in fact a directive/procedural control, and is frequently used in conjunction with the requirement to lock a computer screen when the user is away from their desk.

Physical controls

Physical controls are those that address physical and environmental threats, and as such always change either the impact or the likelihood of the risk. Physical controls may contain a technology element, but this is invariably unrelated to the information itself, merely pertaining to a technical solution to a physical risk.

An example of a physical control (actually detective/physical) is a CCTV system to monitor the perimeter of a building.

Technical controls

Technical or logical controls refer generally to those controls that are directly related to the technology that underpins (or is) the information-supporting infrastructure. Technical controls may be implemented either in hardware or in software, and frequently both.

An example of a corrective/technical control is that of configuring a virtual local area network (VLAN) environment in order to segregate live and test networks.

EXAMPLES OF CRITICAL CONTROLS AND CONTROL CATEGORIES

The following sections list the chief controls suggested by:

- the Centre for Internet Security Controls Version 8;
- ISO/IEC 27001:2017;
- National Institute for Standards and Technology (NIST) Special Publication 800-53 Revision 5.

The Centre for Internet Security Controls Version 8

A number of organisations have published a list of the 18 most critical security controls, based upon the Centre for Internet Security Controls Version 8.¹

¹ See <https://www.cisecurity.org/controls/>.

While this is not a comprehensive list of controls, it does provide a good starting point for organisations that have conducted their risk assessments but are unsure where to begin with risk treatment. These controls are described more fully in [Appendix D](#).

No.	Control title
	Basic controls
1	Inventory and Control of Enterprise Assets
2	Inventory and Control of Software Assets
3	Data Protection
4	Secure Configuration of Enterprise Assets and Software
5	Account Management
6	Access Control Management
7	Continuous Vulnerability Management
8	Audit Log Management
9	Email and Web Browser Protections
10	Malware Defences
11	Data Recovery
12	Network Infrastructure Management
13	Network Monitoring and Defence
14	Security Awareness and Skills Training
15	Service Provider Management
16	Application Software Security
17	Incident Response Management
18	Penetration Testing

ISO/IEC 27001:2017 controls

Reference	Title (number of controls)
A.5	Information security policies (2)
A.6	Organisation of information security (7)
A.7	Human resource security (6)
A.8	Asset management (10)
A.9	Access control (14)
A.10	Cryptography (2)
A.11	Physical and environmental security (15)
A.12	Operations security (14)

A.13	Communications security (7)
A.14	System acquisition, development and maintenance (13)
A.15	Supplier relationships (5)
A.16	Information security incident management (7)
A.17	Information security aspects of business continuity management (4)
A.18	Compliance (8)

Although the primary ISO standard for information risk management is ISO/IEC 27005:2018, it contains no detailed information on suitable tactical or operational controls for risk treatment, restricting itself instead to the strategic level only. Instead, ISO/IEC 27001:2017 provides a comprehensive list of 114 separate operational level controls, grouped into 14 categories in its [Appendix A](#). A more detailed description of the controls can be found in ISO/IEC 27002:2017 in its sections 5–18.

These control categories are expanded and described more fully in [Appendix D](#).

NIST Special Publication 800-53 Revision 5

Identifier	Family (number of controls)
AC	Access Control (24)
AT	Awareness and Training (5)
AU	Audit and Accountability (16)
CA	Security Assessment and Authorisation (8)
CM	Configuration Management (11)
CP	Contingency Planning (12)
IA	Identification and Authentication (12)
IP	Individual Participation (6)
IR	Incident Response (10)
MA	Maintenance (6)
MP	Media Protection (8)
PA	Privacy Authorization (4)
PE	Physical and Environmental Protection (22)
PL	Planning (9)
PM	Program Management (32)
PS	Personnel Security (8)
RA	Risk Assessment (9)
SA	System and Services Acquisition (20)
SC	System and Communications Protection (42)
SI	System and Information Integrity (19)

Although the primary NIST publication on information risk management is Special Publication 800-30, it contains no detailed information on risk treatment or the selection of controls. NIST Special Publication 800-53 Revision 5² lists 283 separate operational level controls, grouped into 20 categories, and maps them against ISO/IEC 27001:2017 controls in its [Appendix I](#).

These control categories are expanded and described more fully in [Appendix D](#).

SUMMARY

In this chapter, we have examined strategic, tactical and operational controls, and have looked in greater detail at some of the more critical controls and control categories as identified in a number of international standards.

Before we can begin applying these controls, we will probably be required to present the results of the information risk management programme to date in order to secure funds and people resources to undertake the risk treatment activities we have selected. This part of the programme begins with developing and presenting one or more business cases.

² Joint Task Force (2020) Security and Privacy Controls for Information Systems and Organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-53, Rev. 5. <https://doi.org/10.6028/NIST.SP.800-53r5>.

8 RISK REPORTING AND PRESENTATION

The risk treatment requirements obtained during an information risk management programme should never be actioned without having the full approval and support of senior management (often at board level), so that the correct levels of funding and people resources can be allocated without causing problems for the organisation.

One of the main avenues for establishing the support of senior management is the preparation and presentation of business cases, which will set out the risks in terms the business will understand and make clear the costs of remediation.

Following the approvals to proceed, the process of detailed decision-making, planning and implementation may begin. There may also be a spin-off in the form of a business continuity exercise, which might include DR planning.

The process of communicating within the information risk management programme is extremely important, and serves a number of purposes. It allows the information risk management programme manager to:

- Maintain a two-way flow of information between the programme manager and those stakeholders who are closely involved in the process of impact, threat and vulnerability assessments.
- Keep the organisation's senior management and other stakeholders informed both of general progress and of specific actions regarding the highest risks encountered or those that are particularly costly to treat.
- Flag up any new risks deemed to be very severe and requiring immediate attention.
- Present business cases requesting approval of recommendations and funding.
- Report on those risks that have been successfully treated and those that remain untreated.

It is often said that senior executives will never understand information risk, but this is not entirely correct. They may not understand the technicalities of information security, but business risk is something they will definitely understand, so the streetwise information risk management programme manager will ensure that all reporting is couched in terms of risk to the organisation and the business benefits to be gained by avoiding, transferring, reducing or accepting it.

BUSINESS CASES

Business cases are a standard vehicle for demonstrating a genuine need to carry out some form of activity that will require senior management approval. They are generally used in those circumstances in which a significant financial spend is proposed (beyond that of day-to-day budgets), and in the case of an information risk management programme will most frequently be brought into play in gaining approval to carry out risk reduction or modification, although some aspects of risk transfer or sharing will also require board-level agreement.

In some situations, the business case might present senior management with a clear and simple 'yes or no' decision, while others might involve a number of options with a recommendation for a specific approach that, in the view of the information risk manager, represents the most appropriate solution combined with good value for money. In the case of the latter, the senior management team will be required to choose their favourite option, and the contents of the business case will heavily influence this choice.

It follows, therefore, that the business case should be as comprehensive and compelling as possible, so that senior management's decision-making process is made completely straightforward and that they make a fully informed choice.

There is no generic set format for a business case. Some organisations have their own template, whereas others allow a free format of presentation. In this section, we suggest some of the essential components of the business case and describe how best to present it.

Many people will be familiar with the name of Robert Maxwell, who ran a vast publishing business empire that included newspapers such as *The New York Times*, *The Daily Mirror*, *The Scottish Daily Record* and *The European*. Whatever his faults, he adopted a very simple approach to business cases: he relied on his senior management team to pull together the best advice and to present this to him in as short a time as possible.

When it came to receiving his formal approval, a single sheet of A4 was all he needed to read, written in a 14-point Courier typeface, with 1.5 line spacing, and a signature line near the bottom of the page followed by the words 'Approved. R Maxwell, Chairman'. Supporting information was always stapled behind this, but he rarely studied it.

Most senior executives do not have the time to read large amounts of detail and, since information security is not usually their strongest point, might find it difficult to follow. What they do need are clear, concise facts: the issue, the proposed solution, the costs, the benefits to be gained and, if necessary, the downsides of not choosing the recommended option.

It is suggested that a business case document should contain the following sections:

- An introductory executive summary – preferably on a single page.
- The benefits to the organisation of undertaking the work.
- A synopsis of the goals and objectives and the main risks threatening the information assets, together with the likely impacts or consequences faced by the organisation if the threats were to materialise.
- A synopsis of the proposed solution, together with reasoning as to how and why this would eliminate the risk or reduce it to a level acceptable to the organisation, together with the timescales for doing so.
- A financial breakdown, showing both capital and operating expenditure required over a three to five year period, with resources for premises, equipment and people clearly identified.
- A high-level project overview, including critical success factors.
- An implementation plan, including resources required, a timeline and key milestones.

Many organisations prefer a personal briefing as well as a business case document, in which case a slide presentation – probably no more than 10 slides – should be prepared and delivered by a programme representative who feels comfortable presenting to very senior managers and who can also answer penetrating questions without the need to refer to detailed notes.

Whichever approach is taken, the person presenting the business case would be well advised to socialise the business case beforehand with as many members of the approving committee as possible, so that it is approved 'on the nod'. This approach has another advantage, in that many of the questions that might be asked during a presentation will either be known or answered beforehand, and any last-minute changes to the business case that will assist in gaining approval can be included.

RISK TREATMENT DECISION-MAKING

The decision-making process for risk treatment follows a logical path. It begins by identifying the strategic option or options that the organisation should take – risk avoidance or termination; risk transfer or sharing; risk reduction or modification; and risk acceptance or tolerance – and this part of the process will have been taken care of during the final stage of risk assessment, risk evaluation.

The next step for each of the chosen strategic approaches is to identify the tactical options. These will depend completely on the strategic approaches, but will be as follows:

- Risk avoidance or termination presents both preventative and detective options, but these are invariably used together, since the preventative course of action will require ongoing (detective) monitoring to ensure that further action is taken if something changes.

- Risk transfer or sharing has both directive and detective options and, again, these are used together, since any shared risk also requires ongoing monitoring to ensure that it is working as expected.
- Risk reduction or modification is the most complex, as it can involve directive, detective, preventative and corrective actions, again with ongoing monitoring.
- Finally, risk acceptance or tolerance follows the detective approach with ongoing monitoring, and it is very important to repeat that no risk, no matter how trivial it might appear, should ever be ignored.

Having identified the tactical risk treatment options, the final stage is to identify the operational options:

- Risk avoidance utilises both detective and preventative controls. The preventative controls will almost always be physical or technical, and the detective controls will always be procedural.
- Risk transfer's detective controls will usually be procedural (in the case of insurance, for example), or may be both technical and procedural in the example of outsourcing operations to a cloud service provider.
- Risk reduction uses all three types of operational control. Preventative controls can be either physical (e.g. security barriers) or technical (e.g. firewalls). Corrective controls can be physical or technical (e.g. ensuring that bug fixes are applied to software), or procedural. Directive controls are instructive, and therefore always procedural, and include policies, processes and works instructions. Detective controls can be physical (e.g. CCTV systems), technical (e.g. intrusion detection software) or procedural (e.g. system activity monitoring).
- Finally, risk acceptance requires just detective procedural controls to provide ongoing monitoring of threats to ensure that the level of risk has not changed.

RISK TREATMENT PLANNING AND IMPLEMENTATION

It is quite conceivable that many of the risks requiring treatment as part of the information risk management programme can undergo treatment as an integral part of the programme. However, some risks might require extensive (or expensive) treatment, and as such may need to be treated as a project or programme of work in their own right.

However, although the implementation may be carried out under a separate project or programme, progress reporting of the implementation should remain part of the original information risk management programme so that the audit trail is complete.

Such a project requires the setting of goals, objectives, scope and milestones, which, given the controls recommended and agreed earlier in the information risk management programme, should be relatively straightforward to define.

The risk treatment plan should commence with the production of a prioritised list of risks for treatment, which includes realistic estimates of the length of time these might take to achieve, the approximate cost of the treatment and the resources required

(including the name of the responsible person) for doing so. By totalling the number of completed risk treatments and the running costs, additional information can be reported to senior management.

Regardless of whether the project is to be managed from within or outside the main information risk management programme, resources, especially people and funding, must have been agreed and committed by the organisation. This will include a suitably qualified alified project manager, who may be a different entity from the information risk management programme manager, particularly if the project is significant in its scope; for example, if the agreed control is for the provision of an entire backup data centre with high-availability standby systems, this would be a major project in its own right, and would certainly require at least one dedicated project manager, if not several.

However, even if the remedial work to implement the agreed controls is relatively minor, each individual control should be considered as a task within an overall project, so that it can have resources assigned to it and be tracked to completion and sign-off.

BUSINESS CONTINUITY AND DISASTER RECOVERY

Occasionally, the controls recommended may be very wide-ranging, such as the need for business continuity management (BCM) and DR arrangements, which are specialist subject areas in their own right. However, it is worth providing a brief description of both approaches.

Business continuity

The concept of BC became better known in 2006 with the introduction of the first full standard, BS 25999-1, the Code of Practice, and then its Specification, BS 25999-2, in 2007. Prior to that, there had only ever been a publicly available specification, PAS 56, published in 2003 and developed from an early Business Continuity Institute (BCI) *Good Practice Guidelines* document.

The two BS 25999 standards were superseded in 2012, and the international standard *ISO 22301:2019 – Societal security – Business continuity management systems – Requirements* now applies instead.

BC is defined as 'The capability of the organisation to continue delivery of products and services at acceptable predefined levels following a disruptive incident' (ISO 22301:2019).

BC applies to a number of key areas within an organisation, and so is considered to be a holistic approach to risk management. It includes:

- The people employed by the organisation, together with its contractors.
- The organisation's premises, whether these are offices, factories, warehouses or other types of building.
- The organisation's processes and procedures.
- The technology that underpins the organisation's activities.

- The organisation's information in both physical and electronic forms.
- The organisation's supply chain.
- The organisation's delivery chain.
- Any other stakeholders that have an interest in the organisation.
- The organisation's responsibilities (if any) in the event of civil emergencies.

At first sight, it would appear that information is just one of these areas, but it actually cuts across all of the remainder, and hence the principles explored in information risk management are fundamental to the discipline of BCM.

The Business Continuity Institute Good Practice Guidelines 2018

Founded in 1994, the BCI has always been at the forefront of business continuity standards development, and was instrumental in the first UK specification PAS 56, published in 2003. Its members have subsequently taken a leading role in the later development of BS 25999 in 2006/7 and ISO 22301 in 2012 and beyond.

Over the years, the Institute has developed a set of good practice guidelines (GPGs) that define the generic approach to BCM in six distinct stages, or so-called Professional Practices (PPs):

PP1 Policy and Programme Management – this is the beginning of the overall BCM life cycle, and defines the organisation's policy for BC: how it will be implemented, managed and tested.

PP2 Embedding – it is important that the culture of BCM is embedded into day-to-day operations within an organisation.

PP3 Analysis – in earlier versions of the *Good Practice Guide*, this was known as Understanding the Organisation, and assesses the organisation's overall objectives, how it functions and the internal and external context within which it operates. It includes the risk assessment process of risk management.


PP4 Design – formerly known as Determining Business Continuity Strategy, this area recommends suitable approaches (both strategic and tactical) to recover from disruptive events and to provide continuity of operations.

PP5 Implementation – this area was previously known as Determining and Implementing a BCM Response, and carries out the recommended and agreed approaches through the development of business continuity plans (BCPs). Together with Design, this area aligns with the risk treatment portion of risk management.

PP6 Validation – validation was originally referred to as Exercising, Maintaining and Reviewing, and deals with the validation of BC plans through tests and exercises to ensure that they are fit for purpose and would be effective in disruptive situations.

Professional Practices 1 and 2 are described as management practices, whereas Professional Practices 3 to 6 are described as technical practices. The BCI's life cycle diagram illustrates this graphically in [Figure 8.1](#).¹

¹ The BCI life cycle diagram is included courtesy of the BCI.

Figure 8.1 The BCI life cycle

These are by no means mandatory requirements, but most BC practitioners – and not only in the UK – will follow them, since they provide considerable assistance when an organisation wishes to become compliant with the standard and to achieve accreditation against it.


Business continuity plans

BCPs produced will normally include:

- IM plans, which deal with the immediate aftermath of business-disrupting incidents and which can include information security incidents as well as civil emergencies, strikes and pandemics.
- BC plans, which take the process following incident management through recovery to a normal, near-normal or new normal state.
- DR plans, which deal mainly with restoring the capability of IT systems.
- Business resumption (BR) plans, which are the final stage of recovery from incidents and which take operations back as closely as possible to the same state as they were in prior to the incident.

Although BC itself is generally thought of as being a form of risk reduction or modification, a BC programme of work may well make use of all forms of strategic, tactical and operational controls in order to achieve its objectives. [Figure 8.2](#) illustrates the generic BC incident timeline.

Once IM, BCM, DR and BR plans have been developed, they must be tested in order to prove their fitness for purpose.

Figure 8.2 The generic business continuity incident timeline

BC introduces some terminology that is not generally used in information risk management. However, when implementing a BC strategy as part of the treatment process for an information risk management programme, it is worthwhile being aware of these terms:

Recovery point objective (RPO). The point to which information used by an activity must be restored to enable the activity to operate on resumption.

Recovery time objective (RTO). The period of time following an incident within which products, services or activities must be resumed or resources must be recovered.

Maximum acceptable outage (MAO). The time it would take for adverse impacts, which might arise as a result of not providing a product/service or performing an activity, to become unacceptable.

Maximum tolerable data loss (MTDL). The maximum loss of information (electronic and other data) that an organisation can tolerate. The age of the data could make operational recovery impossible, or the value of the lost data is so substantial as to put business viability at risk.

Maximum tolerable period of disruption (MTPD). The time it would take for adverse impacts, which might arise as a result of not providing a product/service or performing an activity, to become unacceptable.

Minimum business continuity objective (MBCO). The minimum level of services and/or products that is acceptable to the organisation to achieve its business objectives during a disruption.

Various types of test may be undertaken:

- Communications tests, in which the contact procedures are invoked to ensure that members of the various teams responsible for managing the situation can be contacted and instructed either to attend the crisis management centre or to join an audio or video conference call.

- Desktop read-throughs, in which the plans are scrutinised by all members of the various response and recovery teams in order to verify that all necessary activities have been identified, that they are in the correct order and that all interdependencies have been considered.
- Scenario-based exercises, in which the BC manager develops an imaginary or real-world event-based scenario for the response and recovery teams to work through as if it were an actual event. Lessons learnt from this type of exercise will often refine the plans as gaps and overlaps are identified.
- Full-scale exercises, again usually scenario-based, in which many or all of the organisation's staff are involved to some extent in order to verify that the plans do actually work in situations as close as possible to real-world events. Such exercises will disrupt the organisation's business activities, and will often only be performed under exceptional circumstances.


BC is invariably conducted as a separate programme of work from that of information risk management, since it may have much wider implications for the organisation, especially in terms of the resources required to operate the programme and to exercise the plans.

Disaster recovery

DR is a specialised subset of BC, and is generally used to refer to the arrangements put in place to provide backup or recovery computing facilities, although it can refer to other forms of technical processing. In our Glossary of Terms, we describe disaster recovery as 'A coordinated activity to enable the recovery of ICT [information, communications and technology] systems and networks due to a disruption'.

Some organisations make use of system hardware normally used for software application testing purposes to provide DR: sometimes on a one-for-one basis, so that the standby hardware is identical to the system being replicated, sometimes on a one-to-many basis, where one backup system can be used to provide DR for a number of live systems. [Figure 8.3](#) illustrates the overall structure of DR operations:

Figure 8.3 Overall structure for disaster recovery



Platform disaster recovery

Platform DR generally involves the use of one or more of the three following types of facility.

Cold standby platforms These consist of bare computer systems and associated communications equipment. They may have an operating system loaded, but little else. The organisation or its outsourced DR partner will be responsible for loading any additional operating systems and applications software required in order to operate the system in the same way as the one it is replicating. In addition, all data must be restored from backup media, and the organisation will need to take into account any patches or software updates that have been issued.

Because these systems are very basic, they represent the lowest cost to an organisation, and take the longest amount of time to bring up to full operation.

Warm standby platforms Warm standby systems invariably have their full operating system and key applications loaded, and may have some backed-up data loaded as well. However, unless the system has been maintained in a fully 'ready' state, the organisation will need to take into account any patches or software updates that have been issued since the system was originally configured. Data will have to be brought fully up to date by restoring from the most recent backups.

Warm standby systems are more expensive to provide than cold standby systems, and can normally be brought into service much more quickly.

Hot standby/high-availability platforms At the top end of the DR range, there are hot standby or high-availability platforms, which are always maintained in a fully ready state from the point of view of operating systems and application software. Data will also be fully up to date, since the system being replicated will copy across all data onto the standby system.

These vary in type and cost, as can be seen from [Figure 8.4](#). Availability is measured in 'nines', with five nines, that is, 99.999%, general availability being the highest, which allows for five minutes' downtime in any 12-month period. Unsurprisingly, higher availability comes with a greatly increased cost, and each 'nine' added would probably increase the cost tenfold.


In cases where two systems operate jointly to deliver the service, data are copied between the live and the standby system in one of two ways:

- Asynchronous replication, in which each block of data is sent from the live system to the standby system, receipt is confirmed, but the live system continues to process data in the meantime.
- Synchronous replication, in which each block of data is sent from the live system to the standby system, but the live system waits before continuing to process data until receipt is confirmed.

Synchronous replication is slightly slower than asynchronous replication, but has greater reliability, since no data can be lost at the point of switchover. The distance

between the live and standby locations cannot currently be greater than around 200 km (125 miles), and typically uses a direct fibre-optic link, which guarantees capacity as well as reliability.

Figure 8.4 Cost versus availability



High-availability systems are by far the most costly to operate, but for organisations such as banks, large online retail organisations, airlines and the like, failure of service and possible loss of information is simply not an option.

In conjunction with platform DR, organisations should take into account four key areas:

- Resilient power, in which uninterruptible power supplies (UPS) ensure that power is always available to the platforms. UPS systems feed power directly from the incoming mains supply during normal operation, and if and when the incoming supply fails, batteries take over and generate the appropriate levels of power through static inverters. In order to make the power supplies completely resilient, standby generation can also be used, but resilience of fuel supplies for these must then be taken into account.
- Cooling systems are also highly necessary to take away excess heat and allow the systems to operate at a comfortable temperature. It is common practice to allow more cooling systems than are necessary to maintain a constant temperature to allow for failures, and this is referred to as N+1, N+2, and so on.

- Since system hardware can run at temperatures higher than staff would find a comfortable environment, some organisations run the equipment rooms of their data centres slightly warmer, which can reduce cooling costs without a detrimental effect on the systems.

Systems and service monitoring is always required, so that remedial action can be taken as soon as there is a failure in any part of the service being provided. In larger organisations, the internal and external data networks are usually monitored in addition to platforms and services.

- Vendor support is the final key requirement for all operations of this kind, simply because the systems vendor organisations can often supply additional or replacement platforms either at the data centre itself or, increasingly, through cloud services.

Data resilience

Although data storage has moved on considerably in recent years, magnetic media of one kind or another remains the most cost-effective technology, although it is rarely the fastest. There are several commonly used methods of providing resilient data storage:

- Redundant Array of Inexpensive Disks (RAID), which uses several different technologies to achieve various levels of resilience, including error correction, striping and mirroring, some of which will permit a faulty disk drive to be exchanged without loss of any data whatsoever. There are currently around 12 different levels of RAID, and different operating systems (for example, Unix, Linux, MacOS and Microsoft Windows) support slightly different combinations of RAID level. In general, the RAID systems that use more disk drives provide greater resilience, but there are advantages and disadvantages of all types in terms of ease of implementation, cost, resilience and performance.
- Direct Attached Storage (DAS), in which RAID arrays are connected directly to a system using a standard RAID controller. DAS is less resilient than other methods, in that access must always be via the system and controller to which it is connected, which presents a potential single point of failure.
- Network Attached Storage (NAS). Unlike DAS storage, NAS storage does not rely on a separate system for connection but instead uses its own proprietary operating system and connects directly to a network. This means that it can be accessed in a more flexible manner, making deployment much more straightforward. With the ongoing increase of local area network (LAN) speeds, this makes NAS a very attractive option. NAS storage also has the advantage of being able to operate at block level as well as at file level, making it independent of the host operating system.
- Storage Area Networks (SANs) make use of highly resilient fibre-optic links between the host systems and their storage arrays, which are also designed to be highly resilient. Operating completely at the block level, they are entirely independent of the host operating system.
- Cloud-based data storage is increasingly popular, providing a fully managed storage service in which the supplier will take responsibility not only for storage capacity, but also for the (hopefully redundant) links into the customer's network.

Cloud-based storage has become so low in cost for a given volume of data stored that it is now very popular at the individual consumer level as well as that of larger organisations.

Application resilience

We are all used to experiencing applications on a home or office computer, and these occasionally fail but usually impact only the computer user. At a corporate level, application failures will affect many users, and those that provide a service for online use (for example online banking applications) can affect very large numbers of people if they fail. For this reason, application resilience is key to such services, and can be delivered in one of three ways:

- Clustered file systems distribute the file system across multiple systems or nodes, each of which holds part of the overall file system, but appears to the end-user as a single entity. They provide a very high throughput, but at high cost of deployment and management.
- Application clusters, which are similar to clustered file systems but distribute just the application software across multiple systems, provide increased performance and availability. The host computer sees the application as a single resource, which requires the use of a so-called 'heartbeat' between all the systems in the cluster so that a node that has failed can be flagged as such and its resources transferred across the remaining online nodes. Application clusters are quite complex to deploy.
- Computer cluster services include two key options. The first is load balancing, in which the network traffic is distributed as evenly as possible across multiple servers that all share a single Internet Protocol (IP) address, and the load balancing software uses a rule set to decide which server should receive the next request. The second is failover cluster services, which has two options. In the first option, all nodes in a cluster have access to shared resources (such as DAS, NAS, SAN or cloud), and a heartbeat similar to that used in application clusters is used to assist in the control of access to nodes. This is sometimes referred to as the 'share everything' approach. In the second option, known as 'share nothing', only one node in a cluster has access to the resources at any one time.

