INFORMATION SECURITY MANAGEMENT PRINCIPLES

Third edition

Andy Taylor, David Alexander,
Amanda Finch and David Sutton

bcs
The Chartered Institute for IT

INFORMATION SECURITY MANAGEMENT PRINCIPLES

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,
First Floor, Block D,
North Star House, North Star Avenue,
Swindon, SN2 1FA, United Kingdom.
T +44 (0) 1793 417 424
F +44 (0) 1793 417 444
(Monday to Friday, 09:00 to 17:00 UK time)
www.bcs.org/contact
http://shop.bcs.org/

INFORMATION SECURITY MANAGEMENT PRINCIPLES

Third edition

Andy Taylor, David Alexander, Amanda Finch and David Sutton

bcs
The Chartered Institute for IT

Published by BCS Learning and Development Ltd, a wholly owned subsidiary of BCS, The Chartered Institute for IT, First Floor, Block D, North Star House, North Star Avenue, Swindon, SN2 1FA, UK.
https://www.bcs.org

Paperback ISBN 978-1-780175-18-8
PDF ISBN 978-1-780175-19-5
ePUB ISBN 978-1-780175-20-1
Kindle ISBN 978-1-780175-21-8

British Cataloguing in Publication Data.
A CIP catalogue record for this book is available at the British Library.

Publisher's acknowledgements
Publisher: Ian Borthwick
Commissioning editor: Rebecca Youé
Production manager: Florence Leroy
Project manager: Sunrise Setting Ltd
Copy-editor: Mary Hobbins
Proofreader: Barbara Eastman
Indexer: John Silvester
Cover design: Alex Wright
Cover image: Steve Mcsweeny
Typeset by Lapiz Digital Services, Chennai, India

CONTENTS

	Figures and tables	vii
	Authors	viii
	Acknowledgements	x
	Abbreviations	xi
	Preface	xvi
1.	**INFORMATION SECURITY PRINCIPLES**	**1**
	Concepts and definitions	1
	The need for, and benefits of, information security	9
	Sample questions	17
2.	**INFORMATION RISK**	**19**
	Threats to, and vulnerabilities of, information systems	19
	Risk management	24
	Sample questions	36
	References and further reading	37
3.	**INFORMATION SECURITY FRAMEWORK**	**39**
	Organisation and responsibilities	39
	Organisational policy, standards and procedures	47
	Information security governance	53
	Information assurance programme implementation	58
	Security incident management	63
	Legal framework	67
	Security standards and procedures	79
	Sample questions	85
	References	87
4.	**SECURITY LIFE CYCLES**	**88**
	The information life cycle	88
	Testing, audit and review	90
	Systems development and support	93
	Sample questions	100
	Reference	101
5.	**PROCEDURAL AND PEOPLE SECURITY CONTROLS**	**102**
	General controls	102
	People security	104
	User access controls	109

	Training and awareness	117
	Sample questions	123
6.	**TECHNICAL SECURITY CONTROLS**	**125**
	Technical security	125
	Protection from malicious software	126
	Networks and communications	132
	Operational technology	144
	External services	147
	Cloud computing	153
	IT infrastructure	158
	Sample questions	164
7.	**PHYSICAL AND ENVIRONMENTAL SECURITY**	**166**
	Physical security	166
	Different uses of controls	174
	Sample questions	175
8.	**DISASTER RECOVERY AND BUSINESS CONTINUITY MANAGEMENT**	**177**
	Relationship between DR/BCP, risk assessment and impact analysis	177
	Resilience and redundancy	179
	Approaches to writing plans and implementing plans	180
	The need for documentation, maintenance and testing	182
	Need for links to managed service provision and outsourcing	184
	Need for secure off-site storage of vital material	185
	Need to involve personnel, suppliers and IT systems providers	186
	Relationship with security incident management	187
	Compliance with standards	188
	Sample questions	188
9.	**OTHER TECHNICAL ASPECTS**	**190**
	Investigations and forensics	190
	Role of cryptography	194
	Threat intelligence	202
	Conclusion	206
	Sample questions	206
	References and further reading	207
	APPENDIX A	**209**
	Activity solution pointers	215
	Sample question answers	230
	Glossary	233
	Index	241

FIGURES AND TABLES

Figure 2.1	The risk management life cycle	26
Figure 2.2	A typical risk matrix	27
Figure 4.1	The data and information life cycle	89
Figure 6.1	The Plan–Do–Check–Act model	144
Figure 9.1	Symmetric key encryption	196
Figure 9.2	Asymmetric key encryption	198
Figure 9.3	Producing a signed message digest	199
Figure 9.4	Verifying message integrity	200

Table 2.1	One possible rating framework for risk assessment	31

AUTHORS

Andy Taylor, after initially teaching in secondary schools, Andy has been involved with information assurance for over 35 years, starting when he served in the Royal Navy in several posts as security officer. He had responsibility for all classified and cryptographic materials in both warships and shore establishments, at times helping to maintain the effectiveness of the nuclear deterrent. After leaving the Royal Navy, he chose a further career in consultancy and was instrumental in achieving one of the first accreditations for a management consultancy against the information security standard ISO 17799 (now ISO/IEC 27001). As an independent information security consultant, he has provided information assurance advice to a wide variety of organisations in the public and private sectors including the Health Service, Home Office, utility regulators, the Prison and Probation Services and web developers. He has developed and delivered a number of specialist security briefings to help educate users in the effective use of information in a secure manner, and has provided induction security training in many organisations. He has been directly involved with the development, establishment and maintenance of several different certification schemes relating to information security including the assessment of individuals and of training. He is a Fellow of both BCS and the Association for Project Management (APM), a Chartered IT Professional and a member of the Chartered Institute of Information Security. He has a passionate interest in maintaining the highest standards of information assurance and helping others to gain expertise in it.