Site recovery

Organisations may choose to replicate a complete equipment site in a physically separated location, usually around 30 miles (48 km) apart. This option is invariably expensive, both to provision and to maintain, but offers organisations the opportunity to provide a fully resilient service, not only to the organisation itself but also to its customers and suppliers.

In practice, organisations that make use of site recovery either do so through the agencies of a third party or will make use of space in other offices, warehouses, factories and the like. This is frequently the case for those organisations that operate a large data centre or telecommunications hub, in which site recovery for one location can be relatively straightforward to provide in another site having similar infrastructure and environmental facilities.

DISASTER RECOVERY FAILOVER TESTING

The testing of DR plans generally follows one of two paths:

- 'Fire drills', which normally refer to the testing of DR plans in which the process of bringing the standby systems into full readiness is tested but stops short of an actual switchover from live to standby systems, since, depending on the type of standby system implemented, it might be disruptive to the organisation's business.
- Full switchover tests, in which not only are the standby systems brought into a state of full readiness, but also a switchover from live to standby systems is performed and the performance of the new system status is verified as being at an acceptable level.

Again, depending on the type of standby system implemented, a full switchover test might be disruptive, but it is the only way in which the organisation can be completely certain that its DR arrangements are fully working. However, if the standby system has been correctly implemented, everything should failover without interruption.

SUMMARY

In this chapter, we have examined the need for and the importance of business cases in gaining support from senior management for the information risk management programme, together with the processes of risk treatment decision-making, planning and implementation. Finally, we have looked at the requirements for business continuity planning, based on the BCI's approach, and, where applicable, various types of solution to disaster recovery planning and testing.

We will look next at communicating, monitoring and reviewing activities.

9 COMMUNICATION, CONSULTATION, MONITORING AND REVIEW

Although we have already discussed the need for business cases, a greater part of the information risk management programme will require input and agreement from many stakeholders, and it is important that the programme team are able to undertake discussions with them in such a way as to ensure the programme's success.

Communication is key here, as is the ability to consult with stakeholders before, during and following the programme's activities.

This chapter deals with the communication and consultation activities required, and also focuses on the need to review risks once they have been treated and the ongoing monitoring of them to ensure that they remain at a satisfactory level.

SKILLS REQUIRED FOR AN INFORMATION RISK PROGRAMME MANAGER

Before we examine the requirements for communication, consultation, monitoring and review within an information risk management programme, let us take a few moments to examine the kinds of skills an information risk programme manager will require.

Business skills

A sound knowledge of the organisation's business is essential when commencing an impact assessment. This will include not only an understanding of the key activities, products and services, but also how the organisation goes about its business, including areas such as sales and marketing, order processing and fulfilment, procurement, manufacturing, IT, finance and human resources.

Technical skills

Although the person conducting the impact assessment will be unlikely to have a detailed technical knowledge of all the processes involved, this is not necessarily a bad thing. A general appreciation of technical issues is, however, a distinct advantage as it permits one to ask 'obvious' questions, the answers to which can occasionally highlight a potential problem.

Interviewing skills

Interviewing skills can be learnt, but if this option is not open, then the interviewer should take a few basic rules into account. Firstly, he or she should explain the purpose

of the discussion clearly, and verify that the interviewee has understood. Secondly, it is essential to have a prepared list of questions.

Often, one question will trigger another that was not on the list, so a degree of flexibility is also required. The interviewer should always listen more than he or she speaks as the objective is for the interviewer to discover information and for the interviewee to provide it.

Interviews should not take more than an hour or so as the interviewee may provide less valuable information if the process is lengthy, and the interviewer will have difficulty in absorbing sufficient information. If time runs out, the interviewer should conclude the discussion, write up the notes and continue at a later date. This also allows the interviewee to review the output and confirm the earlier discussion.

A pocket voice recorder is a really useful tool when conducting interviews, as it allows the interviewer to listen again to what was actually said, rather than what was noted down, which might be slightly different.

Analytical skills

Analysis of the output from interviews is required in order to ensure that all areas of the organisation have been covered and that the findings are also consistent across the whole organisation.

This is particularly important in situations where one activity or process can be dependent upon one or more others – the whole chain of events must be recorded, otherwise what might be a critical part may be overlooked.

Presentation skills

Presentation of the risk assessment results may well be to a senior management team who will look critically at the recommendations for work to be carried out and the resources (especially financial) required. For this reason, the final presentation must be clear, compelling, succinct and accurate, and information risk managers may find that they can benefit from training in this area.

Interpersonal skills

The skills we have already described are necessary in order for the information risk management programme manager to undertake the role, but more important than all of these are the interpersonal skills on which all the others depend. Interpersonal skills include:

- Verbal skills, which are all about what we say to other people and the way in which we say it.
- Non-verbal skills, which include what both we and other people do not say, but indicate in gestures and body language.
- Listening skills, or the way in which we hear, see and interpret both the verbal and non-verbal information given by others.

- Negotiation skills, or the way in which we work with other people to agree a mutually acceptable outcome.
- Problem solving skills, where we work with other people to identify and define problems, explore possible solutions and make recommendations.
- Decision-making skills, where we explore and analyse multiple options to make sound decisions.
- Assertiveness skills, which include communicating our values, ideas, beliefs, opinions, needs and wants but without aggressiveness.

COMMUNICATION

One of the most important components of any information risk management programme is that of communication. As with any project or programme, nothing useful will ever be delivered unless those involved communicate effectively with one another.

An experienced programme manager will be aware of the possibility of a 'lone wolf' member of the team who will be tempted to work independently of everyone else, and who may not be willing to share information or progress with others or occasionally to take on board what others have already done. This will invariably lead to duplication of effort and mistakes being made as assessments, analyses and evaluations are carried out, and possibly to recommendations made and decisions implemented in isolation and without taking full account of the larger picture.

Communication is also a two-way process. Information provided must be acknowledged, and information requested must be delivered. However, bottlenecks can and will occur, especially in larger organisations and in those with relatively rigid hierarchical management chains.

It is not uncommon for the progress of the message through the management chain to stop at some point, especially if there is bad news to be told; or for information to be either diluted or embellished, especially if the news could jeopardise someone's career or if a particular project is under threat.

The means and route of communication should have been agreed in the early stages of the programme. I would suggest that a simple reporting format is always the best, since it presents the information in a way in which all levels of management can understand without the need for large quantities of detailed information. These may mean little to a senior manager or director, but if further information is required, this can either be attached, if the need has been anticipated, or can be requested at the time.

One of the greatest risks over and above those we are considering in an information risk management programme is that of miscommunication. Verbal reports should really be considered as an informal briefing, although they are frequently taken as being both formal and factual: in response to an innocent 'How's it going?' question, the unwary information risk management programme manager might answer 'It's going well', and this could suddenly turn into a boardroom view that there are no issues.

Verbal communication is something to be extremely wary of providing unless written or graphical notes can back it up. Also, the potential audience can make a great deal of difference – briefing an internal audience might include concerns regarding particular information sets or projects, but it might well be wholly inappropriate to allow these to become known to external stakeholders. In such circumstances, it is always best to follow up even the briefest enquiry with a quick email that expands on and clarifies this, so that the enquirer's expectations are met and there is no doubt left as to what 'It's going well' actually means.

A common means of reporting is the so-called 'traffic light' method, in which any individual risk's status can be expressed as:

- GREEN – green risk status indicates that the risk has been successfully treated and is now at or below the level deemed acceptable by the organisation's risk appetite.
- AMBER – risks reported as Amber are currently in the process of being treated, and it is the view of the information risk management programme manager that treatment will ultimately be successful.
- RED – reporting a risk as Red can indicate a number of possibilities:
 - that the risk has not yet been treated;
 - that the risk is potentially untreatable;
 - that treatment has been unsuccessful and the residual risk remains above the level deemed acceptable by the organisation's risk appetite;
 - that the level of the risk is above that deemed acceptable to the organisation, but the recommendation is to accept the risk at that level.

It is also possible to show in a separate column whether or not treatment of the risk is on schedule, for example:

- Green status could represent the view that treatment was either on or ahead of schedule.
- Amber might be interpreted to mean that the risk was in danger of falling behind schedule, especially if activities were dependent upon the treatment of another risk that was already behind schedule.
- Red would mean that treatment of a risk was definitely behind schedule.

Further, the overall status or the information risk management programme can be summarised by representing the number of outstanding risks in any risk category in a similar way, and by agreeing in advance the percentage of risks that fall into each of the three categories, for example:

- Green – 0–35 per cent of individual risks at Red status.
- Amber – 36–65 per cent of individual risks at Red status.
- Red – more than 65 per cent of individual risks at Red status.

Alternatively, the format for reporting can use more descriptive wording, such as 'Business critical – requiring immediate attention' instead of 'Red'; 'Serious – requiring timely attention' instead of 'Amber'; and 'On track' instead of 'Green'. The choice of the exact wording should be agreed beforehand with senior management.

Whatever the method chosen to report the progress or status of an information risk management programme, the reporting manager should ensure that the content is jargon-free, since some of the readership may either be completely unfamiliar with the jargon used, or may interpret its meaning differently from that intended.

The report should ensure that the reader is fully informed about the topic, using clear and concise language and providing a balanced interpretation of the information risk management programme status, since to be too over- or under-optimistic can result in problems further down the line.

Where possible, the report must supply the reader with all the answers they might reasonably want regarding the topic. It is, of course, impossible to anticipate all the questions that might arise, but the information risk management programme manager should be able to anticipate the most likely questions, and ensure that answers to those asked but not anticipated are included in subsequent reports.

In cases where a decision will be required as a direct result of submitting the report, the wording should leave the reader in no doubt about the recommended choice, that a firm decision is required, the date by which a decision is required and what the next steps will be following the decision.

If presentation of the report is to be in the form of slides, then the information risk management programme manager should ensure that the presentation format is clear and free of over-ornate typefaces and graphics, so that the format does not dilute the message.

Senior managers and boards of directors, especially, invariably have a limited amount of time in which to view the presentation, so brevity is essential. The need for decisions and the next steps should always be summarised on the final presentation slide.

CONSULTATION

Consultation is simply another form of communication, but one in which specific questions are asked and answers given – as opposed to information provided on a routine basis or when requested.

Consultation begins right at the start of the information risk management programme. The programme manager must consult at all levels both within and outside the organisation in order to understand:

- The goals, scope and objectives of the programme.
- The key roles and responsibilities needed to undertake the programme.
- How governance of the programme will be managed.

- The internal and external contexts within which the organisation operates.
- The organisation's risk appetite for various types and classes of risk.
- What risk treatment criteria the organisation wishes to set for risk avoidance or termination, risk transfer or sharing, risk reduction or modification and risk acceptance or tolerance.
- The threshold for costs of risk treatment, above which a business case will be required.

That done, the programme manager can begin to:

- Plan the programme.
- Make contact with the most appropriate people or groups within and outside the organisation.
- Identify the organisation's information assets and agree their owners.
- Obtain impact or consequence information from them.
- Obtain threat and vulnerability information.

Much, if not all, of this information will originate from individual directors, managers and specialist staff within the organisation, although some may have to be obtained from external sources such as suppliers, consultants and security-related organisations.

The consultation process itself will require similar skills to those described in [Chapter 4](#) – Risk Identification.

RISK REVIEWS AND MONITORING

If we think back to the PDCA cycle we discussed in [Chapter 2](#), we can see that the Check part of the cycle covers the requirement for monitor and review. This is an important part of the information risk management programme, since it will enable us to verify that treatment has been carried out correctly and that it has been successful. It enables us to:

- Verify that the controls we have put in place have been effective, both in cost and functionality.
- Identify those areas where controls have been less successful, and to take remedial action to resolve these.
- Consider new information assets that have been added to the organisation's portfolio.
- Identify and assess new or changed threats or hazards.
- Identify areas where the vulnerabilities have changed or where new vulnerabilities have been identified.
- Identify areas where the likelihood of risks occurring have increased or decreased.

- Assess and evaluate new or altered risks, and to take them through the risk management process.
- Revise impact assessments for information assets whose value to the organisation has changed due to changing business circumstances.
- Take account of information security incidents that have resulted in harm to the organisation's information assets.
- Consider the use of other methodologies as a means of conducting the information risk management programme.
- Be mindful that some information assets exhibit considerable interdependence, and therefore that if one risk changes, other dependent risks will probably also change.

Risk reviews

There may be some temptation to assume that once a risk has been treated, its treatment has been completely successful and that we can forget about it for the time being since no further treatment would appear to be required. This is a very bad idea, since at this stage we do not know for certain either that the risk has truly been treated successfully or that the treatment has not introduced additional risks.

The first stage of the risk review commences with a revisit of the original risk assessment process to verify that either the impact or the likelihood has been reduced and to calculate a revised risk level. If this calculation indicates that the hoped-for level of risk has actually been attained and that any residual risk is within the organisation's risk appetite, then the treatment can be considered to have been completely successful and recorded as such, together with a date for a follow-up review.

If, however, the level of risk attained has not achieved this, but has been reduced below that of the organisation's risk appetite, treatment can be considered to have been partially successful, and a decision can be made as to whether or not to introduce additional controls in an attempt to reduce the level of risk still further.

If it transpires that the resulting level of risk remains above that of the organisation's risk appetite, then treatment will have been unsuccessful, and a complete reassessment of possible controls will be required. All of this must then be recorded in the risk register to ensure that there is a sound audit trail.

The second stage of the risk review involves an examination of whether or not the original risk treatment introduced any additional risks, or whether it has had an impact on other risks that we have been treating and has either raised or reduced the level of risk in them.

If the overall number of risks treated as part of the information risk management programme is large, there may be too many to examine in one review and it might be more convenient either to review a proportion of the risks in several stages or to review only those risks that were identified as being high or very high, so that the most critical risks are reviewed first.

Risk reviews are best conducted as a group or team activity, since, although an individual manager may have all the requisite skills to carry out the work, a wider perspective can result in more accurate reviews and assessments.

Risk monitoring

Risk monitoring takes a slightly different approach from risk review. Whereas risk review occurs immediately or soon after risk treatment, risk monitoring is an ongoing process that verifies the risk status at regular intervals.

It is also important to monitor the information risk management programme and activities themselves in order to ensure that other factors that have changed are taken into account, such as the organisation's:

- Business strategy and direction.
- Levels of risk tolerance and risk appetite.
- Criteria for impact, risk evaluation and risk acceptance.
- Resources available to engage in the information risk management programme.
- Approach to the controls it chooses when assessing treating risks.
- Internal or external context, including the competitive context, where this is appropriate, as well as the legal and regulatory context.

The periodic basis for risk monitoring should be based on two sets of criteria:

- Firstly, the higher the level of residual risk, the more frequent the monitoring should be – clearly, higher level risks should be monitored more frequently than lower level ones.
- Secondly, for those risks whose threat level does not change quickly – for example, legal and regulatory changes – the risk monitoring may be less frequent, whereas for those risks whose threat level changes more frequently – for example, viruses and software vulnerabilities – the monitoring should be much more frequent.

There will often be a temptation to only monitor the risks annually, but, unless they exhibit either a very slow-changing threat level or the residual risk is very low, this should be avoided, since many attributes of the organisation and its internal and external contexts may have changed during 12 months and some old, altered or new risks could easily be overlooked. As with risk reviews, all risk monitoring output should be recorded in the risk register.

SUMMARY

In this chapter, we have shown how there is an ongoing requirement to communicate and consult with stakeholders at all levels, both within and outside the organisation, and how all risks should be regularly reviewed and monitored to ensure that they remain at a level deemed by the organisation to be within its risk appetite.

10 THE NCSC CERTIFIED PROFESSIONAL SCHEME

Anyone who wishes to work in the information security environment, either within a government organisation or as a contractor to one, must be accredited to do so, regardless of any additional security clearances that may be required.

The certification, known as the Certified Cyber Professional (CCP) scheme, is not merely a qualification, but a full certification awarded to individuals who are able to demonstrate an application of their skills, knowledge and expertise to one of three approved certification bodies. At the time of writing, these are:

- the APM Group;¹
- BCS, the Chartered Institute for IT;²
- the NCSC website.³

The Chartered Institute of Information Security (CII Sec), CREST and RHUL consortium⁴ contains comprehensive details of the scheme, which encompasses six different role areas:

- Accreditor. Accreditors provide impartial assessment of risks to which an information system may be exposed, and accredit such systems on behalf of the organisation's senior management.
- Communications security officer/Crypto custodian. This role manages cryptographic systems and includes Payment Card Industry Data Security Standard (PCI/DSS) compliance.
- Cyber security/information assurance (IA) architect. This role is designed to develop or review system architectures so that they fit the business requirements for security, mitigate risks, conform to security policies and balance information risk against the cost of countermeasures.
- Cyber security/IA auditor. Auditors assess compliance with security objectives, policies, standards and processes.
- IT security officer/information security system manager/information security system officer. This role is to provide governance, management and control of IT security.

¹ See <https://apm.org.uk/>.

² See <https://bcs.org>.

³ See <https://ncsc.gov.uk/information/about-certified-professional-scheme>.

⁴ See <https://www.ciisec.org/>.

- Security and information risk advisor. This role is to provide business-driven advice on security and information risk that is consistent with cyber security policy, standards and guidance.

Note that at the time of writing, the titles and descriptions in the NCSC document are as above, but the titles on the web page are slightly different.

Each role may be assessed at one of three levels, except for the penetration tester, which has four levels and includes a Principle level. Otherwise, they are:

- **Practitioner** – this is the entry level to Certified Professional and is suitable for individuals who work on routine IA tasks under supervision.
- **Senior practitioner** – this level is suitable for individuals who work independently on complex projects and who normally lead a team of IA professionals or lead or oversee the work of other IA professionals.
- **Lead practitioner** – this level is suitable for highly experienced individuals working at senior levels in an organisation, who provide advice and/or leadership on complex strategic IA issues.

It does not follow that professionals who are very experienced at senior level will meet the role description for lead practitioner.

Two different areas of knowledge and expertise are assessed:

- The CIISec (formerly the Institute of Information Security Professionals (IISP) skills matrix for information security technical skills, more information on which is provided in a later section of this chapter and is also available to CIISec members at <https://www.ciisec.org>.
- The Skills Framework for the Information Age (SFIA) levels of responsibility for autonomy, influence, complexity and business skills. Many of the practitioner, senior practitioner and lead practitioner roles align with SFIA levels 2, 4 and 6. More detail is provided in the next section, and is also available from <https://www.sfia-online.org>.

Each of the three levels for each of the six roles is described in greater detail in *NCSC Certification for Cyber Security/IA Professionals*, currently Issue 5.4, dated November 2018, and available from https://www.ncsc.gov.uk/files/CCP-Certification_for_Cyber_Security_IA_Professionals_5-4.pdf.

In the main body of the document, each role has a brief role purpose description and a statement of responsibilities, followed by a 'headline' statement, which outlines the key responsibilities of the role for each of the three practitioner levels. This is followed by a table listing the indicative IISP skill levels for the three role levels. Annex A of the document contains a very detailed description of the CIISec skill definitions.

Each of the skill areas begins with a statement of the CIISec Principle, followed by the knowledge requirements for that skill area.

Following on from this, each skill subset (for example, Governance under Information Security Management) is described in terms of CIISec example skills and, as a NCSC supplementation, the attainment expected for each of the four skill levels.

The three certification organisations operate slightly different schemes with different costs, and prospective candidates are recommended to view the various options on the main CCP page at <https://ncsc.gov.uk/information/about-certified-professional-scheme>.

In addition to completing an application form that itemises and describes their SFIA and CIISec skill sets, candidates are strongly advised to have gained a thorough understanding of the following documents, the first of which is summarised in [Chapter 11](#):

- HMG Security Policy Framework, which can be downloaded from <https://www.gov.uk/government/publications/security-policy-framework>.
- The Cyber Essentials Scheme, details of which can be found at <https://www.ncsc.gov.uk/cyberessentials/resources>.

SFIA

Established in 2003, SFIA was designed as a system aimed at IT professionals to match their skills against business requirements. The current SFIA Framework (SFIA 7, dated May 2018)⁵ is available at <https://www.sfia-online.org/en>.

The scheme consists of five business-related areas:

- autonomy;
- influence;
- complexity;
- knowledge;
- business skills.

These are assessed at seven levels of responsibility:

- Level 1 Follow;
- Level 2 Assist;
- Level 3 Apply;
- Level 4 Enable;
- Level 5 Ensure and advise;
- Level 6 Initiate and influence;
- Level 7 Set strategy, inspire and mobilise.

⁵ A beta release of SFIA version 8 was due to be published at the end of June 2021, with the full release in September 2021.

The levels used in the CCP scheme are mostly 2, 4 and 6, but in some cases level 3 is used.

The 102 skills are subdivided into the following areas:

- information strategy;
- advice and guidance;
- business strategy and planning;
- technical strategy and planning;
- enterprise IT governance;
- strategic planning;
- information governance;
- information systems coordination;
- information security;
- information assurance;
- analytics;
- data visualisation;
- information content publishing.

THE CIISEC SKILLS FRAMEWORK

The CIISec developed its information security skills framework as a means of assessing prospective members prior to interview. At the time of writing, it is currently at version 2.4, dated November 2019. The skills are rated at six levels:

1. Basic knowledge.
2. Knowledge and understanding.
3. Junior practitioner.
4. Practitioner.
5. Senior practitioner.
6. Principal/lead practitioner.

The skills themselves are in 10 distinct groups, listed below.

A Information security governance and management

There are seven subcategories:

- A1 Governance.
- A2 Policy and standards.

- A3 Information security strategy.
- A4 Innovation and business improvement.
- A5 Behavioural change.
- A6 Legal and regulatory environment and compliance.
- A7 Third party management.

B Threat assessment and information risk management

There are three subcategories:

- B1 Threat intelligence, assessment and threat modelling.
- B2 Risk assessment.
- B3 Information risk management.

C Implementing secure systems

There are three subcategories:

- C1 Enterprise security architecture.
- C2 Technical security architecture.
- C3 Secure development.

D Assurance, audit, compliance and testing

There are four subcategories:

- D1 Internal and statutory audit.
- D2 Compliance monitoring and controls testing.
- D3 Security evaluation and functionality testing.
- D4 Penetration testing and conducting simulated attack exercises.

E Operational security management

There are two subcategories:

- E1 Secure operations management.
- E2 Secure operations and service delivery.

F Incident management, investigation and digital forensics

There are three subcategories:

- F1 Intrusion detection and analysis.
- F2 Incident management, incident investigation and response.
- F3 Forensics.

G Audit, assurance and review

There are three subcategories:

- G1 Data protection.
- G2 Privacy.
- G3 Identity and access management (IAM/IdM).

H Business continuity management

There are three subcategories:

- H1 Business continuity and disaster recovery planning.
- H2 Business continuity and disaster recovery management.
- H3 Cyber resilience.

I Information security research

There are two subcategories:

- I1 Research.
- I2 Applied research.

J Management, leadership, business and communications

There are three subcategories:

- J1 Management, leadership and influence.
- J2 Business skills.
- J3 Communication and knowledge sharing.

K Contributions to the information security profession and professional development

There are three subcategories:

- K1 Contributions to the community.
- K2 Contributions to the information systems profession.
- K3 Professional development.

SUMMARY

This completes the chapter on SFIA and the CII Sec Skills Framework, which is generally used by non-governmental organisations. The next chapter deals with those documents and standards that are used more within UK government departments.

11 HMG SECURITY-RELATED DOCUMENTS

In this chapter, we will examine a number of key UK government documents that are the most relevant to information security.

HMG SECURITY POLICY FRAMEWORK

The most recent HMG Security Policy Framework at the time of writing is version 1.1, dated May 2018. It is an unclassified, publicly available document that can be downloaded from <https://www.gov.uk/government/publications/security-policy-framework> and includes government security classifications. It outlines the following areas:

- overarching principles;
- security outcomes;
- good governance;
- culture and awareness;
- risk management;
- information;
- technology and services;
- personnel security;
- physical security;
- preparing for and responding to security incidents.

It then provides additional detail on the following three areas of policy priority:

- information security;
- physical security;
- personal security and national security vetting.

THE NATIONAL SECURITY STRATEGY

This document, dating from October 2010, can be downloaded from <https://www.gov.uk/government/publications/the-national-security-strategy-a-strong-britain-in-an-age-of-uncertainty>. It identifies four distinct areas:

- the strategic context;
- Britain's distinctive role;
- risks to our security;
- our response.

CONTEST, THE UNITED KINGDOM'S STRATEGY FOR COUNTERING TERRORISM

This document, dated June 2018, although not directly related to information risk management, does cover some interesting and useful areas of security. It can be downloaded from <https://www.gov.uk/government/publications/counter-terrorism-strategy-contest-2018>.

It follows on from the national security strategy document by expanding upon the strategic context, which includes the threat from terrorism and strategic factors, and the response areas, especially the Prevent, Pursue, Protect and Prepare functions. It continues with a section on implementation, and finishes with a description of roles and responsibilities.

THE MINIMUM CYBER SECURITY STANDARD

This short document, dated June 2018, is downloadable from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/719067/25062018_Minimum_Cyber_Security_Standard_gov.uk__3_.pdf. It is aimed specifically at government departments, and has five distinct sections covering Identify, Detect, Protect, Respond and Recover.

THE UK CYBER SECURITY STRATEGY 2016–2021

This document can be downloaded from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/567242/national_cyber_security_strategy_2016.pdf. Its principal sections cover:

- The strategic context, including threats and vulnerabilities.
- Our national response, including roles and responsibilities.
- Defend, including active cyber defence.
- Deter, including reducing cyber-crime and terrorism.
- Develop, including cybersecurity skills.

UK GOVERNMENT SECURITY CLASSIFICATIONS

The current definitive document, version 1.1 dated May 2018, describes government security classifications and can be downloaded from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/715778/May-2018_Government-Security-Classifications-2.pdf. The publication consists of three tiers:

- **OFFICIAL**, which includes the majority of information created or processed by the public sector.
- **SECRET**, which includes very sensitive information, justifying heightened protective measures to defend against threats.
- **TOP SECRET**, which includes the government's most sensitive information, requiring the highest levels of protection from threats.

The publication sets out four principles:

1. That all information collected, stored, processed, generated or shared in order to deliver government services requires an appropriate level of protection.
2. That all government employees, contractors and service providers have a duty of confidentiality and a responsibility to safeguard any government information or data to which they have access, irrespective of whether or not it carries protective marking, and that they must be appropriately trained to do so.
3. That access to all sensitive information should only be granted on the basis of a genuine need to know, and an appropriate personnel security control.
4. That all information assets exchanged with external partners must be adequately protected in line with any relevant legislative or regulatory requirements.

The three individual classifications are then defined in greater detail, including baseline security outcomes and protective marking requirements.

The document continues by defining the special handling instructions that might be required, such as descriptors, which further specify the sensitivity of information; code words, which are used to identify specific assets or events; and prefixes and national caveats, which might be used to restrict visibility of information.

Next, the document lists and briefly describes the legal framework that encompasses the classification system, and references four specific Acts of Parliament:

- the Official Secrets Act 1989;
- the Data Protection Act 2018;
- the Freedom of Information Act 2000;
- the Public Records Act 1967.

It then goes into considerable detail on: threat modelling and security outcomes; working with HMG assets, including comprehensive guidelines for the handling, storage, transfer

and disposal of information assets; protecting assets and infrastructure, including security principles; and a summary of technical controls.

SUMMARY

This completes the chapter on UK government documents and standards. The following appendices will discuss taxonomies, typical impacts, threats, hazards, vulnerabilities and controls, and will then examine a number of information risk management methodologies, provide you with some useful templates, take an overview of the HMG cybersecurity guidelines, and finally, provide references and further reading on the subject.

APPENDIX A – TAXONOMIES AND DESCRIPTIONS

Taxonomies are simply ways of ordering or classifying information and can help us to understand concepts through either diagrams or written explanations. For clarity, this appendix includes both forms for the following areas:


- 1. an overall taxonomy of information risk;
- 2. typical impacts or consequences.

It should be noted that these are simply the author’s interpretation, and are not necessarily complete in terms of all possibilities, or to the deepest level of abstraction.

INFORMATION RISK

Information risk is the combination of the impact or consequence of a threat or hazard on an information asset and the likelihood or probability of its doing so. [Figure A.1](#) illustrates the key components.

Figure A.1 An overall taxonomy of information risk



Impact or consequence

The impact or consequence of an event is the successful result of one or more threats acting upon one or more vulnerabilities of an information asset. They are categorised as follows:

- Primary impacts are those that result from the event itself, when a business function is detrimentally affected or unable to continue.
 - Direct primary impacts; for example, if a customer database is hacked and personal information is stolen, the organisation will lose control of that valuable information resource.
 - Indirect primary impacts; indirect impacts are those that may occur as a consequence of the direct impact. In the above example, the Information Commissioner may levy a fine against the organisation for failing to adequately protect the information.
- Secondary impacts are those that result from responding to or recovering from the event.
 - Direct secondary impacts include such things as customers purchasing their products or services from another supplier.
 - Indirect secondary impacts include such things as fines imposed for failing to file statutory returns on time because the information is unavailable.

The section Typical Impacts or Consequences lists and describes the various types in more detail.

Likelihood or probability

Likelihood expresses the possibility that an event may occur, but places no certainty on it doing so. Probability, on the other hand, expresses a greater degree of certainty, in that it is based on mathematical or statistical information computed by or gathered by research. The two terms are sometimes used interchangeably, but it should be remembered that likelihood is a qualitative view, whereas probability is a quantitative view.

Threats or hazards

Some threats are malevolent in origin, such as hacking and social engineering, while others are non-malevolent, such as environmental threats and simple failures.

The likelihood will be influenced by:

- The value of the information asset to an attacker – in cases where an attack is from a malevolent source, the value of the information asset to the attacker will contribute to the lengths to which the attacker will go in order to carry out the threat.
- The complexity of the attack – some threats can easily be carried out, especially where automated tools are available to the attacker. Others will require

considerably more technical expertise, and are unlikely to be undertaken by so-called 'script kiddies'.

- The risk of detection – in the case of some types of attack, the attacker must spend significant amounts of time in carrying out the attack. The risk of detection increases with time, and those attacks that can be carried out more quickly may be adopted more frequently.
- Previous attacks – those threats or hazards that are known to have already occurred will have a considerable bearing on the likelihood of an attack succeeding, since details of them may be available to an attacker.

Appendix B lists and describes various types of threats and hazards.

Vulnerabilities are weaknesses in or surrounding the information asset, which a threat might exploit in order to compromise the information asset. Vulnerabilities may be: physical, such as inefficient locks; technical, such as poorly configured firewall rules; or procedural, such as a lack of segregation of duties. They have the following contributing factors:

- Existing controls – the presence and strength of existing controls will determine the probability that an information asset will be able to resist an attack.
- The motivation of the attacker – this might be either revenge or financial gain for the attacker, or conversely it might be financial disadvantage for the information asset owner.
- The technical skills required to successfully attack on an information asset.
- The tools required in order to carry out a successful attack, many of which are freely available on the internet.
- Success rate – the previous success rate of a type of attack, if this is known.

Appendix C lists and describes various types of vulnerabilities.

TYPICAL IMPACTS OR CONSEQUENCES


Figure A.2 illustrates and describes some of the possible impacts or consequences that might arise as a result of a successful threat against an information asset.

Operational impacts

Operational impacts are generally experienced rapidly by the organisation. Most are very obvious – for example, when information to which they expect to have access is no longer available or which they can plainly see has been dramatically altered.

Very often, direct operational impacts will result in subsequent indirect financial impacts, so an inability to meet a service contract may well result in lost orders or in claims for contractual damages. Operational impacts include:

Figure A.2 Typical impacts or consequences



- loss of data (confidentiality, integrity and availability);
- loss of premises and equipment;
- order backlogs;
- productivity losses;
- industrial action;
- reduced competitive capability;
- inability to meet service contracts;
- inability to progress new business or developments;
- damage to third-party relations;
- impaired management control;
- loss of customers to competitors (churn).

Financial impacts

Unsurprisingly, financial impacts or consequences are normally those that gain the greatest attention within the organisation. It is frequently against a backdrop of possible financial loss that the costs of remedial actions will be compared. While this is certainly correct, it is also true for other types of impact as well.

Many of these impacts – for example, lost sales immediately following the event – will be felt very quickly, while others – for example, increased insurance premiums – may not manifest themselves until a later date, possibly some considerable time after the costs of the event have been counted.

Financial impacts may also not be as noticeable to the whole organisation – for example, staff may not be aware of the financial implications of an event at all, and have no appreciation of the position in which the organisation finds itself until they read about it in the media or find that pay increases and bonuses are reduced. Financial impacts include:

- loss of current and future business opportunities;
- increased cost of borrowing;
- cancellation of contracts;
- contractual penalties;
- loss of cash flow;
- replacement and redevelopment costs;
- loss of share price;
- increased insurance premiums;
- loss of tangible assets.

Legal and regulatory impacts

As with reputational impacts, legal and regulatory impacts can have serious repercussions on an organisation, and the handling of these is best dealt with by a specialist team within the organisation, who may communicate information regarding an event through the corporate communication department. Legal and regulatory impacts include:

- warnings or penalties from the sector regulator;
- fines for late submission of company accounts;
- fines for late payment of taxes;
- breach of contract damages;
- fraud and other criminal acts.

Reputational impacts

Reputational impacts are almost always highly detrimental to the organisation. For this reason, many organisations employ communication specialists who are skilled in countering negative publicity and putting a positive spin on any bad news. In such organisations, most staff are advised not to talk directly to the media, but to pass enquiries through to the corporate communication department. Reputational impacts include:

- stock market confidence;
- competitors taking advantage;
- customer perception;
- public perception;
- industry and institutional image;
- confidential information made public.

Wellbeing of staff and the public-at-large

Although more rare, safety incidents are generally highly visible outside the organisation, and occasionally have an effect on the public-at-large. More common, however, are any events that may have an adverse effect on the organisation's staff, and these can also cascade into financial and operational secondary impacts. Wellbeing impacts include:

- pandemics;
- safety, health risks and injuries;
- stress and trauma;
- low morale of staff.

APPENDIX B – TYPICAL THREATS AND HAZARDS

Threats and hazards cause impacts or consequences to occur on one or more assets by taking advantage of one or more vulnerabilities. The list in [Figure B.1](#) may not be exhaustive, but should provide a starting point.

MALICIOUS INTRUSION (HACKING)

Hacking is a generic term applied to many forms of unpleasant behaviour, although it began as a description of what people did in order to find out how computers worked and how to improve their performance. Hacking almost invariably results in a breach of confidentiality, integrity or availability as hackers use software tools to intercept and decrypt legitimate information, and either steal it, change it or deny access to it in some way.

Since the introduction of the Computer Misuse Act in 1990, hacking has been treated as a crime, since it invariably involves accessing a computer without the owner's permission to do so.


Denial of service

Occasionally, hacking is used to deliver so-called DoS attacks, designed to prevent legitimate access to systems, often to make a political point or as revenge for a perceived or real injustice. Associated with hackers are so-called 'hacktivists' who sometimes deface website pages or mount DoS attacks. Their attacks follow no set pattern or target, but the end result becomes an availability issue not only for the organisation hosting the targeted website, but also potentially for its users.

Unauthorised access

Most people think of hackers as working from outside the organisation, and although this may be generally correct, there is no reason why someone within the organisation should not be considered a hacker, since their ability to access information from within is much easier than for someone on the outside. It is not unknown for criminal gangs to 'plant' well-trained technical staff and security specialists inside target organisations in order to steal information or systems hardware, and such people might be sufficiently well resourced to evade detection during the recruitment process.

Figure B.1 Typical threats and hazards



As with hacking, following the introduction of the Computer Misuse Act in 1990, the unauthorised access of someone else's computer became a criminal offence, since it invariably involves accessing a computer without the owner's permission to do so, and unauthorised use from within an organisation can be considered an offence under the Act just as it can from an external attack.

Unauthorised access to systems and information is not restricted to criminal hackers, and although most prosecutions have been for activities that one would consider unsociable, even people who are supposed to be beyond reproach have been prosecuted for offences under the Act – some years ago an ex-police superintendent was convicted of using the Police National Computer System to acquire information about his ex-wife's new partner.

Unauthorised network scanning

Likewise, once hackers penetrate unsecured wireless networks, they will invariably carry out a scan of systems, and both wired and wireless networks, as part of obtaining information that will aid them in attaining their targets.

Once connected to an organisation's internal network, hackers will invariably try to steal user credentials such as user IDs and passwords and, if they can also steal cryptographic keys, they may be able to decrypt information they recover from systems.

Interception

Hackers should not be viewed as unsophisticated individuals. Very often, they are extremely experienced, equipped with highly sophisticated tools and often financially funded by criminal enterprises, allowing them to intercept wireless and wired network traffic, especially in organisations that promote the wider use of Wi-Fi within their premises. Apart from intercepting wireless traffic, hackers may also introduce their own wireless access points in an attempt to lure unsuspecting users onto insecure networks.

Session hijacking

When a user establishes an internet connection to another computer, the remote computer invariably places a cookie on the user's computer, which permits continuity of the session without the need for the user to re-authenticate at every new request or transaction. This is sometimes known as a 'magic cookie', and if an attacker can obtain this, he can then hijack the user's session and masquerade as the legitimate user.

Session hijacking is quite a complex process and requires an attacker with considerable skill and motivation.

Website modification

Website modification is frequently an attack carried out as revenge for a perceived injustice or to demonstrate an attacker's technical expertise. Weaknesses in the website code allow skilled attackers to take control either of the website page or even the entire

website in order to substitute information of their own choosing. Individuals and groups of hackers have carried out this type of attack on numerous occasions, sometimes to the great embarrassment of governments and businesses alike.

Software modification

Attackers who wish to carry out attacks on a large number of prospective targets will often modify existing software by adding their own malicious code – malware – and offer the software through spam emails or other means. The malware might cause a number of different results – for example, to turn the target user's computer into a botnet client or to steal personal information.

Data modification

Attackers who access websites and visibly change the web pages are just one side of the coin. Instead of changing something obvious, such as the website owner's identity, they might modify other data on the website, which might not be so obvious and might provide a visitor with incorrect or misleading information, often inserting links to other websites containing malware or where the user's security credentials can be obtained.

Decryption

The unauthorised decryption of encrypted information is a threat that is used in combination with, for example, the theft of business information and the theft of cryptographic keys or the brute-force attacking of encrypted information. An example of this is in cases where an attacker will intercept and record encrypted wireless network traffic over a period of time, and either decrypt it having acquired or guessed the correct key, or will subject it to a lengthy attack in order to recover the key, after which further traffic could be decrypted 'on the fly'.

Credential theft

The theft of user credentials permits an attacker to carry out spoofing, masquerading or impersonation attacks against targets. User credentials may be as simple as user ID and password combinations gathered from social engineering or phishing activities. In many cases, attackers will not use the credentials themselves, but will sell them on in bulk to criminal gangs.

ENVIRONMENTAL THREATS

Environmental threats are almost always concerned with availability, since they affect the environment in which a system or information resides. Those impacts or consequences that occur as a result of natural events – for example, severe weather – are often referred to as hazards in order to distinguish their motivation from those of malicious threats.

Natural hazards

Severe weather can knock out power and communications networks, denying us electronic access to our information, and can also have an adverse effect on communications networks that remain operational, since many people stay at home during adverse weather, which increases the load on the internet and delays access to information.

Natural hazards can also include pandemics, in which a large number of key staff may be taken ill and unable to work, resulting in an impact to service delivery, for example. In the case of the coronavirus pandemic in 2020, large numbers of people were encouraged to work from home where possible, and others who were unable to work because their workplaces were closed also stayed at home. This greatly increased the 'normal' load on broadband and mobile networks, and as a consequence overloaded many organisations' websites.

Accidental and malicious physical damage

While physical damage might not at first appear to be a threat to information, instances of accidental damage to underground communications cables is quite commonplace, and in recent years thieves have targeted copper cables for their scrap value. Either instance may well result in an enforced denial of access unless the communications service providers are able to install a fully redundant infrastructure.

A typical example of this was in April 2009, when a tunnel-boring machine managed to directly drill into a tunnel near Ilford that carries communications cables. The cables were wrapped around the machine and pulled out in both directions.

It is also commonplace for thieves to mistakenly target fibre-optic cables instead of copper cables, the impact of which can be even more severe, since fibre-optic cables tend to carry significantly greater quantities of traffic.

Fire

Fire is another serious hazard. Although much information can be backed-up in alternative locations, some cannot. Take, for example, the fire in May 2014 at the Glasgow School of Art, in which many students' entire final year work was destroyed. Either the smoke or the water used in extinguishing the fire may have ruined even work that was not actually destroyed by the flames. Ironically, during the restoration work following the fire, another fire broke out in June 2018, almost totally destroying the Mackintosh building. Art is a form of information that is often overlooked, and in many cases is very fragile and irreplaceable.

Communications jamming or deliberate interference

Communications jamming or deliberate interference is a less common form of threat to information, but can be highly successful if deployed in a very specific location, and it is an unfortunate fact of life that much of this cannot legally be prevented, since wireless networks operated in unlicensed bands of the radio spectrum.

Many of these hazards affect a wide geographic area, and can cause serious disruption to multiple organisations rather than to a specific organisation or system.

Communications failures

Failures in communications networks can have considerable adverse consequences in the area of availability. Their causes are varied, ranging from accidental damage by utility companies digging up underground cables, and cable theft, especially for the copper content in many cables, to utility failures in and between communications centres, and severe weather, which can quickly disrupt both wired and wireless traffic.

Power failures

Fortunately in the UK generally, wide-area power failures are relatively uncommon, since the national grid is extremely robust. However, localised power failures are more frequent due to the nature of network design, since the level of resilience decreases at the lower levels of the network.

In some countries, however, power failures over wide areas are commonplace and regularly restrict the availability of access to information sources.

ERRORS AND FAILURES

Errors and failures fall neatly into two categories – those made by users and technical staff, and those things that simply fail. Neither form is regarded as being malevolent, even though some user and technical errors are caused by lack of attention or poor training. Despite the view of many technicians that both hardware and software is designed to cause them grief, there is no evidence to suggest that this is actually the case. Examples of error threats are discussed below.

Software failures

Software failures are relatively uncommon, since most software is tested exhaustively before it is distributed and manufacturers usually fix bugs very quickly, as customer satisfaction is normally high on their list of priorities. Software failures are more common when two applications, normally from different manufacturers and which usually work well together, cease to function in the expected manner, usually when one or the other has been updated without robust interdependency testing.

Software interdependencies

As organisations' information infrastructures become increasingly complex, the interdependencies between software applications grow exponentially, and it sometimes only takes a minor change in one to cause a serious downstream problem in one or more others, hence the need for full regression testing becomes of prime importance.

In 2012, a UK mobile telecoms provider began the rollout of a new customer database and, despite their best efforts to verify the software before rollout, the software failed under load and almost half of their 22 million customers were without service for several hours. There have also been numerous examples in recent years of high street banks updating their software, only to find that failures have locked customers out of their online accounts and prevented scheduled transactions from taking place.

System overloads

System overloads are far more frequent, and it is often the case that websites are unable to cope with sudden increases in demand for sales or services. In May 2014, the Ticketmaster website selling tickets for the 2014 Commonwealth Games in Glasgow put an additional 100,000 tickets on sale, and suffered a major overload problem resulting in people unable to access the site and people who thought they had obtained tickets but had not.

During the 2020 coronavirus pandemic, people were encouraged to use online shopping rather than visit supermarkets. This resulted in the main food retail organisations experiencing major overloads on their websites.

Hardware failures

Hardware failures are much more common than either of the previous examples. Although systems hardware has become increasingly reliable in recent years, failures do still occur and, unless the organisation has invested in resilient systems, duplication of key elements or DR, customers can find themselves unable to access the services they require, resulting in delays, frustration and ultimately bad press for the organisation concerned.

User errors

Errors made by users form a significant threat to information, since it is very easy using commonly available office applications to change files or filenames, either deliberately or accidentally, to delete files or to change file contents, all with no means of control, and users frequently do not realise their mistakes until long after the event.

Although backup and restore processes should address this, sometimes the backup media has been damaged or overwritten, and also the user may be reluctant to admit their mistake and request the restore of a file.

This type of failure especially affects shared storage resources, including both internal shared drives and external cloud services, where any one of the users sharing the storage facility can inadvertently or deliberately delete or modify files.

Technical staff errors

In theory, technical staff errors should be less frequent than user errors, since technical staff should have received appropriate training for their role. However, mistakes can and do occur, and it is an unfortunate fact that technical staff may well have greater system

access privileges, and can do more damage with a single keystroke or mouse click than an ordinary user can.

Internal and external software errors

Generally, all software applications that are used by both large and small organisations have been thoroughly tested for functionality prior to their release, so failures are relatively uncommon but still can and do happen. More prevalent (and potentially more dangerous) in computational software, errors can at best lead to poor decision-making, and at worst could have life-threatening consequences if, for example, an application recommended the dosage of a drug to be significantly greater than would be safe.

Change failures

Change failures are often identical to technical staff errors, in that changes are incorrectly made to systems, applications and information. However, the other option here is that changes might have been incorrectly specified, with the possible result that the change either simply did not work as expected, or that the change made matters worse.

This particular type of threat reinforces the need for thorough testing following a change and also, later in the information risk management process, to examine the results of risk treatment to verify that it has been effective and that it has not introduced additional problems.

SOCIAL ENGINEERING

Social engineering is a technique used by hackers and other ne'er-do-wells to acquire information, generally about access to systems (both electronic and physical) so that their hacking activities are simplified. Social engineering comes in several forms – not only the traditional approach where a hacker attempts to engage with a user by conversation (usually over the telephone or by email), but also by disguising malware as legitimate software and web links and by copying the style, naming conventions and language of a target organisation. For example, they may send a user an email that appears to originate from their bank, but in which embedded web links take the user to the hacker's own website. Examples of social engineering threats are discussed below.

Spoofing, masquerading or impersonation

Spoofing, masquerading or impersonation is a very common approach to social engineering. It is often achieved by electronic means, in which an email that purports to come from an organisation with whom the subject may have had dealings is sent to a number of recipients, but in fact may lead them to carry out an action that captures some of their personal information. Alternatively, it may take them to a website containing (usually false) offers of goods or services, or may infect their computer with malware (see a more detailed description in the Malware section).

Phishing

Phishing is the act of trying to conduct a scam of some kind by contacting people, usually by text message or email, and either offering them something (frequently at a ridiculously low price) or telling them that they owe money in some way. The recipient clicks on a link in the message and is taken to the scammer's website where they are relieved of their money. Phishing can either be used in the context described above, or as 'spear phishing', in which specific individuals are targeted rather than the general population.

Spoofing, masquerading or impersonation by phishing may also take a more personal approach, in which direct contact is made to the victim inviting them to take some action or other, again usually with unfortunate consequences.

Whatever the method used, the main objective is usually to obtain information that can be used to commit crime – to steal money or to order goods and services to be paid for by another person.

Spam

Spam is a technique frequently used to carry out phishing attacks, by enticing the target to click on a malicious web link or to provide information to the attacker. Many internet service providers (ISPs) use sophisticated systems to detect known spam email and to quarantine it in a 'spam folder', which the user can then examine and decide whether or not the message is genuine.

Attackers who make use of spam do not care if 99 per cent of their messages are trapped and deleted in this way, since the remainder will reach their destination, and a percentage of these will still result in a successful attack.

Disclosure

Disclosure of information is generally accidental in nature, and most social engineering is designed to fool the target into disclosing personal or sensitive information without realising that they have done so.

A typical example of this might be where the attacker calls the target on the telephone and purporting to be their bank requests the first and third digits of their security PIN for identification. They then might say there was noise on the line and ask for the second and fourth digits, thus gaining the full PIN.

Most social engineering attacks rely on the fact that target individuals have a fundamental belief that they are either talking to or exchanging information with a genuine organisation.

MISUSE AND ABUSE

Whereas hacking is usually deemed to originate from outside an organisation, misuse normally originates from within. The net result may well be the same for either approach

but, in the case of misuse, the internal user or technician has the added advantage of already being on the right side of the organisation's firewall and security systems, may have access to the required passwords and have suitable access privileges. For this reason, threats from internal attackers potentially present significantly greater levels of likelihood of success than those of external attackers. Examples of misuse threats include the following scenarios.

Modification of system access privileges

Whenever hackers succeed in penetrating a network, one of their first actions will be the modification of system access privileges so that they can explore systems and steal or change information. Naturally, this is much more easily accomplished if the attacker is already within the organisation – actually connected to the network. Systems administrators are generally well placed to undertake this kind of activity, and this frequently takes place as a result of poor control of passwords and the reluctance to change default user IDs and passwords on new systems.

Unauthorised systems activity

Misuse itself can cover a variety of activities, and includes unauthorised system activity, in which users probe systems on an organisation's network to find out what information or software is available to them. It can involve the abuse of other privileges, such as the ability to send email using the organisation's domain name, or, more commonly, abuse the availability of the internet to download material such as pornography, pirated software, music and films. Inappropriate use of the organisation's internet connectivity can also result in the organisation becoming infected by malware.

Software theft and business information theft

Misuse can also cover the sending of proprietary information to persons outside the organisation, including the theft of the organisation's software, the use of the organisation's facilities to run a business and the posting of derogatory or abusive comments on social networking websites.

Finally, a form of misuse that has hit the headlines is that of confidential information or laptop computers containing confidential information being left unattended – in taxis or on trains, for example. In 2008, secret government documents detailing the UK's policies towards fighting global terrorist funding, drugs trafficking and money laundering were discovered on a London-bound train and handed to a national newspaper – and there have also been numerous reports of government security officials having unencrypted laptops stolen from vehicles.

PHYSICAL THREATS

Many physical threats are also carried out by employees – many will have access to systems and equipment that they can easily remove from the organisation's premises without the fear of discovery, whereas an external attacker would have to pass through the organisation's layers of physical security in order to do so.

Unauthorised access

While unauthorised access to an organisation's premises in itself is not really a threat, it is what an intruder could achieve once within the perimeter that should concern us. The problem is compounded in cases where a legitimate visitor is left unattended in an organisation's premises and is able to wander at will within a supposedly secure environment.

Theft of computers and portable devices

Desktop computers are quite difficult to steal, but laptops, tablet computers and smartphones are relatively easy to slip into a briefcase or pocket and be taken from a building before the alarm can be raised. However, even the theft of larger computers and racks of servers has been known, and criminal gangs have stolen significant numbers of high-end systems and networking equipment to sell on the black market. With any of these systems, laptops, tablets or smartphones go the information they hold and quite possibly also details of access permissions to other systems or services.

Theft of authentication devices

The theft of authentication devices is less common, but used in conjunction with a user ID and password or PIN, may allow an attacker easy access to systems, often over the internet, so that no physical presence on the organisation's premises is necessary.

MALWARE

The term 'malware' is used to refer to any form of malicious software that can be used to attack an information system. Examples of malware include software entities that result in the collection of, damage to or removal of information. Such software is almost always concealed from the user, often self-replicating – attaching itself to an executable program – and can spread to other systems when the user unwittingly activates it.

Some malware makes no attempt to conceal its existence, but appears to the user as legitimate software. Its purpose, however, is usually very similar in that it may collect, damage or remove information when the user activates what they believe is a legitimate program. Examples of malware include the methods discussed below.

Viruses

Viruses must always be attached to another piece of software or data – often legitimate – and may be activated by the user opening an email attachment or executing a program to which the virus has been attached. In the early days of computer viruses, most were harmless, but gradually they have been developed to become highly sophisticated and malevolent. Some will encrypt a user's hard drive so that it can no longer be accessed, and the user will subsequently receive a ransom demand for money to unlock the encryption. Other users will simply have their personal information stolen.

Worms

Worms are very similar to viruses in terms of what they are designed to achieve, but do so in a rather different way. Worms do not require other software or data to spread, seeking out other targets over networks to which they are connected – and they can do so extremely quickly. In 2003, the Slammer worm is reputed to have infected 75,000 computers within 10 minutes.

Back doors

Back doors provide a means for an attacker to access the computer and use it for their own purposes without the need to undergo any authentication checks. Back doors can be used to turn the computer into a 'botnet client' that can be used under remote control by an attacker to send out spam email or to launch a distributed denial of service (DDoS) attack.

Trojan horses

Trojan horses are much more successful than any other method of attacking a target. These are often disguised as legitimate software or hidden inside compromised files, which users are lured into downloading and/or running. They successfully avoid security countermeasures because users having accounts with administrator privileges allow the Trojan to run.

Another highly successful means of infection is by the use of compromised websites. Trojans can download themselves without the user needing to click on any links on the web page, and the simple act of visiting an infected web page can be sufficient. An increasing number of criminal organisations are making full use of sophisticated Trojans to attack target systems in order to capture information.

Rootkits

Rootkits are more complex software applications that hijack the user's operating system and make themselves invisible both to the user and to any security software. They frequently still perform all operations that the user has requested, but they can also make duplicate copies of sensitive information such as user IDs, passwords and account details and then transmit them to another computer. Rootkits are often used to enable financial fraud or identity theft.

Spyware

Spyware is a common example of the use of cookies by websites, some of which are designed to be persistent and to track and report the user's web usage back to a third party without the user's knowledge. Some spyware can log the user's keystrokes and search for specific information such as bank account login credentials. Other versions have been known to install software that dials premium rate telephone numbers on those computers that are still connected to modems in order to generate revenue for the attackers. Spyware can also be installed by software that performs a legitimate service such as freeware.

Active content

Active content is the mechanism by which Trojan horses can be downloaded to a computer through its internet browser. Modern web applications use active code to perform complex tasks within the web page to improve the user's experience. There is no doubt that they are extremely good at achieving this, but they are also ideal for installing malware on a target computer. Failure to set the appropriate level of security in the browser will allow the compromised code to be installed and to run itself on the target without the user's knowledge.

An example of this type of attack is where a banner advert runs on a well-known, usually reliable and frequently visited website, in which the Hypertext Markup Language (HTML) code for the banner has been supplied by a third-party advertiser. The attacker adds the Trojan horse software into the banner HTML code and people view the website thinking it trustworthy because of the reputation of the organisation, little realising that the advert is infecting their computer. The payload can be any of the forms of malware described earlier in this section.

Botnet clients

Botnet clients are systems that have been infected with a particular form of malware that allows a botnet controller to make use of the resource for whatever purpose they have in mind. This might be in order to mount a phishing attack to gain information, to send out large quantities of spam with much the same aim in view or to mount a DoS attack on one or more websites.

Ransomware

Ransomware is simply another variety of Trojan that downloads onto a computer and generally encrypts user files or the entire hard drive. The malware then notifies the user that their machine can be unlocked after paying a ransom fee.

A particular example of ransomware was that of the WannaCry attack that took place across 150 countries in 2017. It was designed to exploit a vulnerability in the Windows operating system, having been allegedly originally created by the United States National Security Agency and leaked by a hacking group.

It is reputed to have affected 230,000 computers globally – including a third of hospital trusts in the UK, and costing the National Health Service as much as £92 million. Users were locked out of their systems and the ransom being demanded by payment in Bitcoin.

APPENDIX C – TYPICAL VULNERABILITIES

Vulnerabilities or weaknesses in or surrounding an asset leave it open to attack from a threat or hazard. This appendix lists a number of typical vulnerabilities, but it should be understood that there are many more and that new vulnerabilities, especially in application software, will be discovered on a daily basis. However, this list, based on BS 7799-3: 2017, provides some generic types and is a good starting point for vulnerability analysis. [Figure C.1](#) illustrates these.

ACCESS CONTROL

Access control has two complementary uses: firstly to permit access to resources for authorised persons, and secondly to deny access to those resources to unauthorised persons. Failures in access control are one of the most common reasons for successful attacks against information assets and are very likely to increase the likelihood of successful attacks against them.

Most failures in access control are either as a result of not following processes or procedures, or as a result of not setting up those processes or procedures in the first place.


The lack of, or poorly written, access control policies

A formal access control policy that is inappropriate for the needs of the organisation, the lack of a suitable policy or one that is not properly communicated to staff will inevitably cause severe repercussions. Access to systems, applications and information should only ever be given on the basis of the user's business need, and should always be approved by their line management.

Failure to change access rights of users changing role or when leaving the organisation

Another vulnerability connected with this area is that of poor access control for users changing roles or leaving the organisation. The access to systems, applications and information is frequently overlooked when an individual changes roles. As we shall see later, a method of combatting this is that of role-based authentication, in which the user gains access by means of job function and identity, rather than by their identity alone.

Figure C.1 Typical vulnerabilities



When a user leaves an organisation, their access control permissions should be treated like any other organisation asset, and revoked on leaving or termination.

Inadequate user password management

A frequent vulnerability is that of poor password management, including the failure to enforce regular password changes, which can be carried out automatically together with a test of password strength. For those circumstances in which the need for security is greater, additional and robust methods of authentication can be provided.

However, more recent advice is that longer, stronger passwords are preferable to regular password changes, since users may be tempted to use passwords that are easier to remember and therefore more likely to be guessed.

The continued use of default system accounts and passwords

An extremely common vulnerability is the continued use of default factory-set accounts and passwords for new and upgraded systems. Many individuals in the hacking community are aware of these and circulate them around the community.

The failure to change or hide wireless network identities or service set identifiers will allow an attacker to pinpoint target networks, and if the default administrator passwords have not been changed, or the security level has been left at the basic wireless encryption protection setting, these provide a highly attractive entry point into an organisation's network.

The use of embedded system accounts and passwords

Worse still than the continued use of default settings, there may sometimes be a tendency to allow one system to connect to another by embedding user IDs and passwords within applications. This is a highly dubious practice, since a change on one system or another can easily result in application failures, and if such a password is discovered by an attacker it can lead them directly into more systems.

The lack of security of mobile devices

Many organisations fail to secure mobile devices, whether these are supplied by the organisation or brought in by the users themselves. Mobile devices generally are relatively insecure and easily lost, mislaid or stolen, making both the device and the network to which it can connect equally vulnerable.

The lack of network segregation

Network segregation is commonplace in larger organisations, in which different networks are constructed according to their use, and possibly according to their confidentiality, integrity and availability requirements. For example, an organisation with a significant research capability might well place this on a separate network from finance or general administration use.

Failure to restrict access to networks differentiated according to use is a very common vulnerability, and may allow people to reach resources to which they have no entitlement.

Some organisations make use of a shared network drive to move or exchange large volumes of information between users or departments. This makes for a rich source of unauthorised access by internal staff, and should be avoided.

Failure to impose a clear desk and clear screen policy

The lack of clear desk and clear screen policy is again another very common vulnerability. Some organisations make it a disciplinary offence for an employee to leave confidential materials in plain view or failing to log out of their workstation when they are away from their desk.

The use of untested software

It is good practice for organisations to test new or updated software, including the testing of patches, before it goes into a production or general use environment. Untested software may not only cause operational issues if it fails to work as expected; in cases where it is used in conjunction with other applications, it can have a knock-on effect resulting in an embarrassing chain of consequences.

Failure to restrict the use of system utilities

Finally, although a relatively minor vulnerability, the failure to restrict the use of system utilities – normally by setting access privileges within the user's profile – can result in users carrying out activities that are detrimental to their own device or to other systems, applications or information within the organisation.

POOR PROCEDURES

When acquiring systems hardware and software, developing software and maintaining both, it is vital to ensure that selection and specification is carried out according to a formal set of criteria that include appropriate security features. Unlike access control failures, this type of vulnerability is rarely noticed immediately but can result in serious consequences at a later time.

The root cause of this is often a failure to correctly specify appropriate criteria prior to acquisition or development, and may result either from a lack of forethought or a desire to cut costs.

The lack of clear functional procurement specifications

The procedures for acquisition and procurement of systems and services are often very detailed and precise. However, in some cases the specifications can omit vital security requirements that have been overlooked simply because they appear to the procurement department to be an obvious requirement.

The lack of clear functional development specifications

The most obvious example of this type of vulnerability is that of the lack of or incomplete specifications for developers. When specifications are unclear, or even non-existent, some applications developers will make their own judgement call on what is required.

Many years ago, I was asked to test a system in which the ability to change the system's settings was specified as being subject to 'a complex key sequence'. The developer had no idea what this actually meant, and coded the command exactly as specified. When one typed in the character string 'a complex key sequence', the system opened up access to the settings.

Failure to validate data entry

Many information vulnerabilities are brought about by a failure of applications to test for correctly formatted data, and hackers can exploit this lack of validation to cause failures, which subsequently allow them to inject inappropriate data, take control of applications or steal information.

The use of undocumented software

We have already discussed the lack of or insufficient software testing under access control vulnerabilities, but this is also a systems acquisition, development and maintenance issue, as is that of poorly documented or undocumented software.

The use of unauthorised software

Organisations are also occasionally very lax regarding the uncontrolled downloading and use of unauthorised software, which includes shareware and freeware. If users require additional software in order to carry out their role, this should be available through the correct channels, so that it can be tested and verified as fit for purpose.

Business continuity and disaster recovery planning

The lack of availability of systems and services constitutes a very serious threat to any organisation, and the failure to produce and test robust BC and DR plans is frequently a root cause of this.

BC and DR are closely allied with information risk management, but are complex areas in their own right and therefore beyond the scope of this book. However, a brief description of these has been provided in [Chapter 8 – Risk Reporting and Presentation](#).

PHYSICAL AND ENVIRONMENTAL SECURITY

Physical security is normally highly visible, both to staff and to potential intruders. Very often, the mere presence of robust physical security is sufficient to deter an intruder, but even so it is important that physical security measures are effective, appropriate and well maintained.

Environmental vulnerabilities tend to be rather more difficult to address, but are generally relatively easy to identify and can either relate to the location or construction of premises (for example, in a flood plain) or to the environmental subsystems (power, air conditioning, smoke and leak detection and fire suppression) that underpin major premises such as large office buildings, factories, warehouses and data centres.

Poor management of access to premises and to areas within them

Security consultants will often offer to try to gain entry to an organisation's premises as part of a security audit or penetration test. The careless use of physical access control to buildings, rooms and offices frequently makes this a simple task – tailgating behind legitimate staff is one of the most common forms of unauthorised entry.

Inadequate physical protection for premises, doors and windows

Despite the number of break-ins, the lack of or poor physical protection for buildings, doors and windows (ground floor especially) remains a very significant vulnerability that both intruders and approved security consultants will use to their advantage.

Unprotected storage

Unprotected storage may be a surprising addition to the list, but some organisations simply throw out paper records in the general waste without first shredding them, allowing an attacker to 'dumpster dive' to retrieve information. The media have been very successful in employing this approach to gather news stories, which are invariably to the detriment and the embarrassment of the organisation concerned.

The use of unsuitable environmental systems, including cooling and humidity control

Proper planning for the environmental conditions within data centres of all sizes is not always observed, with the result that computer rooms become seriously overheated, affecting availability.

The inadequate control of humidity and extremes of temperature should be a serious concern to organisations, as should the susceptibility of equipment to voltage and frequency fluctuations and the loss of power. Additionally, leak and fire detection are vital in sensitive locations such as data centres.

The location of premises in areas prone to flooding

Unfortunately, for many organisations the location of their premises in areas susceptible to flooding has become very apparent. Locations that have never before been flooded within living memory have felt the impact of severe weather in late 2019 and the early part of 2020.

The uncontrolled storage of flammable materials

An organisation's vulnerability to damage by fire is another cause for concern. Older buildings built largely of flammable materials and those that contain stores of highly

flammable stock are chief among these. Again, this reinforces the need for suitable fire detection systems.

The location of premises in proximity to hazardous materials or facilities that process them

Finally in this section, the location of premises in close proximity to hazardous material processing or storage facilities must be considered. One only has to look back to the explosions at the Buncefield Oil Terminal near Hemel Hempstead in December 2005 to observe that, although they were located several hundred metres away from the terminal boundary, several buildings were irreparably damaged by the heat from the fires that ensued.

COMMUNICATIONS AND OPERATIONS MANAGEMENT

Along with access control failures, failures of operations management and communications systems rank high among the vulnerabilities that can be successfully exploited, whether deliberately or accidentally. Many of these are due to process failures – again, either through failure to observe them or failure to have them in the first place.

The failure to ensure the appropriate segregation of duties where necessary

The failure to segregate duties where needed can allow attackers to take advantage of access to information that they might not normally have. This ties back into access control, in which access to information might benefit from being role dependent.

Inadequate network monitoring and management, including intrusion detection

Inadequate network management, including the monitoring of hacking and intrusion attacks, will mean that successful attacks and intrusions are overlooked, and little or nothing is known about their occurrence until it is too late and the damage has been done.

The use of unprotected public networks

Many attacks are caused by unprotected public network connections, which allow an intruder to gain easy access to an organisation's network, including the use of shared computers in public environments, such as internet cafés, and the use of unauthorised and possibly unsecured wireless access points.

The uncontrolled use of users' own wireless access points

Occasionally, users of an organisation's networks will discover ways of subverting the organisation's security procedures and will attempt to connect their devices to parts of the network to which they have no entitlement. One way in which this is achieved is by connecting in a 'rogue' wireless access point to which they have unrestricted access.

One of the main issues with this is that the security settings of such wireless access points might not be as strict as those of the organisation itself, and, while the users may be able to access the network, so might an attacker.

Poor protection against malware and failure to keep protection up to date

Malware protection software, especially antivirus software that is not kept up to date, will make an attacker's job much easier. Attackers will take advantage of any means of access available to them, and often are aware of vulnerabilities in applications and operating systems long before a fix is available. Delays in updating these applications leaves an organisation wide open to attack.

The lack of a patching and updating regime

In the same way as the regular updating of malware protection software, the failure to install manufacturers' software patches will leave operating systems and application software open to attack. Naturally, all manufacturers' patches should be tested in a controlled environment prior to general release.

Inadequate and untested backup and restore procedures

Most organisations nowadays carry out regular backups of important information and user data. However, it is far more rare for them to verify that these backups are actually fit for purpose and that information can be successfully recovered from the backup media. This again presents a serious vulnerability, since backup media that does not fulfil its objective is just as bad as having no backup regime at all.

Improper disposal of 'end of life' storage media

Once storage media have reached the end of their useful lives, they should be properly disposed of. There are numerous stories in the press regarding people who have bought second-hand computers only to find that the hard drives still contain sensitive or personal information that had not been securely removed prior to the sale. Some organisations will not allow magnetic media of any kind to be resold, and insist that disposal is irreversible.

The lack of robust BYOD policies

The concept that an organisation's staff can bring their own devices into the organisation's network has become very popular, since it can reduce the hardware costs to an organisation. However, the lack of appropriate policies for its use and the lack of enforcement can bring about serious breaches of security, especially in situations where other members of a user's family have access to the same device.

In 2010, one organisation was badly affected by a virus that was brought in on a user's own personal computer. The machine had been used over a weekend by the user's teenage son, who had unwittingly accessed a website that contained infected software. The resulting infection took the organisation's entire IT department several days to clear up, and the user (a senior manager) was cautioned. Unfortunately, the same thing

happened the following week, and the user was then banned from bringing in his own machine.

Inadequate change management procedures

Inadequate change control can lead to software and patches being rolled out to the user population, new systems, services and network connections being made and redundant systems removed without full consideration (and risk assessment) of the consequences. In smaller networks, change control can easily be vested in one or two people on a part-time basis, but as an organisation's network grows, it may be necessary to employ a full-time team with representatives from multiple business units.

The lack of audit trails, non-repudiation of transactions and email messages

In some sectors, it is vital that online transactions and email correspondence is subject to detailed logging and non-repudiation. In many applications, this audit trail is built in to the application software, and in the event of a dispute regarding 'who did what', or 'who said what', those organisations that are able to produce clear evidence in their favour will have a greater chance of success than those who do not.

The lack of segregation of test and production systems

Those organisations that employ large-scale systems and application testing prior to rollout are open to problems if they fail to separate test and operational facilities, since users may inadvertently connect to a test system, resulting in failed transactions.

The uncontrolled copying of business information

Operational management should limit the uncontrolled copying of information by users who have no need to access it – again, this is also largely an access control issue (shared transfer drives being an example of this), but the identification of such activity may fall into a different management area.

PEOPLE-RELATED SECURITY FAILURES

While some incidents are caused by the failure of systems, software and the supporting infrastructure, the root cause of most incidents is through people who either fail to follow a process correctly (or indeed at all) or who follow a process that is flawed. Much of this is accidental, some remains deliberate, but, whichever the type, most of it is avoidable to some greater or lesser degree.

The insufficient or inappropriate security training of technical staff

Operational support staff may make errors of judgement or mistakes due to insufficient or inadequate training, not only of the technical functions of their role, but also of the need for security within their day-to-day activities. This is especially true of those

members of technical and security staff who have the ability to cause serious issues, whether deliberate or not.

The lack of appropriate security awareness training for users

User errors in protecting the organisation's information assets are generally the result of a lack of security awareness, and the organisation may not have considered that even basic security training for users is essential if they are to carry out their role properly and protect the organisation from harm.

The lack of monitoring mechanisms, including intrusion detection systems

We have already covered the need for robust monitoring and intrusion detection systems in organisations, but it is worth restating the need for these, not only to detect unauthorised activity from outside the organisation, but also from within, as this carries a significantly greater likelihood of success.

The lack of robust policies for the correct and appropriate use of systems, communications, media, social networking and messaging

The lack of robust policies for the correct and appropriate use of systems, communications, media, social networking and messaging remains a key vulnerability for many organisations, since without these there is little or no indication of what constitutes 'reasonable use' by staff, and what they might expect to happen if they fail to behave in an ethical manner.

The failure to review users' access rights whenever they change roles or leave the organisation

Also covered under access control vulnerabilities and frequently overlooked is that of the failure to review users' access rights whenever they change roles or leave the organisation, as this may present them with an opportunity to take advantage of situations in which the segregation of duties would normally be implicit.

The lack of a procedure to ensure the return of assets when leaving the organisation

On a similar theme, the failure to ensure the return of assets when leaving the organisation presents a very real opportunity for staff to take away physical assets – for example, laptops, tablets, smartphones and authentication devices – which, combined with the previous vulnerability, might leave the organisation open to future unauthorised access. Further, and equally important, would be their continued ability to access business information to which they have no entitlement.

Unmotivated or disgruntled staff

Unmotivated or disgruntled staff may take advantage of their role within the organisation to steal or otherwise damage valuable information assets as well as physical assets. This kind of vulnerability presents a difficult problem to the organisation, since, by the

time their feelings become known, the damage may already have been done. Hence, many other factors must be taken into account when finding ways to address this.

Unsupervised work by third-party organisations or by staff who work outside normal business hours


Many organisations outsource or otherwise make use of third-party service providers to carry out specific functions for the business. Inevitably, these companies and the staff who work for them are provided with access to the outsourcing organisation's network and systems. It is especially important that work carried out by these organisations or by regular internal staff who work outside normal business hours does not go unsupervised.

APPENDIX D – INFORMATION RISK CONTROLS

It is often wrongly assumed that a single control of any kind is sufficient to resolve a risk. In fact, it is frequently the case that more than one control is required, and these may often be controls of different types. It is common that a risk may have been reduced by some means, but leaving some level of risk that is shared with a third party before the residual risk is accepted.

There are three levels of control: strategic, tactical and operational. [Figure D.1](#) illustrates the overall structure of controls.

Figure D.1 Information risk controls



STRATEGIC CONTROLS

Strategic controls come in four flavours:

- Avoid or terminate – avoiding or terminating the risk can mean either stopping doing the activity, in which case there may well be some residual risk, or not commencing an activity, in which case the organisation may be left with an unsolved problem that the activity was intended to address.

- Reduce or modify – reducing or modifying the risk involves the application of suitable controls that result in a lower level of risk once they have been applied. There may remain some residual risk following treatment.
- Transfer or share – transferring or sharing the risk moves treatment of the risk to a third party who will take action if the risk materialises. Insurance is a common form of risk transfer. However, the organisation that transfers the risk still retains ownership of it.
- Accept or tolerate – when all other options have been discounted, acceptance of the risk is the final choice. This will also be the case when any of the other three options result in some degree of residual risk.

It must be remembered, however, that ignoring a risk is not the same as accepting it, and should never be an option.

TACTICAL CONTROLS

There are also four types of tactical control:

- Detective controls – detective controls are intended to identify and provide some form of alert when a threat is actually having a detrimental effect on an information asset or may be about to do so.
- Preventative controls – preventative controls are intended to stop a threat from having a detrimental effect on an information asset before the threat has any opportunity to do so.
- Directive controls – directive controls are intended to provide instruction on how to stop a threat from having a detrimental effect on an information asset or how to avoid activities that could initiate a detrimental result.
- Corrective controls – corrective controls are intended to prevent a threat from further detrimental activity, to recover from such activity or to prevent it from recurring.

OPERATIONAL CONTROLS

There are just three types of operational control:

- Procedural or people controls – procedural controls dictate the way in which actions must be taken and include such things as segregation of duties, change control mechanisms and the ongoing monitoring of risk.
- Physical or environmental controls – physical controls protect or change the environment in which information is stored and processed and include such things as locked doors and CCTV systems.
- Technical or logical controls – technical controls cover the technology-related aspects of information risk management and include such things as antivirus software and firewalls.

The following sections list the chief controls suggested by:

- the Centre for Internet Security Controls Version 8;
- ISO/IEC 27001:2017;
- NIST Special Publication 800-53 Revision 5.

THE CENTRE FOR INTERNET SECURITY CONTROLS VERSION 8

A number of organisations have published a list of the 18 most critical security controls. The list is based upon the Centre for Internet Security Controls Version 8. While this is not necessarily a comprehensive list of controls, it does provide a good starting point for organisations that have conducted their risk assessments but are unsure where to begin with risk treatment. The document may be downloaded from: <https://www.cisecurity.org/controls/>.

- Inventory and Control of Enterprise Assets (includes five sub-controls)
 - Organisations should actively manage (inventory, track and correct) all enterprise assets (end-user devices, including portable and mobile, network devices, non-computing/IoT devices and servers) connected to the infrastructure physically, virtually, remotely and those within cloud environments to accurately know the totality of assets that need to be monitored and protected within the enterprise. This will also support identifying unauthorised and unmanaged assets to remove or remediate.
- Inventory and control of software assets (includes seven sub-controls)
 - Organisations should actively manage (inventory, track and correct) all software (operating systems and applications) on the network so that only authorised software is installed and can execute, and that unauthorised and unmanaged software is found and prevented from installation or execution.
- Data Protection (includes 14 sub-controls)
 - Develop processes and technical controls to identify, classify, securely handle, retain and dispose of data.
- Secure Configuration of Enterprise Assets and Software (includes 12 sub-controls)
 - Establish and maintain the secure configuration of enterprise assets (end-user devices, including portable and mobile, network devices, non-computing/IoT devices and servers) and software (operating systems and applications).
- Account Management (includes six sub-controls)
 - Use processes and tools to assign and manage authorisation of credentials for user accounts, including administrator accounts as well as service accounts, to enterprise assets and software.
- Access Control Management (includes eight sub-controls)
 - Use processes and tools to create, assign, manage and revoke access credentials and privileges for user, administrator and service accounts for enterprise assets and software.

- Continuous Vulnerability Management (includes seven sub-controls)
 - Develop a plan to continuously assess and track vulnerabilities on all enterprise assets within the enterprise's infrastructure in order to remediate and minimise the window of opportunity for attackers. Monitor public and private industry sources for new threat and vulnerability information.
- Audit Log Management (includes 12 sub-controls)
 - Collect, alert, review and retain audit logs of events that could help detect, understand or recover from an attack.
- Email and Browser Protections (includes seven sub-controls)
 - Improve protections and detections of threats from email and web vectors, as these are opportunities for attackers to manipulate human behaviour through direct engagement.
- Malware Defences (includes seven sub-controls)
 - Prevent or control the installation, spread and execution of malicious applications, code or scripts on enterprise assets.
- Data Recovery (includes five sub-controls)
 - Establish and maintain data recovery practices sufficient to restore in-scope enterprise assets to a pre-incident and trusted state.
- Network Infrastructure Management (includes eight sub-controls)
 - Establish, implement and actively manage (track, report, correct) network devices, in order to prevent attackers from exploiting vulnerable network services and access points.
- Network Monitoring and Defence (includes 11 sub-controls)
 - Operate processes and tooling to establish and maintain comprehensive network monitoring and defence against security threats across the enterprise's network infrastructure and user base.
- Security Awareness and Skills Training (includes nine sub-controls)
 - Establish and maintain a security awareness programme to influence behaviour among the workforce to be security conscious and properly skilled to reduce cybersecurity risks to the enterprise.
- Service Provider Management (includes seven sub-controls)
 - Develop a process to evaluate service providers who hold sensitive data or are responsible for an enterprise's critical IT platforms or processes, to ensure these providers are protecting those platforms and data appropriately.
- Application Software Security (includes 14 sub-controls)
 - Manage the security life cycle of in-house developed, hosted or acquired software to prevent, detect and remediate security weaknesses before they can impact the enterprise.
- Incident Response Management (includes nine sub-controls)
 - Establish a programme to develop and maintain an incident response capability (for example, policies, plans, procedures, defined roles, training and communications) to prepare, detect and quickly respond to an attack.

- Penetration Testing (includes five sub-controls)
 - Test the effectiveness and resiliency of enterprise assets through identifying and exploiting weaknesses in controls (people, processes and technology) and simulating the objectives and actions of an attacker.

ISO/IEC 27001:2017 CONTROLS

Although the primary ISO standard for information risk management is ISO/IEC 27005, it contains no detailed information on suitable tactical or operational controls for risk treatment, restricting itself instead to the strategic level only. Instead, ISO/IEC 27001:2017 provides an annex containing a comprehensive list of 114 separate operational level controls, grouped into 14 categories.

ISO/IEC 27001, 27002 and 27005 may be purchased from the BSI online shop at <https://shop.bsigroup.com/>.

A more detailed description of the controls can be found in ISO/IEC 27002:2017 in its sections 5 to 18. The categories and their associated controls are summarised below.

A.5 Information security policies (2 controls)

The information security policy controls specify that organisations should define, approve and communicate information security policies to all stakeholders both within and outside the organisation. They also state the need to review these policies at intervals, or if anything changes that might impact on the policies.

Policies for information security	Review of the policies for information security
-----------------------------------	---

A.6 Organisation of information security (7 controls)

This area of controls introduces the need for defined responsibilities within the organisation, along with the segregation of duties in order to prevent misuse and abuse. Interestingly, this area introduces the concept that information security should be considered in all organisational projects, regardless of whether or not they are related to information security.

Additionally, this area deals with the need to address the security of mobile devices and teleworking options.

Information security roles and responsibilities	Segregation of duties
Contact with authorities	Contact with special interest groups
Information security in project management	Mobile device policy
Teleworking	

A.7 Human resource security (6 controls)

This area deals with three distinct topics. Firstly, those controls required prior to employment, including background checks and the employment terms and conditions that relate to information security. Secondly, those controls that apply during a period of employment, including the requirement to adhere to the organisation's information security policies, the need for security awareness training and the disciplinary procedures. Finally, those controls that apply to any employee's change of responsibilities or when they leave the organisation.

Screening	Terms and conditions of employment
Management responsibilities	Information security awareness, education and training
Disciplinary process	Termination or change of employment responsibilities

A.8 Asset management (10 controls)

Asset management controls are concerned with the identification of the organisation's information assets, allocation of ownership of them, their information classification, labelling and handling, their acceptable use and their return when employees leave the organisation.

Additionally, this area specifies controls concerned with the management, transfer and disposal of media such as memory sticks and DVDs.

Inventory of assets	Ownership of assets
Acceptable use of assets	Return of assets
Classification of information	Labelling of information
Handling of assets	Management of removable media
Disposal of media	Physical media transfer

A.9 Access control (14 controls)

One of the larger topic areas, access controls require a policy that provides for access to network resources to which users have been authorised, and the registration and deregistration and access provisioning processes. It goes on to cover the ongoing management of user access rights and their revocation or modification.

Additionally, access control covers the access to systems and application functions, password management and the access to privileged system utility software and source code.

Access control policy	Access to networks and network services
User registration and de-registration	User access provisioning
Management of privileged access rights	Management of secret authentication information of users
Review of user access rights	Removal or adjustment of access rights
Use of secret authentication information	Information access restriction
Secure log-on procedures	Password management system
Use of privileged utility programs	Access control to program source code

A.10 Cryptography (2 controls)

Cryptography controls include the provision of a cryptographic policy and the process for cryptographic key management.

Policy on the use of cryptographic controls	Key management
---	----------------

A.11 Physical and environmental security (15 controls)

Physical and environmental controls begin with those for controlling the physical perimeter of premises together with controls for restricting entry to the premises and areas within it, including secure areas such as computer rooms and less secure areas such as loading bays. It also includes the need to protect the premises from environmental threats and hazards.

The controls continue by considering the siting of equipment, the utilities that support it and the processes for removing it from the premises. Finally, they cover the need for clear desks and screens.

Physical security perimeter	Physical entry controls
Securing offices, rooms and facilities	Protecting against external and environmental threats
Working in secure areas	Delivery and loading areas
Equipment siting and protection	Supporting utilities
Cabling security	Equipment maintenance
Removal of assets	Security of equipment and assets off-premises
Secure disposal or re-use of equipment	Unattended user equipment
Clear desk and clear screen policy	

A.12 Operations security (14 controls)

Operational security controls focus on formal operating procedures, the need for change and capacity management and the separation of systems used for testing from those in the live environment. They go on to cover malware protection, event logging and the need to synchronise system clocks, the management of vulnerabilities and the requirement to restrict the installation of unauthorised software on systems.

Documented operating procedures	Change management
Capacity management	Separation of development, testing and operational environments
Controls against malware	Information backup
Event logging	Protection of log information
Administrator and operator logs	Clock synchronisation
Installation of software on operational systems	Management of technical vulnerabilities
Restrictions on software installation	Information system audit controls

A.13 Communications security (7 controls)

Communication security controls deal with the provision of security for networks and network services, and in particular they highlight the need to segregate networks carrying different security classifications of traffic.

They continue by describing the need to manage the transfer of information across and between organisations and include electronic communications such as email and social networking and the requirement for non-disclosure agreements.

Network controls	Security of network services
Segregation in networks	Information transfer policies and procedures
Agreements on information transfer	Electronic messaging
Confidentiality or non-disclosure agreements	

A.14 System acquisition, development and maintenance (13 controls)

This area deals both with the requirement for information security specifications to be included in the procurement process and how application services and their transactions passing over public networks require protection.

It continues by examining the need for development rules and change control procedures, technical reviews and secure design principles, whether development takes place within the organisation or outside it. Finally, it covers the security and system acceptance testing to be carried out on all systems.

Information security requirements analysis and specification	Securing application services on public networks
Protecting application services transactions	Security development policy
System change control procedures	Technical review of applications after operating platform changes
Restrictions on changes to software packages	Secure system engineering principles
Secure development environment	Outsourced development
System security testing	System acceptance testing
Protection of test data	

A.15 Supplier relationships (5 controls)

Supplier relationships are key to many organisations' day-to-day operations, and consequently the controls in this area relate to security policies and agreements between the organisation and its suppliers. In addition, the controls include the ongoing monitoring of the service delivery and how changes are carried out.

Information security policy for supplier relationships	Addressing security within supplier agreements
Information and communication technology supply chain	Monitoring and review of supplier services
Managing changes to supplier services	

A.16 Information security incident management (7 controls)

Dealing with security incidents requires its own set of controls, which include the allocation of responsibilities for the IM team and the reporting of both incidents and weaknesses within the organisation. The controls then focus on how incidents are assessed and dealt with, how lessons are learnt that can reduce future risk, and how the evidence (whether physical or electronic) of an incident should be collected, handled and stored.

Responsibilities and procedures	Reporting information security events
Reporting information security weaknesses	Assessment of and decision on information security events
Response to information security incidents	Learning from information security incidents
Collection of evidence	

A.17 Information security aspects of business continuity management (4 controls)

In this section, the standard considers the controls required for the planning and implementation of continuity of availability of information services, which include BC, DR and the redundant operation of systems.

Planning information security continuity	Implementing information security continuity
Verify, review and evaluate information security continuity	Availability of information processing facilities

A.18 Compliance (8 controls)

The controls for compliance deal with legal and regulatory requirements and the protection of essential records, how personally identifiable information is handled and, in cases where cryptography is used, how this is managed to ensure compliance with legislation.

Finally, the compliance section deals with independent reviews of the information security position within an organisation, and how individual parts of an organisation comply with its defined security policies.

Identification of applicable legislation and contractual requirements	Intellectual property rights
Protection of records	Privacy and protection of personally identifiable information
Regulation of cryptographic controls	Independent review of information security
Compliance with security policies and standards	Technical compliance review

NIST SPECIAL PUBLICATION 800-53 REVISION 5

Although the primary NIST publication on information risk management is Special Publication 800-30, it contains no detailed information on risk treatment or the selection of controls. However, NIST Special Publication 800-53 Revision 5 lists more than 300 separate operational level controls, grouped into 20 categories in its [Appendix F](#), and also maps them against ISO/IEC 27001 controls in its [Appendix H](#). The document can be downloaded free of charge from <https://csrc.nist.gov/publications/PubsSPs.html>.

The categories and their associated controls are summarised below.

AC Access Control (24 controls)

Policy and procedures	Account management
Access enforcement	Information flow enforcement
Separation of duties	Least privilege
Unsuccessful log-on attempts	System use notification
Previous log-on notification	Concurrent session control
Device lock	Session termination
Suspension and review – Control incorporated into AC-2 and AU-6	Permitted actions without identification or authentication
Automated marking	Security and privacy attributes
Remote access	Wireless access
Access control for mobile devices	Use of external systems
Information sharing	Publicly accessible content
Data mining protection	Access control decisions
Reference monitor	

AT Awareness and Training (6 controls)

Policy and procedures	Literacy training and awareness
Role-based training	Training records
Contacts with security groups and associations	Training feedback

AU Audit and Accountability (16 controls)

Policy and procedures	Event logging
Content of audit records	Audit log storage capacity
Response to audit logging process failures	Audit record review, analysis and reporting
Audit record reduction and report generation	Time stamps
Protection of audit information	Non-repudiation
Audit record retention	Audit record generation
Monitoring for information disclosure	Session audit
Alternate audit logging capability	Cross-organisational audit logging

CA Security Assessment and Authorisation (9 controls)

Policy and procedures	Control assessments
Information exchange	Security certification
Plan of action and milestones	Authorisation
Continuous monitoring	Penetration testing
Internal system connections	

CM Configuration Management (14 controls)

Policy and procedures	Baseline configuration
Configuration change control	Impact analysis
Access restrictions for change	Configuration settings
Least functionality	System component inventory
Configuration management plan	Software usage restrictions
User-installed software	Information location
Data action mapping	Signed components

CP Contingency Planning (12 controls)

Policy and procedures	Contingency plan
Contingency training	Contingency plan testing
Contingency plan update – Control incorporated into CP-2	Alternate storage site
Alternate processing site	Telecommunications services
System backup	System recovery and reconstitution
Alternate communications protocols	Safe mode
Alternative security mechanisms	

IA Identification and Authentication (12 controls)

Policy and procedures	Identification and authentication (organisational users)
Device identification and authentication	Identifier management
Authenticator management	Authentication feedback
Cryptographic module authentication	Identification and authentication (non-organisational users)
Service identification and authentication	Adaptive authentication
Re-authentication	Identity proofing

IR Incident Response (10 controls)

Policy and procedures	Incident response training
Incident response testing	Incident handling
Incident monitoring	Incident reporting
Incident response assistance	Incident response plan
Information spillage response	Integrated information security analysis team

MA Maintenance (7 controls)

Policy and procedures	Controlled maintenance
Maintenance tools	Non-local maintenance
Maintenance personnel	Timely maintenance
Field maintenance	

MP Media Protection (8 controls)

Policy and procedures	Media access
Media marking	Media storage
Media transport	Media sanitisation
Media use	Media downgrading

PE Physical and Environmental Protection (23 controls)

Policy and procedures	Physical access authorisations
Physical access control	Access control for transmission
Access control for output devices	Monitoring physical access
Visitor control	Visitor access records
Power equipment and cabling	Emergency shutoff
Emergency power	Emergency lighting
Fire protection	Environmental controls
Water damage protection	Delivery and removal
Alternate work site	Location of system components
Information leakage	Asset monitoring and tracking
Electromagnetic pulse protection	Component marking
Facility location	

PL Planning (9 controls)

Policy and procedures	Security and privacy plans
System security plan update	Rules of behaviour
Privacy impact assessment – Control incorporated into RA-8	Security-related activity planning – Control incorporated into PL-2
Concept of operations	Security and privacy architectures
Central management	Baseline selection
Baseline tailoring	

PM Programme Management (32 controls)

Information security programme plan	Information security programme leadership role
Information security and privacy resources	Plan of action and milestones process
System inventory	Measures of performance
Enterprise architecture	Critical infrastructure plan
Risk management strategy	Authorisation process
Mission and business process definition	Insider threat programme
Security and privacy workforce	Testing, training and monitoring
Security and privacy groups and associations	Threat awareness programme
Protecting controlled unclassified information on external systems	Privacy programme plan
Privacy programme leadership role	Dissemination of privacy programme information
Accounting of disclosures	Personally identifiable information quality management
Data governance body	Data integrity board
Minimisation of personally identifiable information used in testing, training and research	Complaint management
Privacy reporting	Risk framing
Risk management programme leadership roles	Supply chain risk management strategy
Continuous monitoring strategy	Purposing

PS Personnel Security (9 controls)

Policy and procedures	Position risk designation
Personnel screening	Personnel termination
Personnel transfer	Access agreements
External personnel security	Personnel sanctions
Position descriptions	

PT Personally Identifiable Information Processing and Transparency (8 controls)

Policy and procedures	Authority to process personally identifiable information
Personally identifiable information processing purposes	Consent
Privacy notice	System of records notice
Specific categories of personally identifiable information	Computer matching requirements

RA Risk Assessment (9 controls)

Policy and procedures	Security categorisation
Risk assessment	Risk assessment update – Control incorporated into RA-3
Vulnerability monitoring and scanning	Technical surveillance countermeasures survey
Risk response	Privacy impact assessments
Criticality analysis	Threat hunting

SA System and Services Acquisition (20 controls)

Policy and procedures	Allocation of resources
System development life cycle	Acquisition process
System documentation	Software usage restrictions – Control incorporated into CM-10 and SI-7
User-installed software – Control incorporated into CM-11 and SI-7	Security and privacy engineering principles
External system services	Developer configuration management
Developer testing and evaluation	Supply chain protection
Trustworthiness	Criticality analysis – Control incorporated into RA-9

Development process, standards and tools	Developer-provided training
Developer security and privacy architecture and design	Tamper resistance and detection
Component authenticity	Customised development of critical components
Developer screening	Unsupported system components
Specialisation	

SC System and Communications Protection (46 controls)

Policy and procedures	Separation of system and user functionality
Security function isolation	Information in shared system resources
Denial of service protection	Resource availability
Boundary protection	Transmission confidentiality and integrity
Transmission confidentiality – Control incorporated into SC-8	Network disconnect
Trusted path	Cryptographic key establishment and management
Cryptographic protection	Public access protections – Control incorporated into AC-2, AC-3, AC-5, AC-6, SI-3, SI-5, SI-7 and SI-10
Collaborative computing devices and applications	Transmission of security and privacy attributes
Public Key Infrastructure certificates	Mobile code
Voice over Internet Protocol – Control withdrawn	Secure name / address resolution service (authoritative source)
Secure name / address resolution service (recursive or caching resolver)	Architecture and provisioning for name / address resolution service
Session authenticity	Fail in known state
Thin nodes	Decoys
Platform-independent applications	Protection of information at rest
Heterogeneity	Concealment and misdirection
Covert channel analysis	System partitioning
Transmission preparation integrity – Control incorporated into SC-8	Non-modifiable executable programs
External malicious code identification	Distributed processing and storage

Out-of-band channels	Operations security
Process isolation	Wireless link protection
Port and I/O device access	Sensor capability and data
Usage restrictions	Detonation chambers
System time synchronisation	Cross-domain policy enforcement
Alternate communications paths	Sensor relocation
Hardware-enforced separation and policy enforcement	Hardware-based protection

SI System and Information Integrity (22 controls)

Policy and procedures	Flaw remediation
Malicious code protection	System monitoring
Security alerts, advisories and directives	Security and privacy function verification
Software, firmware and information integrity	Spam protection
Information input restrictions – Control incorporated into AC-2, AC-3, AC-5 and AC-6	Information input validation
Error handling	Information management and retention
Predictable failure prevention	Non-persistence
Information output filtering	Memory protection
Fail-safe procedures	Personally identifiable information quality operations
De-identification	Tainting
Information refresh	Information diversity
Information fragmentation	

SR Supply Chain Management (12 controls)

Policy and procedures	Supply chain risk management plan
Supply chain controls and processes	Provenance
Acquisition strategies, tools and methods	Supplier assessment and reviews
Supply chain operations security	Notification agreements
Tamper resistance and detection	Inspection of systems or components
Component authenticity	Component disposal

APPENDIX E – METHODOLOGIES, GUIDELINES AND TOOLS

The *Collins English Dictionary* defines a methodology as a way of proceeding or doing something, especially a systematic or regular one.

The discipline of risk management has its fair share of methodologies, some of which are described here.

METHODOLOGIES

CORAS

CORAS is an open-source risk management tool available from SourceForge without the additional scope included in Sherwood Applied Business Security Architecture (SABSA; see below). It consists of eight distinct steps,¹ which follow the generic risk management principles:

- Step 1 is the initial preparations for a risk analysis. The main objectives are to understand what the target is and what the size of the analysis will be.
- Step 2 is to establish the overall goals of the analysis and the target to be analysed. The objectives are to achieve a common understanding of the target and the major areas of concern.
- Step 3 aims to ensure a common understanding of the target of analysis, including its focus, scope and main assets. This step will conduct a rough, high-level analysis to identify major threat scenarios, vulnerabilities and enterprise level risks.
- Step 4 aims to ensure that the background documentation for the rest of the analysis, including the target, focus and scope, is correct and complete and is approved. Step 4 furthermore includes deciding the risk evaluation criteria for each asset. This analysis step concludes the context establishment.
- Step 5 is risk identification using structured brainstorming. Risk identification involves a systematic identification of threats, unwanted incidents, threat scenarios and vulnerabilities with respect to the identified assets. The results are documented by means of CORAS threat diagrams.
- Step 6 aims to determine the risk levels of the risks that are represented by the identified unwanted incidents.

¹ See Lund, M.S., Solhaug, B. and Stølen, K. (2011) *Model-Driven Risk Analysis: The CORAS Approach*. Berlin, Heidelberg: Springer-Verlag.

- Step 7 examines which of the identified risks are acceptable, and which of the risks must be further evaluated for possible treatment.
- Step 8 is concerned with the identification and analysis of treatments. The risks that are found to be unacceptable are evaluated to find means to reduce them. Since treatments can be costly, they are assessed with respect to their cost-benefit, before a final treatment plan is made.

The platform-independent CORAS risk management tool is available as a free download from <https://www.coras.sourceforge.net/>, as is a guided tour of the CORAS method.

Factor Analysis of Information Risk (FAIR)

FAIR is a framework developed in the early 2000s and now maintained by the FAIR Institute (<https://www.fairinstitute.org>).

FAIR follows a similar route to risk assessment as other methodologies, but uses slightly different terminology together with its method of deriving the likelihood of an event. It carries out the risk assessment in a different order and consists of four stages:

Stage 1 – Identify scenario components (much the same as asset identification)

1. Identify the asset at risk.
2. Identify the threat community under consideration. The threat community is used to group different threat sources together, such as internal staff or the hacking community.

Stage 2 – Evaluate the loss event frequency (the likelihood)

3. Estimate the probable threat event frequency that, within a given timeframe, a threat agent will act against an asset.
4. Estimate the threat capability or the probable level of force that a threat agent is capable of applying against an asset.
5. Estimate the strength of existing controls, which, over a given timeframe, are measured against a baseline level of force.
6. Derive the vulnerability, or the probability, that an asset will be unable to resist the actions of a threat agent.
7. Derive the loss event frequency, or the probable frequency, within a given timeframe, that a threat agent will inflict harm upon an asset.

Stage 3 – Evaluate probable loss magnitude (the impact or consequence)

8. Estimate worst-case loss.
9. Estimate probable loss.

Stage 4 – Derive and articulate risk (the risk analysis)

10. Derive and articulate risk, by combining the probable frequency and probable magnitude of loss.

Although FAIR covers risk assessment, it does not address an organisation's risk appetite or tolerance and makes no mention of risk treatment, and so is only really useful for the earlier stages of an information risk management programme.

Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE)

OCTAVE originates from the Carnegie Mellon Software Engineering Institute.

The OCTAVE method uses a three-phased approach to examine organisational and technology issues, assembling a comprehensive picture of the organisation's information security needs. It is aimed at organisations with around 300 or more people. The method is organised into a progressive series of workshops, and the process is supported with guidance, worksheets and questionnaires.

In phase 1 the organisation builds asset-based profiles, identifying assets, threats, current practices, organisational vulnerabilities and security requirements. This is achieved in four distinct processes:

1. Identify senior management knowledge.
2. Identify operational area management knowledge.
3. Identify staff knowledge.
4. Create threat profiles.

Moving to the second phase, which involves two processes, the organisation will identify infrastructure vulnerabilities:

1. Identify key components.
2. Evaluate selected components and establish technical vulnerabilities.

In the third and final phase, the organisation will develop security strategies and plans:

1. Conduct risk analysis.
2. Develop protection strategy.

The method then uses a catalogue of practices to apply appropriate controls.

The strategic practices are:

- security awareness and training;
- security strategy;
- security management;
- security policies and regulations;
- collaborative security management;
- contingency planning and disaster recovery.

The operational practices are:

- Physical security – physical security plans and procedures; physical access control; monitoring and auditing physical security.
- Information technology security – system and network management; system administration tools; monitoring and auditing IT security; authentication and authorisation; vulnerability management; encryption; security architecture and design.
- Staff security – incident management; general staff practices.

OCTAVE-S

Whereas OCTAVE relies on a progressive series of workshops involving managers from different levels within the organisation, OCTAVE-S, which is designed for organisations with fewer than 100 people, uses the skills and experience of a smaller team of people (typically three to five in number) who have extensive knowledge of the organisation.

It also differs from OCTAVE in that the OCTAVE-S worksheets and guidance already contain security concepts, which permits less experienced information risk management practitioners to assess a very broad range of risks, many of which may be unfamiliar to them.

Finally, OCTAVE-S is less demanding on information relating to the organisation's information infrastructure, since smaller organisations tend to have less capability to use vulnerability tools.

OCTAVE Allegro

OCTAVE Allegro is a more streamlined version of OCTAVE, and takes a slightly different approach from OCTAVE in that, although it makes use of progressive workshops, the focus is more on the information assets themselves, the use to which they are put, stored, transported and processed and how they are impacted by threats, vulnerabilities and disruptions.

The method consists of four stages:

Stage 1 – Establish drivers

1. Establish risk measurement criteria.

Stage 2 – Profile assets

2. Develop information asset profile.
3. Identify information asset containers.

Stage 3 – Identify threats

4. Identify areas of concern.
5. Identify threat scenarios.

Stage 4 – Identify and mitigate risks

6. Identify risks.
7. Analyse risks.
8. Select mitigation approach.

Information on all three OCTAVE methodologies can be found at <https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=309051>.

SABSA

Developed by John Sherwood in 1995 and published in 1996 as *SABSA: a method for developing the enterprise security architecture and strategy*, the SABSA framework has evolved as a 'best practice' method for delivering cohesive information security solutions to enterprises. It is a six-layer model covering all four parts of the IT life cycle: strategy, design, implementation, and management and operations.

This makes SABSA a very powerful tool that is not limited just to risk management. It is designed to ensure that the security needs of enterprises are met and that security services are designed, delivered and supported as an integral part of an IT management infrastructure and it provides guidance for aligning architecture with business value.

SABSA looks at security architecture from several perspectives:

- The business view, referred to as the contextual security architecture, which examines the business assets, risks, processes, organisation, geography and time dependencies.
- The architect's view, referred to as the conceptual security architecture, which examines the strategic and risk management objectives and the organisation's high-level roles and responsibilities.
- The designer's view, referred to as the logical security architecture, which examines the policies and logical security services such as authentication, confidentiality and integrity protection, non-repudiation and system assurance, people, their roles and responsibilities, and the security domains.
- The builder's view, referred to as the physical security architecture, which examines the business data model, rules, security mechanisms, people dependencies and security technology infrastructure.
- The tradesman's view, referred to as the component security architecture, which examines the standards, tools and products used to implement the overall security architecture.
- Cutting across all these is the service manager's view, referred to as the security service management architecture, which examines the delivery management, operational risk management, personnel management and environmental management of the security infrastructure.

Each of these security architectures is mapped against a series of basic questions: what? why? how? who? where? and when? to form the SABSA Matrix.

SABSA mimics the PDCA model as its life cycle, but names the parts slightly differently as 'Strategy and planning'; 'Design'; 'Implement'; 'Manage and measure'. It then examines the use of business attributes to provide a link between the organisation's business requirements and the technology and process design, either in the form of ICT business attributes or general business attributes:

ICT attributes:

- user;
- management;
- operational;
- risk management;
- legal and regulatory;
- technical strategy;
- business strategy.

High-level general business attributes:

- financial;
- physical;
- human;
- process;
- strategic;
- system.

In terms of the generic information risk management method, SABSA also includes the processes to provide consultation and communication, referred to as 'communicate', and monitor and review, referred to as 'assure', and has the capability to do this at four distinct levels.

More information on SABSA is available from <https://www.sabsa.org/>.

OTHER GUIDELINES AND TOOLS

BS 7799-3

BS 7799-3:2017 – Information security management systems. Guidelines for information security risk management.

Recently updated, BS 7799-3 summarises the information risk management process extremely well, with numerous references to the ISO IEC 27001:2017 standard. It is well worth obtaining a copy as background reference material.

Its main sections follow the standard risk management process and are a revision of the earlier ISO/IEC 27005 standard.

- An overview of the information security risk assessment and risk treatment process in section 4, including the information security risk treatment process diagram (see [Chapter 1](#) of this book).
- A description on communication and consultation in section 5.
- A description of context establishment in section 6, with examples of:
 - logarithmic likelihood scales;
 - logarithmic consequence scales;
 - indicator scales.
- Section 7 covers risk identification and analysis, and contains an example of scenarios that give coverage of the controls in ISO/IEC 27001:2017, Annex A.
- Section 8 discusses information security risk treatment.
- Section 9 describes the verification of necessary controls, and includes a diagrammatic representation of the cross-checking processes.
- Section 10 briefly discusses approval of the risk treatment controls.
- Section 11 briefly describes operation of the overall organisational process.
- Section 12 examines monitoring, audit and review.

BS 7799-3:2017 can be obtained in PDF or hard copy formats from the BSI online shop at <https://www.bsigroup.com/Shop>.

NIST SP800-30

Guide for Conducting Risk Assessments – NIST Special Publication 800-30 Revision 1

While there are some minor differences in the approach of SP800-30 from those described elsewhere in this book, it remains an extremely comprehensive and detailed information risk management standard.

One of the most useful sections is its final two-page Appendix L – Summary of tasks:

Step 1 Prepare for risk assessment

- Identify purpose.
- Identify scope.
- Identify assumptions and constraints.
- Identify information sources.
- Identify risk model and analytic approach.

Step 2 Conduct risk assessment

- Identify threat sources.
- Identify threat events.
- Identify vulnerabilities and predisposing conditions.
- Determine likelihood.
- Determine impact.
- Determine risk.

Step 3 Communicate and share risk assessment results

- Communicate risk assessment results.
- Share risk-related information.

Step 4 Maintain risk assessment

- Monitor risk factors.
- Update risk assessment.

The first chapter introduces the content and organisation of the remainder of the standard, which consists of:

Chapter 2 The fundamentals

- Risk management process.
- Risk assessment.
- Key risk concepts.
- Application of risk assessments.

Chapter 3 The process

- Preparing for the risk assessment.
- Conducting the risk assessment.
- Communicating and sharing risk assessment information.
- Maintaining the risk assessment.

Appendix A References

Appendix B Glossary

Appendix C Acronyms

Appendix D Threat sources

Appendix E Threat events

Appendix F Vulnerabilities and predisposing conditions

Appendix G	Likelihood of occurrence
Appendix H	Impact
Appendix I	Risk determination
Appendix J	Informing risk response
Appendix K	Risk assessment reports
Appendix L	Summary of tasks (listed above)

As with all NIST standards, they may be downloaded free of charge from <https://doi.org/10.6028/NIST.SP.800-30r1>.

Risk assessment tools

The internet lists numerous risk management software tools, many of which are not especially well-suited to use in information risk management, some of which are aimed at the larger enterprise and others that can be better used by smaller organisations.

It is not intended to endorse or make any recommendations as to which (if any) of the tools are best suited to information risk management use, but instead to provide some suggestions as to the key attributes you may wish to consider when selecting one. Please note that these do not include the 'usual' considerations, such as cost, support capability and so on.

1. Does the tool address any or all of the standards to which the organisation is working or aims to reach?
2. Does the tool provide a complete risk management overview, or is it limited to risk assessment only (i.e. no risk treatment)?
3. Does the tool contain predefined:
 - types of asset;
 - types of impact;
 - threats;
 - vulnerabilities;
 - controls;
 and can additional ones be user-defined?
4. Can the tool import an asset list from a spreadsheet or database?
5. Does the tool permit the user to break down the impact assessment by confidentiality, integrity and availability?
6. Does the tool permit more than one threat or vulnerability for each asset?
7. Does the tool permit more than one control for each risk identified?
8. Does the tool provide output in the form of a risk register?
9. Does the tool provide output in graphical form?

10. Is the tool scalable to enterprise level?
11. Is the tool single-user or multi-user?
12. Can the tool be run on multiple operating systems (for example, Windows, MacOS, Unix and Linux), and are there mobile applications (for example, iPhone, iPad, Android) that can interwork with it?
13. Are you able to download a trial version of the tool?

APPENDIX F – TEMPLATES

The following templates may be used as a basis for carrying out risk assessments:

- impact assessment;
- threat/hazard assessment;
- vulnerability assessment;
- existing controls assessment;
- risk register.

Impact assessment template

Impact/Asset			Date	
Asset owner			Analyst	
Asset location			Reference	
	Primary impacts		Secondary impacts	
	Direct impacts	Indirect impacts	Direct impacts	Indirect impacts
Operational				
Score				
Legal and regulatory				
Score				
Reputational				
Score				
People-related				
Score				
Financial				
Score				
Total financial impact				
Total impact rating				

VH = Very high H = High M = Medium L = Low VL = Very low

Threat assessment template

Date	Threat description	Reference	
Hacking			
Environmental threats and hazards			
Errors and failures			
Social engineering			
Misuse and abuse			
Physical threats			
Malware			
Operating systems affected			
Applications affected			
Information types affected			
Previous attack history (if known)			
Previous success rate (if known)			
Attack motivation (if known)			

Vulnerability assessment template

Date	Reference
Vulnerability description	
Access control	
Systems acquisition, development and maintenance	
Physical and environmental	
People-related	
Communications and operations	
Operating systems affected	
Applications affected	
Information types affected	
Previous attack history (if known)	
Previous success rate (if known)	
Attack motivation (if known)	

Existing controls assessment template

Date		Reference	
Description of controls			
Asset name			
Asset location			
Preventative controls	Physical or environmental		
	Technical or logical		
Detective controls	Physical or environmental		
	Technical or logical		
	Procedural or people		
Directive controls	Procedural or people		
Corrective controls	Physical or environmental		
	Technical or logical		
	Procedural or people		

Risk register template

[illegible]

* Treatment: X = Avoid; S = Share; R = Reduce; A = Accept

APPENDIX G – HMG CYBERSECURITY GUIDELINES

We have stated already that information risk management is not only about cybersecurity, but that it encompasses other areas, especially including the risks associated with people who, at the end of the day, are actually the cause of many of the information security problems. That said, cybersecurity will remain a key part of the information risk management programme for many organisations, and it would be highly remiss to ignore it.

We have already (in [Chapters 10 and 11](#)) discussed the way in which the UK government deals with information risk management in its own environment. In June 2014 the government launched a new scheme to improve and promote cybersecurity, its primary objective being 'to make the UK a safer place to conduct business online'.¹

Firstly, let us take a very brief look at what cybersecurity actually is. The UK's Cyber Security Strategy 2016–2021 defines cyberspace as:

the interdependent network of information technology infrastructures that includes the Internet, telecommunications networks, computer systems, internet-connected devices and embedded processors and controllers. It may also refer to the virtual world or domain as an experienced phenomenon, or abstract concept.

We can therefore suggest that cybersecurity is the art or science of protecting this infrastructure against accidental or deliberate loss or harm.

There are two separate government initiatives:

- HMG Cyber Essentials Scheme² from the Department for Digital, Culture, Media & Sport (DCMS).
- 10 Steps to Cyber Security,³ produced by the National Cyber Security Centre (NCSC).

HMG CYBER ESSENTIALS SCHEME

The Cyber Essentials Scheme defines a set of controls that, when properly implemented, will provide organisations with basic protection from the most prevalent forms of

¹ <https://www.ncsc.gov.uk>.

² See <https://www.gov.uk/government/publications/cyber-essentials-scheme-overview>.

³ See <https://www.ncsc.gov.uk/collection/10-steps-to-cyber-security>.

threats coming from the internet. In particular, it focuses on threats that require low levels of attacker skill and are widely available online.

Risk management is the fundamental starting point for organisations to take action to protect their information. However, given the nature of the threat, the government believes that action should begin with a core set of security controls, which all organisations – large and small – should implement. Cyber Essentials defines what these controls are.

The scheme provides for two distinct levels of certification:

- **Cyber Essentials** certification is awarded on the basis of a verified self-assessment. An organisation undertakes their own assessment of their implementation of the Cyber Essentials control themes via a questionnaire, which is approved by a senior executive such as the chief executive officer (CEO). This questionnaire is then verified by an independent certification body to assess whether an appropriate standard has been achieved, and certification can be awarded. This option offers a basic level of assurance and can be achieved at low cost.
- **Cyber Essentials Plus** offers a higher level of assurance through the external testing of the organisation's cybersecurity approach. Given the more resource-intensive nature of this process, it is anticipated that Cyber Essentials Plus will cost more than the foundation Cyber Essentials certification.

The scheme recommends the use of controls in five separate areas, as follows.

Securing the internet connection

By installing a firewall between the organisation's network and the internet, incoming traffic can be analysed to establish whether or not it should be allowed into the network.

Securing all devices and software

By checking the settings of new software and devices, the level of security should be raised, for example by disabling any functions that are not required.

All devices and accounts should be password protected. These should be easy for the user to remember, but difficult for an attacker to guess. Default passwords should always be changed at the first opportunity.

Two-factor authentication will increase the level of security still further.

Control access to data and services

Accounts with administrative privileges should only be allocated to those administrators who have a genuine business need to have them. Where a user does not have this requirement, they should be allocated a standard, non-privileged account.

Users should not be allowed to make changes to their operating system or application software, nor should they be allowed to install new software, which may contain malware.

Protect against viruses and other malware

Operating systems should be protected against malware and other viruses by installing suitable antivirus software. This can be further enhanced by the process of 'whitelisting', which allows only those programs that have been tested and approved for use to be run on a device.

Furthermore, the use of applications that support 'sandboxing' should be preferred. These operate within a controlled environment that has restricted access to other devices and network areas.

Keep devices and software up to date

By keeping operating systems and applications up to date by applying the manufacturer's latest patches, the opportunity for an attacker to intrude will be much reduced.

There are two main supporting documents currently available:

- *Cyber Essentials: Requirements for IT Infrastructure* – downloadable from <https://www.ncsc.gov.uk/files/Cyber-Essentials-Requirements-for-IT-infrastructure.pdf>.
- *Cyber Essentials Plus: Illustrative Test Specification* – downloadable from <https://www.ncsc.gov.uk/files/Cyber-Essentials-Plus-Illustrative-Test-Specification-April-2020.pdf>.

10 STEPS TO CYBER SECURITY

The 10 Steps to Cyber Security⁴ advice originates from the NCSC, which is part of GCHQ.

The measures detailed in the cybersecurity advice collectively represent a good foundation for effective information risk management. The degree of implementation of these steps will vary between organisations depending on their risks to the individual business.

This Crown Copyright material is included here under the UK Open Government Licence.⁵ The 10 areas are discussed below.

Risk management regime

Organisations should introduce and operate a regime of information risk management, which must be supported both by an effective governance structure and by the organisation's board and senior management team. The organisation must communicate clearly its approach to risk management and must develop appropriate policies, procedures and practices. The aim is to ensure that all employees, contractors and suppliers are aware of the regime and the constraints to which they must adhere.

⁴ See <https://www.ncsc.gov.uk/collection/10-steps-to-cyber-security>.

⁵ See <http://www.nationalarchives.gov.uk/doc/open-government-licence/version/3/>.

Secure configuration

Organisations must have an approach to identify baseline technology builds and processes in order to ensure that configuration management can improve the security of systems. Organisations should ensure that all unnecessary functionality is either removed or disabled from systems, and that known vulnerabilities are fixed in a timely manner, usually via patching. Failure to do so is likely to result in increased risk of compromise of systems and information.

Network security

In order to reduce the chances of network attacks succeeding, organisations should create and implement simple policies and appropriate procedures. Many organisations' networks span multiple sites and the use of remote working, and the use of cloud services makes it difficult to define fixed network boundaries. Instead of focusing on physical connections, organisations should consider where their data are stored and processed, and where an attacker might have the opportunity to access them.

Managing user privileges

Users who are provided with unnecessary system privileges or data access rights may more easily impact, misuse or compromise information and services on the organisation's networks. Non-administrative users should be provided with the minimal level of system privileges and rights needed for them to undertake their role. The granting of elevated or administrative privileges should be carefully controlled and managed.

User education and awareness

In order that users may play their role in their organisation's security, it is important that both the security rules and the technology provided enable users to fulfil their role as well as help to keep the organisation secure. Delivery of suitable security expertise awareness programmes and training help to establish a culture of security.

Incident management

Organisations should establish appropriate incident management policies and processes in order to improve resilience, support the BC function, improve customer and stakeholder confidence and potentially reduce the impact of incidents. Where necessary, organisations should employ recognised sources of specialist incident management expertise.

Malware prevention

The term malware covers any code or content that could have a malicious or undesirable impact on systems. Any exchange of information may carry a degree of risk that malware might be present, and which could seriously impact the organisation's systems and services. Anti-malware policies can help to reduce the risk as part of an overall 'defence in depth' approach.

Monitoring

In order to detect actual or attempted attacks on systems and business services, organisations should introduce system monitoring in order to respond effectively to attacks. Additionally, monitoring permits organisations to ensure that systems are being used appropriately, and is often a capability required in order to comply with legal or regulatory requirements.

Removable media controls

Malware is frequently introduced by means of removable media, which can also enable the accidental or deliberate export of sensitive data. Organisations should have a strict policy regarding the need for users to make use of removable media and should apply appropriate security controls to limit its use.

Home and mobile working

Mobile working and remote system access exposes risks that must be managed. Organisations should introduce risk-based policies and procedures supporting mobile working or remote access to systems by all users and service providers. Users must be trained on the secure use of their mobile devices in whatever environments they are working.

APPENDIX H – REFERENCES AND FURTHER READING

I have attempted to include a number of useful references and items for further reading in this section. Where articles and documents are downloadable, I have provided the appropriate URL, but you should be aware that some organisations make regular changes to their websites and these links cannot be guaranteed. Where appropriate, I have provided a brief synopsis of the material listed.

The reference material falls into the following areas:

- primary UK legislation;
- good practice guidelines;
- other reference material;
- the NCSC Certified Cyber Professional scheme;
- other UK government publications;
- risk management methodologies;
- UK and international standards.

PRIMARY UK LEGISLATION

Data Protection Act 2018. Her Majesty's Stationery Office. Available from <https://www.gov.uk/government/collections/data-protection-act-2018>.

The Computer Misuse Act 1990.¹ Her Majesty's Stationery Office. Available from <https://www.legislation.gov.uk/ukpga/1990/18/contents>.

The Police and Criminal Evidence Act 1984. Her Majesty's Stationery Office. Available from <https://www.legislation.gov.uk/ukpga/1984/60/contents>.

The Official Secrets Act 1989. Her Majesty's Stationery Office. Available from <https://www.legislation.gov.uk/ukpga/1989/6/contents>.

The Freedom of Information Act 2000. Her Majesty's Stationery Office. Available from <https://www.legislation.gov.uk/ukpga/2000/36/contents>.

¹ This Act is under review at the time of writing.

The Regulation of Investigatory Powers Act 2000. Her Majesty's Stationery Office. Available from <https://www.legislation.gov.uk/ukpga/2000/23/contents>.

The Copyright, Designs and Patents Act 1988. Her Majesty's Stationery Office. Available from <https://www.legislation.gov.uk/ukpga/1988/48/contents>.

Control of Major Accident Hazards Regulations 2015. Health and Safety Executive. Available from <https://www.legislation.gov.uk/uksi/2015/483/contents/made>.

Civil Contingencies Act 2004. Her Majesty's Stationery Office. Available from <https://www.legislation.gov.uk/ukpga/2004/36/contents>.

GOOD PRACTICE GUIDELINES

Good practice guidelines 2018. The Business Continuity Institute. Available from <https://www.thebci.org/product/good-practice-guidelines-2018-edition---download.html>.

Synopsis: This document is considered by many to be the definitive work on BCM. It has undergone several iterations during its lifetime, and has been consistent in its approach, while keeping in step with developments and standards. The document is free to BCI members.

OTHER REFERENCE MATERIAL

The Traffic Light Protocol (TLP). European Network and Information Security Agency (ENISA). Available from <https://www.enisa.europa.eu/topics/csirts-in-europe/glossary/considerations-on-the-traffic-light-protocol>.

Synopsis: The TLP is used internationally to represent the sensitivity of information and under what circumstances it may be shared with other parties. It is not designed to replace an information classification scheme, but is used especially when sharing security-related information with partners. The description of the TLP also appears in ISO/IEC 27010:2015.

The Capability Maturity Model. Carnegie Mellon University.

Synopsis: The CMM was originally developed as a tool for assessing the performance of organisations developing software for the US Department of Defense. It can be adapted for many uses, and many organisations have adopted it for such areas as BCM and information risk management. See <https://www.itgovernance.co.uk/capability-maturity-model>.

Centre for Internet Security Controls Version 8. The Centre for Internet Security. Available from <https://cisecurity.org/controls/>.

Synopsis: Developed originally by the US military community, the list of critical security controls was compiled in order to reduce the data losses experienced by the American defence industrial base. Use is free under creative commons, and many organisations worldwide have now adopted the controls as good practice recommendations.

The CII Sec Skills Framework. The Chartered Institute of Information Security (CII Sec). Available from https://www.ciisec.org/Skills_Framework.

Synopsis: The CII Sec Skills Framework was developed by the Institute as a means of assessing potential members for their skills and experience in the information security industry. The framework consists of four levels of skill – awareness, basic application, skilful application and expert – and assesses skills in information security management, information risk management, implementing secure systems, information assurance methodologies and testing, operational security management, IM, audit, assurance and review, BCM, information systems research and teamwork and leadership. The framework is used for assessment in the CCP scheme.

CII Sec have produced two further useful documents:

*The CII Sec Roles Framework,*² which sets out a typical set of skills expected of information security and information assurance professionals in the effective performance of their roles.

*The CII Sec Knowledge Framework,*³ which expands upon the widely used CII Sec Skills Framework allowing information security professionals to have a consistent view of cybersecurity and information security.

A structured approach to enterprise risk management (ERM) and the requirements of ISO 31000. The Federation of European Risk Management Associations. Available from <https://www.ferma.eu/app/uploads/2011/10/a-structured-approach-to-erm.pdf>.

Synopsis: Following the publication of the *ISO 31000 Risk Management Principles and Guidelines* in 2009, this document has been produced by the Institute of Risk Management (IRM), the Association of Insurance and Risk Managers (Airmic) and the Public Risk Management Association (PRIMA), and provides up-to-date guidance on the implementation of ERM in the context of the new ISO standard.

The Standard for Information Assurance for Small and Medium Sized Enterprises (IASME). Available from <https://iasme.co.uk>.

Synopsis: IASME operate a number of schemes, including Cyber Essentials (partnered with the NCSC) and IASME Governance, which may be used as an alternative to ISO/IEC 27001 and includes the Cyber Essentials assessment together with GDPR requirements.

The Open Group Standard, Risk Taxonomy (O-RT), Version 2.0, October 2013.

NCSC CERTIFIED PROFESSIONAL SCHEME

Several of these documents are described in greater detail in [Chapter 10](#).

The CCP scheme. Available from: <https://ncsc.gov.uk/information/certified-cyber-professional-assured-service>.

² https://www.ciisec.org/Roles_Framework.

³ https://www.ciisec.org/Knowledge_Framework.

Synopsis: The NCSC is the information security arm of GCHQ and runs the CCP scheme, described as ‘a certification which is awarded to those who demonstrate their sustained ability to apply their skills, knowledge and expertise in real-world situations’. It makes use of the IISP Skills Framework⁴ and the SFIA Framework for the accreditation process. Along with the APM Group and the CIISec, CREST and RHUL, BCS, the Chartered Institute for IT is approved by NCSC to carry out certification.

OTHER UK GOVERNMENT PUBLICATIONS

Several of these documents are described in greater detail in [Chapter 11](#) – HMG Security-related Documents.

HMG Security Policy Framework version 1.1 (May 2018). Government Digital Service. Available from <https://www.gov.uk/government/publications/security-policy-framework>.

Synopsis: This document begins by stating the overarching principles for government security and the security outcomes that are required.

It continues by describing the means by which these should be achieved, including: good governance; culture and awareness; risk management; information; technology and services; personnel security; physical security; and preparing for and responding to security incidents.

Government Security Classifications. The current definitive document, version 1.1 dated May 2018, describes the current government security classifications and can be downloaded from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/715778/May-2018_Government-Security-Classifications-2.pdf.

Synopsis: The updated classification scheme simplifies the earlier scheme and reduces the number of categories to three: Official, Secret and Top Secret.

HM Treasury, *Orange Book: Management of Risk – Principles and Concepts*. (October 2004). Government Digital Service. Available at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/866117/6.6266_HMT_Orange_Book_Update_v6_WEB.PDF.

Synopsis: Not as prescriptive as the other government documents, the *Orange Book* is an excellent guide to risk management from HM Treasury. Now over 10 years old, it still has considerable value and uses plain language to describe the risk management process.

The HMG Cyber Essentials scheme. Available at <https://www.gov.uk/government/publications/cyber-essentials-scheme-overview>.

⁴ See <https://www.ciisec.org/>.

The 10 Steps to Cyber Security. Available at <https://www.ncsc.gov.uk/collection/10-steps-to-cyber-security>.

RISK MANAGEMENT METHODOLOGIES

These methodologies are described in greater detail in [Appendix E – Methodologies, Guidelines and Tools](#):

CORAS Risk Assessment Platform. SourceForge. Available from <https://coras.sourceforge.net/index.html>.

See also: Mass Soldal Lund, Bjørnar Solhaug and Ketil Stølen (2011) *Model-Driven Risk Analysis: The CORAS Approach*. Berlin, Heidelberg: Springer-Verlag.

FAIR (Factor Analysis of Information Risk). Risk Management Insight. Available from <https://www.fairinstitute.org>.

The OCTAVE method (operationally critical threat, asset, and vulnerability evaluation).

The OCTAVE-S method – designed for use by smaller organisations.

The OCTAVE Allegro method, a streamlined approach for information security assessment and assurance. Carnegie Mellon University.

Details of all three OCTAVE methodologies are available from the Carnegie Mellon University at <https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=309051>.

SABSA. The SABSA Institute. Available from <https://www.sabsa.org/>.

UK AND INTERNATIONAL STANDARDS

The first part of this section lists and briefly describes some of the ISO/IEC 27000 series standards, especially those that are part of the so-called ISMS Family of Standards, including those in the 2703x range, which specify control-specific requirements.

Vocabulary standards

ISO/IEC 27000:2020 – Information technology – Security techniques – Information security management systems – Overview and vocabulary.

Synopsis: This standard includes three main sections: terms and definitions (some of which are replicated earlier in this book), terminology and definitions. It then describes the purpose and approach to establishing an information security management system (ISMS), and finally itemises and describes the 27000 family of standards that are relevant to an ISMS.

ISO Guide 73:2009 – Risk management – Vocabulary – Guidelines for use in standards.

Synopsis: ISO Guide 73:2009 provides the definitions of generic terms related to risk management. It aims to encourage a mutual and consistent understanding of, and a coherent approach to, the description of activities relating to the management of risk and the use of uniform risk management terminology in processes and frameworks dealing with the management of risk.

Requirement standards

ISO/IEC 27001:2017 – Information technology – Security techniques – Information security management systems – Requirements

Synopsis: This standard is recognised as the main document for the establishment, implementation, operation, monitoring, review, maintenance and improvement of an ISMS. It is against this standard that organisations can be audited and certified. The annex to the document lists 114 separate controls in 14 categories, including the control objectives. This list is greatly enhanced in ISO/IEC 27002:2017, which expands the detail of each control by providing implementation guidance (see below).

ISO/IEC 27006:2015 – Information technology – Security techniques – Requirements for bodies providing audit and certification of information security management systems

Synopsis: Aimed those bodies who provide certification, this standard guides auditors through the process of conducting an ISMS audit against ISO/IEC 27001. In particular, it provides guidance on the complexity of audit scope depending on the size of the organisation concerned, areas of competence required by auditors in this area, the likely time it will take to undertake an audit, and, for all 114 controls, details whether they are organisational or technical controls, whether they can be verified through system testing and whether visual inspection of documentation is required. It also provides further guidance on areas where such information is likely to be found.

The standard makes several references to ISO/IEC 17021:2015 – Conformity assessment – Requirements for bodies providing audit and certification of management systems.

Guideline standards

ISO/IEC 27002:2017 – Information technology – Security techniques – Code of practice for information security management

Synopsis: This standard contains the recommended best-practice controls for treatment of risks in an ISMS. Each of the 114 controls listed in ISO/IEC 27001:2017 is repeated, together with implementation guidance.

ISO/IEC 27003:2017 – Information technology – Security techniques – Information security management system implementation guidance

Synopsis: As the title suggests, this standard provides guidance on the establishment, implementation, operation, monitoring, review, maintenance and improvement of an ISMS. It includes a useful section on the preparation of business cases, and continues with the inputs required, outputs delivered and guidance on the various stages of the ISMS.

It contains a detailed checklist, together with information regarding roles and responsibilities, internal audit, policy structure and monitoring and measuring.

ISO/IEC 27004:2016 – Security techniques – Information security management – Measurement

Synopsis: Based on the PDCA model, this standard provides suggestions as to how the effectiveness of some information risk controls can be measured, and provides a number of templates for achieving this.

ISO/IEC 27005:2018 – Information security risk management (based on and incorporating ISO/IEC 13335-2:1997 – Management of information and communications technology security)

Synopsis: As the standard for information security risk management, ISO/IEC 27005 is an essential part of the ISO/IEC 27000 series, and encourages organisations to follow the well-established process to ensure that their information security management system reflects and manages the risks to the organisation's information.

ISO/IEC 27007:2020 – Information technology – Security techniques – Guidelines for information security management systems auditing

Synopsis: While ISO/IEC 27006 is aimed at certification bodies, this standard is aimed more at the internal audit community, and provides advice and guidance on the audit criteria for each area, any relevant standards that apply, the evidence the auditors should look for and practical guidance on how to approach the audit. It makes a number of references to ISO 19011:2018 – Guidelines for auditing management systems.

ISO/IEC TS 27008:2019 – Information technology – Security techniques – Guidelines for auditors on information security controls

Synopsis: Whereas the previous standard deals with the process for auditing an ISMS, this technical report examines the audit approach to verifying individual controls. It provides examples of: the requirement of a number of controls; some additional technical explanations of the reason for the control's use; the security implementation standard expected together with any additional technical supporting information; the practice guidance; evidence assumed; and the method to be adopted. It also provides guidance on initial information gathering that falls outside the IT environment.

ISO/IEC 27013:2015 – Information technology – Security techniques – Guidance on the integrated implementation of ISO/IEC 27001 and ISO/IEC 20000-1

Synopsis: This standard deals with the joint implementation of both ISO/IEC 27001 and ISO/IEC 20000-1 (Information technology – Service management – Part 1: Service management system requirements), which may already have been implemented by an organisation. The standard allows for three options: to implement ISO/IEC 27001 when ISO/IEC 20000-1 has already been implemented; to implement both ISO/IEC 27001 and ISO/IEC 20000-1 together; to align both the ISO/IEC 27001 and ISO/IEC 20000-1 implementations.

Although the two standards to which ISO/IEC 27013 refers have been updated since its publication in 2015, there is little content that is either omitted or changed.

ISO/IEC 27014:2020 – Information technology – Security techniques – Governance of information security

Synopsis: This standard deals with how the senior management of an organisation can manage the relationship between organisation governance, governance of information technology and governance of information security.

ISO/IEC 27016:2014 – Information technology – Security techniques – Information security management – Organisational economics

Synopsis: This is a technical report as opposed to a standard, and provides guidance on how organisations can use the impact assessment information to create business cases and then better compare the costs of protection with the potential costs of an incident.

Sector-specific guideline standards

ISO/IEC 27010:2015 – Information technology – Security techniques – Information security management for inter-sector and inter-organisational communications

Synopsis: ISO/IEC 27001 deals extremely well with the requirements for an ISMS within organisations, but does not specifically address the concept of information shared between organisations, either on a one-to-one or a many-to-many basis, especially in situations where both are cognisant of the sensitivity of the information being shared. This standard addresses this subject area, discusses the concept of anonymising information in order to protect its source and deals with the concept of trust in information sharing communities. Additionally, it is the first ISO standard to include the use of the TLP, which is discussed in [Chapter 2](#) of this book.

ISO/IEC 27011:2020 – Information technology – Security techniques – Information security management guidelines for telecommunications organisations (based on ISO/IEC 27002)

Synopsis: This is a copy-and-paste of the International Telecommunication Union (ITU) Standard X.1051 of 2008. It includes sections on: the organisation of information security; asset management; human resources security; communications and operations management; access control; information systems acquisition; development and maintenance; information security incident management; BCM; and compliance. The document finishes by describing some additional controls that are relevant to the telecommunications sector together with some additional implementation guidance.

Control-specific guideline standards

ISO/IEC 27031:2011 – Information technology – Security techniques – Guidelines for information and communication technology readiness for business continuity

Synopsis: This standard is one of the first to examine the world of DR, and addresses the concept of ICT readiness for business continuity. It includes references to five other relevant ISO/IEC standards: 18044, 27000, 27001, 27002 and 27005.

ISO/IEC 27032:2012 – Information technology – Security techniques – Guidelines for cybersecurity

Synopsis: This standard relates specifically to cybersecurity, which it defines as a 'complex environment resulting from the interaction of people, software and services on the Internet by means of technology devices and networks connected to it, which does not exist in any physical form'.

It would be easy to assume that this dealt solely with technology issues, but it does deal with all types of operational controls: people/procedural, physical/environmental and technical/logical.

ISO/IEC 27033 – Information technology – Security techniques – Network security

Synopsis: This group of six standards relate to network security. The first (2015) provides an overview and concepts; the second (2012) provides guidelines for the design and implementation of network security; the third (2010) addresses reference networking scenarios – threats, design techniques and control issues; the fourth (2014) provides guidance on securing communications between networks using security gateways; the fifth (2013) examines securing communications across networks using VPNs; and the sixth (2016) describes the threats, security requirements, security control and design techniques associated with wireless networks.

ISO/IEC 27034-1:2011 – Information technology – Security techniques – Application security – Overview and concepts

ISO/IEC 27034-2:2015 – Information technology – Security techniques – Application security – Organisation normative framework

ISO/IEC 27034-3:2018 – Information technology – Security techniques – Application security – Application security management process

ISO/IEC 27034 – Information technology – Security techniques – Application security – Application security validation (currently under review)

ISO/IEC 27034-5:2017 – Information technology – Security techniques – Application security – Protocols and application security control data structure

ISO/IEC 27034-6:2016 – Information technology – Security techniques – Application security – Case studies

ISO/IEC 27034-7:2018 – Information technology – Security techniques – Application security – Assurance prediction framework

Synopsis: This series of standards is intended to assist organisations in the integration of security into the processes used for managing their applications, and introduces definitions, concepts, principles and processes, but stops short of providing guidelines for secure software application development.

ISO/IEC 27035-1:2016 – Information technology – Security techniques – Information security incident management – 1 Principles of incident management

ISO/IEC 27035-2:2016 – Information technology – Security techniques – Information security incident management – 2 Guidelines to plan and prepare for incident response

ISO/IEC 27035-3:2016 Information technology – Security techniques – Information security incident management – Guidelines for ICT incident response operations

Synopsis: This standard is aimed at helping organisations to deal with information security incidents by introducing a structured approach, including: planning and preparation; detection and reporting; assessment and decision-making; response; and learning lessons. It provides some examples of information security incidents and their causes, lists a number of information security incident classification types, provides templates for recording and reporting incidents and addresses the legal and regulatory aspects of information security IM.

ISO/IEC 27036:2016 – Information technology – Security techniques – Information security for supplier relationships

Synopsis: This standard comes in four parts: overview and concepts; requirements; guidelines for information and communication technology supply chain security; and guidelines for security of cloud services.

ISO/IEC 27037:2016 – Information technology – Security techniques – Guidelines for identification, collection, acquisition and preservation of digital evidence

Synopsis: Although some organisations do not wish the fact that they have suffered an information security incident to be made public, many are now ensuring that they are able to collect evidence of such an incident in a manner that will enable a successful prosecution. This standard provides guidelines that will enable organisations to undertake this in a manner that is considered forensically acceptable.

ISO/IEC 27038:2016 – Information technology – Security techniques – Specification for digital redaction

Synopsis: Occasionally, organisations are required through legal and regulatory or other routes to release documents for public consumption; for example, under the Freedom of Information Act 2000. If these documents contain some sensitive information, but are otherwise able to be made public, then it is possible to redact the sensitive information.

This standard provides guidelines for the principles and processes associated with redaction, together with the characteristics of software redaction tools and the requirements for redaction testing.

ISO 22301:2019 – Societal security – Business continuity management systems – Requirements

Synopsis: This standard replaced BS 25999-2 in 2012, and is now regarded worldwide as the definitive business continuity standard. Since its publication, many organisations

that were previously accredited to BS 25999-2 have been successfully re-accredited to ISO 22301.

Other relevant standards

ISO/IEC 15504-2:2003 – Software engineering – Process assessment – Part 2: Performing an assessment

Synopsis: While ISO/IEC 15504-2 is aimed primarily at process assessment in software engineering, the principles apply equally to information risk management. The capability levels described are very similar to those in the Capability Maturity Model from the CMMI Institute, and also to those used in COBIT 5.

ISO 31000:2018 Risk management – Principles and guidelines

Synopsis: ISO 31000 is the international standard for risk management. By providing comprehensive principles and guidelines, this standard helps organisations with their risk analysis and risk assessments. ISO 31000 applies to most business activities, including planning, management operations and communication processes. While all organisations manage risk to some extent, this international standard's best-practice recommendations were developed to improve management techniques and ensure safety and security in the workplace at all times.

IEC 31010:2019 – Risk management – Risk assessment techniques

Synopsis: While not aimed specifically at information risk management, this standard goes into the generic risk management process in considerable detail. Much of the document is made up of the detailed techniques for undertaking the whole risk management programme, and each area covered is itemised and described in six different categories: overview; use; inputs; process; outputs; and strengths and weaknesses.

BS 7799-3:2017 – Information security management systems – Guidelines for information security risk management

Synopsis: This standard is one upon which this book was based, and follows the structure of context establishment, risk identification, risk analysis, risk evaluation and risk treatment, together with communication and consultation and monitoring and review.

BS 31100:2011 – Risk management – Code of practice and guidance for the implementation of BS ISO 31000

Synopsis: BS 31100:2011 outlines a risk management process that can be used and interpreted so that each group within an organisation works in a manner that increases consistency and communication across the business.

PAS 555:2013 – Cyber security risk – Governance and management – Specification

Synopsis: Threats to an organisation's cybersecurity present a critical challenge in terms of scale, complexity and impact – with business assets such as corporate and customer data, IP, brand and reputation at risk. It is crucial that an organisation understands and manages its exposure to cybersecurity threats.

The Risk Management Standard. Available from <https://www.theirm.org/what-we-do/what-is-enterprise-risk-management/irms-risk-management-standard/>.

Synopsis: The Risk Management Standard was originally published by the IRM, Airmic and Alarm in 2002.

Guide for conducting risk assessments – NIST Special Publication 800-30 Revision 1 (September 2012). Available from https://csrc.nist.gov/publications/nistpubs/800-30-rev1/sp800_30_r1.pdf.

Synopsis: This NIST standard deals with the fundamentals of risk management: the risk management process; risk assessment; key risk concepts; and the application of risk assessments. It then examines the risk management process in greater detail: preparing for the risk assessment; conducting the risk assessment; communicating and sharing risk assessment information; and maintaining the risk assessment.

In its appendices, the standard includes: references; a glossary of terms; acronyms; threat sources; threat events; vulnerabilities and predisposing conditions; likelihood of occurrence; impact; risk determination, informing risk response; risk assessment reports; and a summary of tasks.

Security and privacy controls for federal information systems and organisations – NIST Special Publication 800-53 Revision 5 (March 2020). Available from <https://doi.org/10.6028/NIST.SP.800-53r5>.

Synopsis: This NIST standard covers the fundamentals of security and privacy controls: assessments within the system development life cycle; the strategy for conducting security control assessments; building an effective assurance case; and assessment procedures. It continues by describing the process: preparing for security control assessments; developing security assessment plans; conducting security control assessments; and analysing security assessment report results.

In its appendixes, the standard includes: references; a glossary of terms; acronyms; assessment method descriptions; penetration testing; an assessment procedure catalogue; security assessment reports; and assessment cases.

All British and ISO standards may be purchased in PDF or hard copy formats from the BSI online shop <https://www.bsigroup.com/Shop> or by contacting BSI Customer Services, for hardcopies only, at tel: +44 (0) 20 8996 9001, email: cservices@bsigroup.com.

It is worth pointing out that by joining BSI, members can enjoy a 50 per cent discount on many British Standard (BS) products, and 10 per cent discount on ISO standards. Visit <https://shop.bsigroup.com/Navigate-by/Membership/> for more details.

APPENDIX I – DEFINITIONS, STANDARDS AND GLOSSARY OF TERMS

It is very helpful in any context, but especially in information risk management, that we have a common understanding of the terminology used. For example, people often refer to risk when they actually mean threat without perhaps realising that there is a distinct difference.

In this section, we provide definitions of all the key terms used in information risk management, most of which originate in ISO Guide 73:2009 – Risk Management – Vocabulary.

We shall then move on to cover the main national and international standards and good practice guidelines used in the management of information risk, and also identify where the reader can obtain them.


Risk management can be significantly more effective with clear and concise definitions:

You can't effectively and consistently manage what you can't measure, and you can't measure what you haven't defined.

The Open Group Standard, 2013

Let us begin though, by taking a very high-level view of the information risk management concepts and the relationships between them.

The diagram in [Figure I.1](#) illustrates the interrelationships between seven key areas of information risk management. We will expand on the individual concepts in due course, but for now it is worthwhile keeping this picture in mind as we work through the remaining chapters of this book.

Figure I.1 Concepts and relationships

DEFINITIONS AND GLOSSARY OF TERMS

Access control: 'the means to ensure that access to assets is authorised and restricted on business and security requirements' (ISO/IEC 27000:2018).

Asset: 'any item that has value to the organisation' (ISO/IEC 27000:2012 – oddly omitted from the more recent versions). Assets can be tangible, such as buildings, systems, people or information, or intangible, such as brand or reputation. IP is also an asset and results from the expression of an idea. IP might be a patent, trademark, copyright, design right, registered design, technical or commercial information. Bizarrely, although IP can be owned, bought and sold, information per se is not considered 'property' in the strictest sense of UK law!

Attack: 'attempt to destroy, expose, alter, disable, steal or gain unauthorised access or to make unauthorised use of an asset' (ISO/IEC 27000:2018). An attack can be a simple event, such as an attempt to break into a computer system, or a more complex event such as a DDoS attack in which multiple systems mount an attack on an information asset. Attacks differ slightly from threats and hazards in that attacks are something that actually happen, whereas threats and hazards only have the potential to cause harm. An attacker is therefore someone who deliberately sets out to cause harm. See also Exploit.

Attribute: 'the property or characteristic of an object that can be distinguished quantitatively or qualitatively by human or automated means' (ISO/IEC 27000:2018).

Audit: 'the systematic, independent and documented process for obtaining audit evidence and evaluating it objectively to determine the extent to which the audit criteria are fulfilled' (ISO/IEC 27000:2018).

Authentication: 'the provision of assurance that a claimed characteristic of an entity is correct' (ISO/IEC 27000:2018).

Availability: 'property of being accessible and usable upon demand by an authorised entity' (ISO/IEC 27000:2018).

Business continuity (BC): 'the capability of the organisation to continue delivery of products and services at acceptable predefined levels following a disruptive incident' (ISO 22300:2018).

Business impact analysis (BIA): 'the process of analysing activities and the effect that a business disruption can have upon them' (ISO 22300:2018).

Communication and consultation: 'the continual and iterative processes that an organisation conducts to provide, share or obtain information, and to engage in dialogue with stakeholders regarding the management of risk' (ISO Guide 73:2009). Such dialogue will be with internal stakeholders, such as senior managers and staff involved in managing information, and also external stakeholders such as outsourcing organisations.

Communication and consultation must happen continuously throughout the life cycle (and beyond) of an information risk management project or programme of work, and must never be viewed as a one-off exercise and it should be remembered that communication and consultation is a two-way process.

Confidentiality: 'the property that information is not made available or disclosed to unauthorised individuals, entities or processes' (ISO/IEC 27000:2018).

Consequence or impact: 'an outcome of an event affecting objectives' (ISO Guide 73:2009). The two terms are both widely used and are completely interchangeable. Consequences and impacts may be direct, for example the loss of a building due to a fire, or indirect, such as the fine imposed for a breach of the Data Protection Act.

Also, there are primary impacts, such as the loss of revenue when a system fails, and secondary impacts, such as the overtime payments to staff for working extra hours to repair or replace such a system.

As shown in this book, consequences and impact may be described in a qualitative or a quantitative manner. Qualitative descriptions are generally in the form of 'trivial', 'minor', 'major', 'severe' or 'critical', but unless they are based on some form of numerical value, do not really provide an insightful assessment of the grim reality.

Quantitative descriptions are much easier to understand and provide a firm basis for comparison and assessment, but are generally much harder to predict with any accuracy unless very detailed analysis is carried out, which can be time consuming.

In practice, the best balance can usually be obtained by providing a quantitative rating for a qualitative term; for example, between £1 million and £10 million represents 'severe', which allows a greater degree of subjectivity while anchoring the assessment in numeric terms. This is sometimes referred to as a semi-quantitative measure.

Context establishment: 'defining the external and internal parameters to be taken into account when managing risk, and setting the scope and risk criteria for the risk management policy' (ISO Guide 73:2009). This will be discussed in greater detail in [Chapter 3](#).

Control: 'a measure that is modifying risk' (ISO Guide 73:2009). Controls can be strategic, tactical or operational. Strategic controls are very high level, such as risk avoidance, transfer, reduction and acceptance; tactical controls determine a general course of action, such as detective, preventative, corrective and directive; and operational controls determine the actual treatment, such as technical or logical, procedural or people, and physical or environmental. However, controls may also be used to monitor processes, ensuring predictability without actually modifying them.

Control objective: 'a statement describing what is to be achieved as a result of implementing controls' (ISO/IEC 27000:2018).

Cybersecurity: Refers specifically to information security as applied to computers, tablet computers, smartphones, computer networks (both public and private) and the wider internet. In this respect it is slightly different from the wider area of information security, which includes non-electronic information as well. Cybersecurity is sometimes also referred to as computer security or IT security.

Data: A collection of values assigned to base measures, derived measures and/or indicators.

Disaster recovery (DR): A coordinated activity to enable the recovery of ICT systems and networks due to a disruption.

Disruption: This term is generally applied to events or incidents that interfere with normal business operations and have a detrimental impact on information or information processing.

Effectiveness: 'the extent to which planned activities are realised and planned results achieved' (ISO/IEC 27000:2018).

Estimation: It is almost impossible to predict either the impact or the likelihood of a threat arising with any degree of accuracy or certainty, so almost all risk assessments are carried out on the basis of estimation. Estimates can be refined and improved over time, and with the hindsight of real events they may even become quite accurate, but initially they will always be little more than an educated guess.

Event: 'the occurrence or change of a particular set of circumstances' (ISO Guide 73:2009). Sometimes these are referred to as incidents and, while there are similarities, there needs to be a differentiation between the various types of change of circumstances.

In terms of information risk, 'events' can vary considerably in scale and severity from so-called 'glitches', lasting perhaps a fraction of a second, through to major incidents that can affect the organisation for weeks or months. In order to place events in a clearer descriptive context, examples are provided in [Chapter 4](#).

Exploit: An exploit is a particular form of attack, in which a tried and tested method of causing impact is followed with some rigour. Exploits are similar in nature to processes, but whereas processes are generally benign, exploits are almost always harmful.

Exposure: 'the extent to which an organisation and/or stakeholder is subject to an event' (ISO Guide 73:2009).

External context: 'the external environment in which the organisation seeks to achieve its objectives' (ISO Guide 73:2009). Again, as the name suggests, the external context takes in factors outside the bounds of the organisation. This will be discussed in greater detail in [Chapter 3](#).

Frequency: 'the number of events or outcomes per defined unit of time. Frequency can be applied to past events or to potential future events, where it can be used as a measure of likelihood or probability' (ISO Guide 73:2009). Most real-world events are not precisely regular in occurrence, the tides and phases of the moon being obvious exceptions, and so any statements of frequency are only really estimates and cannot be relied upon for accuracy.

Hazard: 'a source of potential harm' (ISO Guide 73:2009). Hazards are generally seen as natural (as opposed to man-made) events, such as flooding, hurricanes or ice storms. See also Threats.

Horizon scanning: A procedure that involves the systematic observation and monitoring of various key drivers of change at the margins of current thinking and planning.

Impact: 'an outcome of an event affecting objectives' (ISO Guide 73:2009).

Information: An organised and formatted collection of data.

Information assurance: The process of ensuring that data are not lost when critical events or incidents occur. It is generally associated with computer, cyber or IT security rather than the somewhat wider meaning of 'information security'.

Information security: The practice of protecting information from unauthorised access, use, disclosure, disruption, modification or destruction. Information security encompasses both physical and electronic information.

Information security event: 'an identified occurrence of a system, service or network state indicating a possible breach of information security policy or failure of safeguards, or a previously unknown situation that may be security relevant' (ISO/IEC 27000:2018).

Information security incident: 'indicated by a single or a series of unwanted or unexpected information security events that have a significant probability of compromising business operations and threatening information security' (ISO/IEC 27000:2018).

Inherent risk: The risk that an activity would pose if no controls or other mitigating factors were in place (the gross risk or risk before controls).

Integrity: 'property of protecting the accuracy and completeness of assets' (ISO/IEC 27000:2018).

Internal context: 'the internal environment in which the organisation seeks to achieve its objectives' (ISO Guide 73:2009). As the name suggests, the internal context is that which exists within the organisation itself. This will be discussed in greater detail in [Chapter 3](#).

Level of risk: 'the magnitude of a risk expressed in terms of the combination of consequences and their likelihood' (ISO/IEC 27000:2018).

Likelihood: 'the chance of something happening' (ISO Guide 73:2009). The terms likelihood and probability are often used interchangeably, but likelihood is a rather general or qualitative term denoting a degree of uncertainty, whereas the quantitative term 'probability' has a more statistical underpinning. The term 'possibility' is generally not used, since many things are possible, and the term gives no indication whether or not the event is actually likely to take place.

Monitoring: 'determining the status of a system, a process or an activity' (ISO 22300:2018).

Non-repudiation: 'the ability to prove the occurrence of a claimed event or action and its originating entities, in order to resolve disputes about the occurrence or non-occurrence of the event or action and involvement of entities in the event' (ISO/IEC 27000:2018).

Objective: 'a result to be achieved' (ISO/IEC 27000:2018).

Organisation: 'a person or group of people that has its own functions with responsibilities, authorities and relationships to achieve its objectives' (ISO/IEC 27000:2018).

Outsource: 'make an arrangement where an external organisation performs part of an organisation's function or process' (ISO/IEC 27000:2018).

Policy: 'the intentions of an organisation as formally expressed by its top management' (ISO/IEC 27000:2018).

Probability: 'the measure of the chance of occurrence expressed as a number between zero and one, where zero is impossibility and one is absolute certainty' (ISO Guide 73:2009). Probability is often expressed as a percentage, and being a quantitative term, is able to express the chance of something happening with a greater degree of accuracy. See also Likelihood.

Processes and procedures: Many organisations do not think of processes as being information assets, but as they are documented in some way and often refer to the use

or production of information, they can be considered as intangible assets. Processes detail how to go about achieving a goal or objective. Procedures, which are a subset of processes, explain how to conduct the individual steps within processes, and therefore take on the same status as intangible assets.

Qualitative risk assessments: These are subjective in nature, and are generally expressed in verbal terms such as 'high', 'medium' and 'low'. While in common use, this is not always an ideal state of affairs, as it may render risk assessments unreliable. This is discussed in greater detail in [Chapter 4](#).

Quantitative risk assessment: These are objective in nature and are generally expressed in numerical terms, such as financial values, percentages and so on. While these provide a more accurate measurement of risk, they are usually more time consuming to undertake. They are discussed in greater detail in [Chapter 4](#).

Requirement: 'the need or expectation that is stated, generally implied or obligatory' (ISO/IEC 27000:2018).

Residual risk: 'the risk remaining after risk treatment' (ISO Guide 73:2009). Once all other risk treatment options have been explored, it is often the case that some (usually small) risk remains. It is normal to accept or tolerate this, since further treatment might either have no effect, or might be prohibitively expensive. Because residual risks are often very small, they are occasionally (incorrectly) overlooked.

Resilience: 'the adaptive capacity of an organisation in a complex and changing environment' (ISO Guide 73:2009). Although this definition refers to organisations rather than to information assets, the definition holds true, in that where an information asset is properly protected, it is able to resist certain threats. However, to make an information asset fully resilient may be a very complex task and require several different methods of protection.

Review: 'an activity undertaken to determine the suitability, adequacy and effectiveness of the subject matter to achieve established objectives' (ISO/IEC 27000:2018).

Risk: 'the effect of uncertainty on objectives' (ISO Guide 73:2009). Risk is the product of consequence or impact and likelihood or probability, and is not the same as a threat or hazard. In the context of information risk management, risk is usually taken to have negative connotations. In the wider context of risk however, it may also be seen in a positive light and may be referred to as 'opportunity'.

Risk acceptance or risk tolerance: 'the informed decision to take a particular risk' (ISO Guide 73:2009). Risk acceptance or tolerance is the final choice in risk treatment once all other possible avenues have been explored. This is not the same as ignoring risks – something that should never be done!

Risk aggregation: 'the combination of a number of risks into one risk to develop a more complete understanding of the overall risk' (ISO Guide 73:2009). Where a number of risks exist in a certain area, it may be possible to treat them all with one or more controls rather than treating them individually. Therefore, for the purposes of risk management, they can be grouped together or aggregated in order to save time and effort.

Risk analysis: 'the process to comprehend the nature of risk and to determine the level of risk' (ISO Guide 73:2009). This is the part where we combine the impact and the likelihood (or probability) to calculate the level of risk and to plot it onto a risk matrix, which allows us to compare risks for their severity and to decide which are in greatest need of treatment.

Risk appetite: 'the amount and type of risk that an organisation is willing to pursue or retain' (ISO Guide 73:2009). Organisations will have differing levels of risk appetite for different types of information; and different types of organisation will have vastly differing levels of risk appetite, depending on their sector.

Risk assessment: 'the overall process of risk identification, risk analysis and risk evaluation' (ISO Guide 73:2009). It includes identification of the information assets and their owners; impact assessment; threat and vulnerability identification; likelihood assessment; risk analysis; production of the risk matrix; and finally risk evaluation.

Risk avoidance or risk termination: 'an informed decision not to be involved in, or to withdraw from, an activity in order not to be exposed to a particular risk' (ISO Guide 73:2009). This is one of the four strategic options for risk treatment. Avoiding the risk should normally remove the risk completely, but may leave the organisation with other challenges.

Risk criteria: 'the terms of reference against which the significance of a risk is evaluated' (ISO Guide 73:2009). Risk criteria will include such things as impact, likelihood, proximity and risk appetite.

Risk evaluation: 'the process of comparing the results of risk analysis with risk criteria to determine whether the risk and/or its magnitude is acceptable or tolerable' (ISO Guide 73:2009). This is the final stage in the risk assessment process, in which all risks plotted onto the risk matrix are evaluated against a set of criteria in order to decide which should receive the highest priority for treatment.

Risk identification: 'the process of finding, recognising and describing risks' (ISO Guide 73:2009). Risk identification includes the identification risk sources, events, their causes and the possible consequences to the information assets.

Risk management: 'coordinated activities to direct and control an organisation with regard to risk' (ISO Guide 73). Risk management is the identification, assessment and prioritisation of risks followed by coordinated and economical application of resources to minimise, monitor and control the probability and/or impact of unfortunate events or to maximise the realisation of opportunities.

Risk matrix: A graphical representation of impact versus likelihood used to assist in the prioritisation of risks.

Risk modification or risk reduction: The process of treating risk by the use of controls to reduce either the consequence/impact or the likelihood/probability. Sometimes the term 'risk treatment' is used in this context, but risk treatment is really a generic term for all four kinds of strategic control. Strangely, ISO Guide 73 does not attempt to define risk modification or reduction, although it does refer to it under the definition of 'control'.

Risk monitoring: 'the continual checking, supervising, critically observing or determining the status in order to identify change from the performance level required or expected' (ISO Guide 73:2009). This is an ongoing process to ensure that risks that change over time – whether for the better or the worse – are reviewed and that appropriate action is taken.

Risk owner: 'a person or entity with the accountability and authority to manage a risk' (ISO Guide 73:2009).

Risk proximity: How far away in time will the risk occur (if it materialises). It can also mean when the risk will occur. While there does not appear to be a standards definition for risk proximity, it remains a vital element of the risk assessment process, since those risks that could manifest themselves sooner will probably require attention before those that are further away in time. Risk proximity is one of the criteria against which risks are evaluated.

Risk reduction: The process of treating risk by the use of controls to reduce either the consequence/impact or the likelihood/probability.

Risk register: 'a record of information about identified risks' (ISO Guide 73:2009). Simple risk registers are often maintained as a spreadsheet, while more complex registers may use a proprietary software package, capable not only of recording the information but also of carrying out some analysis or evaluation.

Risk reporting: 'a form of communication intended to inform particular internal or external stakeholders by providing information regarding the current state of risk and its management' (ISO Guide 73:2009).

Risk review: 'the activity undertaken to determine the suitability, adequacy and effectiveness of the subject matter to achieve established objectives' (ISO Guide 73:2009). Risk reviews capture not only risks and their treatments, but also the whole process by which risk management is undertaken, the status of information assets and the organisation's risk appetite.

Risk termination: 'an informed decision not to be involved in, or to withdraw from, an activity in order not to be exposed to a particular risk' (ISO Guide 73:2009).

Risk tolerance: 'an organisation or stakeholder's readiness to bear the risk after risk treatment in order to achieve its objectives' (ISO Guide 73:2009). Risk tolerance is sometimes also viewed as being the same as risk acceptance. The difference is that risk acceptance takes place when no other form of risk treatment is suitable, whereas risk tolerance takes place after other forms of risk treatment have taken place, and there is some residual risk.

Risk transference or risk sharing: 'a form of risk treatment involving the agreed distribution of risk with other parties' (ISO Guide 73:2009). One of the risk treatment options is to transfer the risk to or to share it with a third party. Transferring or sharing the risk, however, does not change ownership of the risk, which remains with the organisation itself regardless of who else shares the risk.

Risk treatment: ‘the process to modify risk’ (ISO Guide 73:2009). While this may be technically correct, treatment may alternatively involve risk transference or sharing; risk avoidance or termination, or risk modification or reduction, which are all methods of treating risk.

Scale: ‘an ordered set of values, continuous or discrete, or a set of categories to which the attribute is mapped’ (ISO/IEC 27000:2014).

Stakeholder: ‘a person or organisation that can affect, be affected by, or perceive themselves to be affected by a decision or activity’ (ISO Guide 73:2009). Stakeholders may be people within or outside the organisation, including customers, suppliers or governments.

Threat: ‘potential cause of an unwanted incident, which can result in harm to a system or organisation’ (ISO/IEC 27000:2018). ISO Guide 73:2009 defines ‘hazards’, but does not refer to threats. While hazards are generally viewed as natural events, threats are usually man-made, whether accidental or deliberate.

Threat sources and threat actors: A threat source is a person or organisation that wishes to benefit from attacking an information asset. A threat actor is a person or organisation that actually mounts the attack. Threat sources often coerce threat actors to attack information assets on their behalf.

Uncertainty: This is the state, even partial, of deficiency of information related to, understanding or knowledge of, an event, its consequence or likelihood. While not an actual term defined in ISO Guide 73:2009, uncertainty is explained in a note below the definition of risk. Uncertainty goes hand in hand with estimation, meaning that many of our assessments will be subject to a greater or lesser degree of uncertainty. In some cases, uncertainty increases as a possible event’s proximity decreases.

Validation: ‘confirmation, through the provision of objective evidence, that the requirements for a specific intended use or application have been fulfilled’ (ISO/IEC 27000:2014).

Vulnerability: ‘the intrinsic properties of something resulting in susceptibility to a risk source that can lead to an event with a consequence’ (ISO Guide 73:2009). Vulnerabilities or weaknesses in or surrounding an asset leave it open to attack from a threat or hazard. Vulnerabilities come in two flavours: intrinsic vulnerabilities, which are something inherent in the very nature of an information asset, such as the ease of erasing information from magnetic media (whether accidental or deliberate); and extrinsic vulnerabilities, which those that are poorly applied, such as software that is out of date due to a lack of patching.

INFORMATION RISK MANAGEMENT STANDARDS

There are a number of useful standards and guidelines available to risk management practitioners. Unfortunately, the BSI and ISO standards can only be purchased, although members of BSI enjoy a discount on many standards. The National Institute for

Standards and Technology (NIST) and Committee on National Security Systems (CNSS) standards are free to download.

The lists below include only the most relevant standards. For a fuller list of all related standards, please see [Appendix H](#).

British Standards Institution (BSI)

<https://bsol.bsigroup.com>

BS 7799-3:2017 – Information security management systems – Guidelines for information security risk management

BS 31100:2011 – Risk management – Code of practice and guidance for the implementation of BS ISO 31000:2018

International Organization of Standardization (ISO)

www.iso.org/iso/home/standards.htm

ISO Guide 73:2009 – Risk management – Vocabulary

ISO/IEC TR 13335-3:1998 – Information technology – Guidelines for the management of IT security – Part 3: Techniques for the management of IT security

ISO/IEC 27000:2020 – Information technology – Security techniques – Information security management systems – Overview and vocabulary

ISO/IEC 27001:2017 – Information technology – Security techniques – Information security management systems – Requirements

ISO/IEC 27002:2017 – Information technology – Security techniques – Code of practice for information security management.

ISO/IEC 27005:2018 – Information technology – Security techniques – Information security risk management

ISO 31000:2018 – Risk management – Principles and guidelines

IEC 31010:2019 – Risk management – Risk assessment techniques

US National Institute of Standards and Technology

<https://csrc.nist.gov/publications/PubsSPs.html>

NIST SP 800-30 Revision 1, September 2012 – Guide for Conducting Risk Assessments.

NIST SP 800-53 Revision 5 Recommended Security Controls for Federal Information Systems and Organizations (2020)

UK Institute of Risk Management

www.theirm.org/

The Risk Management Standard (2002)

Information Security Forum (ISF)

www.securityforum.org

The Standard of Good Practice for Information Security (2020)

Permission to reproduce extracts from ISO/IEC 27000:2018, ISO 22300:2021, and ISO Guide 73:2009 has been granted by BSI. ISO standards can be obtained in PDF or hard copy formats from the BSI online shop: www.bsigroup.com/Shop or by contacting BSI Customer Services for hardcopies only: Tel: +44 (0) 20 8996 9001, Email: cservices@bsigroup.com.

INDEX

Page numbers in italics refer to figures or tables.

- 10 Steps To Cyber Security 18, 187, 189–91, 196
- access control 62, 64, 143, 145–6, 148–9, 151–2, 156, 159–60, 164, 166, 174, 199, 205
- accreditation 13, 23, 29, 97, 195
- AI (artificial intelligence) 10–11
- application clusters 103
- application resilience 103
- assessment
 - impact 18, 38, 39, 42–53, 54, 72, 73, 91, 105, 111, 125, 167, 168, 179, 182, 199
 - of likelihood 71–4
 - of risk 6, 15, 17, 20, 23, 31, 37, 42, 66, 71, 74, 77, 79, 81, 88, 89, 93, 96, 106, 111, 117, 168, 172–3, 177–80, 202–3, 207, 210–12, 214
 - threat 54–60, 61, 62, 66, 117, 183
 - vulnerability 60–6
- BCI (Business Continuity Institute) Good Practice Guidelines 2018 96, 97, 98
- BCI life cycle 96, 97
- BCM (business continuity management) 8, 13, 27, 41, 89, 95–9, 118, 163, 193, 194, 199, 201
- BCP (business continuity plan) 96, 97–9
- business cases 17, 80, 90, 91, 92–3, 104, 105, 197, 199
- business continuity 21, 41, 91, 95–9, 104, 118, 147, 199, 201, 206
- BYOD (bring your own device) policies 63, 65, 144, 150–1
- CCP (Certified Cyber Professional) scheme 17, 113–15, 116, 194–5
- Centre for Internet Security Controls Version 8 87–8, 156–8, 193
- chain of consequence 42, 43, 43, 146
- change failures 56, 58, 131, 137
- CII Sec (Chartered Institute of Information Security) 113, 114, 115, 116–19, 194, 195
- cloud-based storage 102–3
- clustered file systems 103
- CMM (Capability Maturity Model) 13, 193
- COBIT 5® 14, 202
- cold standby platforms 100
- communication 9, 29, 35, 46, 105, 107–9, 118, 129, 161, 176, 177, 206
- communications and operations management failures 65, 149–51
- computer cluster services 103
- confidentiality 5, 6, 8, 10, 20–2, 24, 32, 37, 45, 55, 72, 122, 128, 130, 145, 161, 169, 175, 179, 206
- consultation 16, 17, 29, 35, 39, 105, 109–10, 176, 177, 202, 206
- controls
 - access 62, 64, 143, 145–6, 148, 149, 151, 152, 156, 159, 160, 164, 166, 174, 199, 205
 - Centre for Internet Security Controls Version 8 87–8, 156–8, 193
 - corrective 70, 82, 86, 94, 155, 185
 - critical 87–90
 - detective 70, 82, 85, 86, 94, 155, 185
 - directive 70, 82, 86, 94, 155, 185
 - existing 16, 18, 66–9, 70, 71, 81, 124, 126, 172, 185
 - ISO/IEC 27001:2017 158–63
 - NIST Special Publication 800–53 Revision 5 163–70
 - operational controls 68, 86, 89, 97, 155–6, 158, 200, 207
 - physical 18, 68, 82, 87, 155
 - preventative 18, 70, 82, 85–6, 94, 155, 185
 - procedural 18, 22, 68, 82, 87, 94, 155
 - risk 154–69, 198
 - strategic 18, 68, 85, 154–6, 207
 - tactical risk management 85–6
 - technical 18, 68, 82, 87, 123, 155, 156, 197
 - testing 117
- CORAS 18, 171–2, 196
- corrective controls 70, 82, 86, 94, 155, 185
- DAS (Direct Attached Storage) 102, 103
- data resilience 102–3
- decision-making 15, 21, 33, 91, 92, 93–5, 107, 137, 201
- detective controls 70, 82, 85, 86, 94, 155, 185
- detrimental situations (events/incidents) 40–1, 40
- directive controls 70, 82, 86, 94, 155, 185
- disaster recovery 24, 95, 99–104, 99, 118, 147, 173, 207
- ENISA (European Network and Information Security Agency) 24, 193

environmental threats 55, 57–8, 61, 74, 87, 125, 133–5, 160, 183
errors and failures 58, 135–7
events/incidents 40–1, 40
existing controls 16, 18, 66–9, 70, 71, 81, 124, 126, 172, 185

failover cluster services 103
FAIR (Factor Analysis of Information Risk) 18, 172–3, 196

GCHQ (Government Communications Headquarters) 2, 189, 195

glossary of terms 204–13

hacking 17, 55, 56, 57, 58, 61, 72, 74, 125, 130–3, 137, 138, 142, 145, 149, 172, 183

HMG Cyber Essentials Scheme 18, 187–9, 195

HMG Security Policy Framework 17, 115, 120, 195

hot standby/high-availability platforms 100–1

identification process (risk) 37–8, 40, 74

impact 7, 8, 12, 13, 17, 31, 33, 34, 35, 37, 38, 60, 124, 208

adverse 98

analysis 31, 40, 165, 206

assessment 18, 38, 39, 42–53, 54, 72, 73, 91, 105, 111, 125, 167, 168, 179, 182, 199

availability 72

BIA (business impact assessment/analysis) 16, 206

confidentiality 72

or consequence 125, 127

direct 44, 52, 124, 125, 182

financial 21, 23, 24, 30, 45–6, 49, 51, 126, 128

indirect 44, 52, 124, 125, 182

integrity 72

legal and regulatory 45, 46, 49, 129

operational 45, 48, 50, 126, 128

people 47, 49, 129

primary 44, 52, 124, 125, 182, 206

reputational 45, 46–7, 49, 129

scales 32, 48

secondary 44, 47, 52, 124, 125, 129, 182, 206

time dependency 51

types 44, 44–7, 179

typical 125, 126, 127

information

assets 5, 6, 8, 12, 16, 21–6, 29, 31, 33, 35, 37–8, 39, 42, 48–51, 53–5, 60, 62, 66, 68–9, 71–2, 77, 79–81, 93, 110, 111, 122–6, 143, 152, 155, 159, 174

assigned attributes 3

assurance 20, 22, 113, 116, 194, 208

classification 5, 16, 22–5, 26, 27, 159, 193

definitions 1, 2

financial 5, 10

governance 20, 22, 116

immutable attributes 3

life cycle 4, 4

owner 25, 38, 72

personal 5, 6, 38, 44, 125, 133, 137, 140, 150

risk controls 154–69, 198

security 5, 10, 11, 20–7, 29, 30, 34, 54, 60, 66, 69, 88–9, 91–2, 97, 111, 113–23, 158–9, 161–3, 166–7, 173, 175–7, 187, 193–9, 201–2, 207–9, 214–15

sensitive 3, 7, 23, 57, 122, 138, 141, 201

system 59, 84, 113, 116, 119, 140, 161, 194, 199, 203, 214

type 3, 5, 38, 39, 61, 67, 81, 183, 184, 211

information risk management 174, 187

governance 30–1

method 176

need for 1–19

NIST publication 90, 163, 177

process 14–16, 15

programme 2, 16, 18, 28–36, 55, 62, 68, 76, 80, 94–5, 98, 105–12, 173, 187

regime 189–91

software 179

standards 89, 158, 176–7, 202, 213–15

technical controls 18, 68, 82, 87, 123, 155, 156, 197

threat assessment 54–60, 61, 62, 66, 117, 183

information risk manager 35, 91–2, 95, 105–10

IOT (Internet of Things) 1, 9–10

ISO/IEC 27001:2017 controls 88–9, 158–63

likelihood (assessment of) 71–4

likelihood or probability 125

load balancing 99, 103

malicious intrusion (hacking)

17, 55, 56, 57, 58, 61, 72, 74, 125, 130–3, 137, 138, 142, 145, 149, 172, 183

malware 17, 55, 58–60, 65, 68, 72, 74, 133, 137, 139, 140–2, 150, 157, 161, 188, 190, 191

MAO (maximum acceptable outage) 98

MBCO (minimum business continuity objective) 98

Minimum Cyber Security Standard 121

misuse and abuse 58–9, 138–9

monitoring 16, 17, 27, 28, 29, 33, 35–6, 63, 65, 66, 81, 93, 94, 102, 105, 110–11, 112, 117, 149, 152, 157, 174, 191, 197, 209, 212

MTDL (maximum tolerable data loss) 98

MTPD (maximum tolerable period of disruption) 49, 98

NAS (Network Attached Storage) 102, 103

National Security Strategy 120–1

NCSC (National Cyber Security Centre) 113–19, 187, 189, 194

NCSC Certified Professional Scheme 17, 19, 113–19, 192, 194–5

NIST Special Publication 800-30 90, 163, 177–80, 203, 214

NIST Special Publication 800-53 Revision 5 controls 87, 89–90, 156, 163–70, 203, 214

NSA (National Security Agency) 2

OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation) 18, 173–4, 196

OCTAVE Allegro 174–5, 196

OCTAVE-S 174, 196

operational controls 68, 86–7, 89, 97, 155–6, 158, 200, 207

operational impact 45, 48, 50, 126, 128

PDCA (Plan-Do-Check-Act) 26, 26–7, 28, 110

people-related security failures 66, 151–3

physical and environmental failures 64–5, 147–9

physical controls 18, 68, 82, 87, 155

physical threats 59, 139–40

platform disaster recovery 100–2

poor procedures 64, 146–7

PP (professional practice) 96

preventative controls 18, 70, 82, 85–6, 94, 155, 185

primary UK legislation 19, 192–3

privacy and secrecy 20–2

procedural controls 18, 22, 68, 82, 87, 94, 155

programme manager skills 105–7

qualitative and quantitative assessments 16, 47–8, 53, 72–4, 125, 206, 210

RAID (Redundant Array of Inexpensive Disks) 102

remote working 1, 11–12, 190, 191

reputational impact 45, 46–7, 49, 129

risk

- acceptance 32, 34, 79, 85, 93, 94, 110, 112, 210, 212
- analysis 16, 71, 74–6, 171, 172, 173, 202, 211
- appetite 8, 9, 29, 31–2, 33, 77, 79, 81, 84, 108, 110, 111, 112, 173, 211, 212
- assessment of 6, 15, 17, 20, 23, 31, 37, 42, 66, 71, 74, 77, 79, 81, 88, 89, 93, 96, 106, 111, 117, 168, 172–3, 177–80, 202–3, 207, 210–12, 214
- avoidance 32, 33, 79, 82, 93, 94, 110, 207, 211, 213
- controls 154–69, 198
- criteria 28, 77, 207, 211
- environment 8, 38
- evaluation 16, 76–80, 93, 112, 171, 202, 211
- event 34, 40
- identification 16, 37–53, 71, 74, 171, 177, 179, 202, 211
- intelligence 30
- matrix 16, 74–6, 75, 76, 211
- mitigation 113, 175
- monitoring 112, 212
- operational controls 17, 86–7, 155–6, 175
- options 82–5
- policy 30
- reduction 32, 33, 79, 84, 92, 93, 94, 97, 162, 110, 211, 212
- register 68, 76–7, 78, 80, 111, 112, 179, 186, 212
- reporting 17, 91–104, 147, 212
- residual 15, 29, 33, 34, 68, 81, 84, 85, 108, 111, 112, 154, 155, 210, 212
- review 110, 111–12, 212
- strategic options 81, 82–5, 154–5
- tactical controls 85–6, 155
- tolerance 112, 210, 212
- transfer 32, 33–4, 79, 83–4, 92, 93, 94, 110, 155, 212, 213
- treatment 15, 16, 17, 31–4, 50, 69, 79, 80, 81–90, 91, 93–5, 96, 110–12, 137, 156, 158, 163, 173, 177, 179, 202, 210–13

RPO (recovery point objective) 98

RTO (recovery time objective) 98

SABSA (Sherwood Applied Business Security Architecture) 18, 171, 175–6, 196

SAN (Storage Area Networks) 102, 103

security classifications 17, 120, 121–2, 161, 195

SFIA (Skills Framework for the Information Age) 17, 114–16, 119, 195

site recovery 99, 103

social engineering 17, 55, 56, 58, 74, 125, 133, 137–8

staff training 34–5

standards 14, 27, 28, 31, 95–6, 113, 114, 196–203, 213–15

strategic controls 18, 68, 85, 154–6, 207

strategic impact assessment 49–50

tactical impact assessment 50

technical controls 18, 68, 82, 87, 123, 155, 156, 197

templates

- existing controls assessment 185
- impact assessment 182
- risk register 186
- threat assessment 183
- vulnerability assessment 184

testing 137, 146, 147, 161, 194, 197

- application 99, 151
- contingency plan 165
- controls 117
- cybersecurity 188
- developer 168
- failover 104
- functionality 117
- incident response 166
- penetration 66, 88, 114, 117, 158, 165, 203
- redaction 203
- regression 135
- security 66, 161, 162
- system acceptance 161, 162

threat

- assessment 54–60, 61, 62, 66, 117, 183
- environmental 55, 57–8, 74, 87, 125, 133–5, 160
- errors and failures 17, 55, 56, 58, 72, 135–7
- hacking 17, 55, 56, 57, 58, 61, 72, 74, 125, 130–3, 137, 138, 142, 145, 149, 172, 183
- malware 17, 55, 58–60, 65, 68, 72, 74, 133, 137, 139, 140–2, 150, 157, 161, 188, 190, 191
- misuse and abuse 58–9, 138–9
- physical 17, 55, 59, 72, 74, 139–40
- social engineering 58, 137–8
- typical threats 56, 125–6, 130–42

time dependency 51

'Traffic Light Protocol' 24–5, 108, 193

typical impacts 125, 126, 127

typical threats 56, 130–42

typical vulnerabilities 143–53

UK Cyber Security Strategy
2016–2021 121

vulnerability

access control 143, 145–6
assessments 16, 18, 54, 55,
60–6, 67, 69, 91, 184

communications and
operations management 65,
149–51

people-related security
failures 66, 151–3

physical and environmental
64–5, 147–9

poor procedures 64, 146–7
typical vulnerabilities 63,
143–53

warm standby platforms 100

INFORMATION RISK MANAGEMENT

A practitioner's guide - Second edition

David Sutton

Organisations rely on information for their day-to-day operations, so the loss or unavailability of information can mean the difference between success and ruin. Information risk management (IRM) is about identifying, assessing, prioritising and treating risks to keep information secure and available.

This book is a practical guide to understanding the principles of IRM and developing a strategic approach to an IRM programme. It's ideal for those studying for the BCS Practitioner Certificate in Information Risk Management; this new edition reflects changes to the syllabus and to the wider discipline.

- Understand the need for information risk management
- Discover the tools and techniques required to conduct a successful IRM programme
- Explore typical threats, hazards and vulnerabilities
- The only textbook for the BCS Practitioner Certificate in Information Risk Management

ABOUT THE AUTHOR

David Sutton's career in IT spans more than 50 years and includes voice and data networking, information security and critical information infrastructure protection. He has been a member of the BCS Professional Certification Information Security Panel since 2005 and has delivered lectures on information risk management and business continuity at the Royal Holloway University of London. He is the author of BCS book *Cyber Security* and co-author of the books *Information Security Management Principles* and *Data Governance*.

You might also be interested in:



This book is essential reading for any risk management practitioner. The author's many years of practical experience in the subject shine through, it is clearly written and easy to follow...Highly recommended, this will be on my bookshelf.

David Alexander, Information Security Group, Royal Holloway, University of London

Information risk management is an integral part of every business and the author presents its lifecycle in an easy-to-follow and well organised format with real-life examples, tools and templates.


Sema Yuce CISM CRISC CISA, Director at Truth ISC Technology and Security Consultancy Ltd

This book should be mandatory reading within any business to understand the scale and scope of the landscape within which their information security and assurance professionals need to operate.

Andrea Simmons PhD FBCS CITP CISM CISSP MA CIPP/E CIPM

**Information Technology,
Business**

Cover photo: Shutterstock © Pat-s pictures



ISBN 978-1-78017-572-0



9 781780 175720

<https://t.me/learningnets>