David Alexander has over 20 years' experience in the field of information security. He has an MSc in Information Security from Royal Holloway, University of London and specialises in advanced network security, information security architectures, cryptographic protocols and the security of operational technology/industrial control systems. He is Senior Security Architect for the Urenco Group. David has worked on the design and assurance of critical national infrastructures around the world and has wide experience of commercial, central government and defence projects. Involved in IT for over 30 years, the first 10 in mainstream IT before he changed sides from 'poacher to gamekeeper', David is a Fellow of BCS and of the Chartered Institute of Information Security, and a Chartered Security Architect and IT professional. He was also one of the first people in the world accredited as Lead Auditor for what is now ISO/IEC 27001, a certification he has maintained through all the versions. As well as working for Urenco, David teaches the Network Security module on the Royal Holloway Information Security MSc programme.

Amanda Finch is the CEO of the Chartered Institute of Information Security and has specialised in information security management since 1991. She has always been an active contributor to the industry and for many years has been dedicated to working towards the discipline being recognised as a profession. Over her career she has been engaged in all aspects of information security management and takes a pragmatic approach to the application of security controls to meet business objectives. Through her work she has developed an extensive understanding of the commercial sector and its particular security needs. In her current role she works with industry, government and academia, assisting all sectors in raising levels of competency and education. Amanda has worked within the retail and banking sectors as well as with the Information Security Forum. She has a Masters degree in Information Security, holds full membership of the Chartered Institute of Information Security with Founder status, and is a Fellow of BCS. In 2007 she was awarded European Chief Information Security Officer of the Year by *Secure Computing* magazine and is frequently listed as one of the most influential women within the industry.

David Sutton's career spans more than 50 years in information and communications technology, incorporating radio transmission, international telephone switching, mainframe computing and data networking. At Telefónica O2 UK he was responsible for ensuring the continuity and restoration of its core cellular networks, and he represented the company in the UK electronic communications industry's national resilience forum. In December 2005 he gave evidence to the Greater London Authority enquiry into the impact of the 7/7 London bombings on mobile telecoms. Since retiring from O2, David has undertaken a number of critical information infrastructure projects for the European Network and Information Security Agency (ENISA), and has developed training material on business continuity and information risk management for InfoSec Skills in addition to authoring books on information security and business continuity. He has been a member of the BCS Professional Certification Information Security Panel since 2005 and a tutor on the distance learning Information Security MSc course at Royal Holloway, University of London. He is a member of the Chartered Institute of Information Security, a Fellow of BCS and a Chartered IT Professional.

ACKNOWLEDGEMENTS

For this third edition, we would like to thank Ian Borthwick for his help in getting this updated edition into print. The cartoons were originally drawn by Ed Brown, so our continuing thanks go to him. We would also like to thank colleagues and clients, families and friends who willingly, or more usually unwittingly, have provided many of the anecdotes, examples and stories with which we have tried to explain some of the principles in this book.

ABBREVIATIONS

2FA	two-factor authentication
4G	International Mobile Telecommunications Advanced or LTE Advanced
5G	fifth generation cellular network telephony
ACL	access control list
ACPO	Association of Chief Police Officers (UK)
ADSL	asymmetric digital subscriber line
AES	Advanced Encryption Standard
ANSI	American National Standards Institute
BCP	business continuity plan
BCS	British Computer Society, The Chartered Institute for IT
BIA	business impact analysis
BS	British Standard
BYOD	bring your own device
CA	certification authority
CAI	computer aided instruction
CAPS	Certified Assisted Products
CAS	Independent Evaluation for Assured Services (UK NCSC)
CASB	cloud access security broker
CBT	computer-based training
CC	Common Criteria (certificate)
CC ITSEC	Common Criteria for Information Technology Security Evaluation Criteria
CCP	Certified Cyber Professional
CCRA	Common Criteria Recognition Arrangement
CCTV	closed-circuit television
CEH	Certified Ethical Hacker (qualification)
CERT	computer emergency response team
CESG	Communications-Electronics Security Group (largely superseded by UK's NCSC)
CFO	chief finance officer

CIISec	Chartered Institute of Information Security
CIO	chief information officer
CISMP	Certificate in Information Security Management Principles
CISO	chief information security officer
CiSP	Cyber Security Information Sharing Partnership
CLEF	Commercial Licensed Evaluation Facility
CMM	Capability Maturity Model
CoCo	code of connection
COSO	Committee of Sponsoring Organizations of the Treadway Commission
COTS	commercial off-the-shelf
CPA	Commercial Product Assurance
CPNI	Centre for the Protection of National Infrastructure
CREST	Council of Registered Ethical Security Testers
CRO	chief risk officer
CSA	Cloud Security Alliance
CTAS	CESG Tailored Assurance Service
CTCPEC	Canadian Trusted Computer Product Evaluation Criteria
CTI	cyber threat intelligence
CVE	Common Vulnerabilities and Exposures database
DCMS	Department for Digital, Culture, Media and Sports
DCS	distributed control system
DDoS	distributed denial of service
DES	Data Encryption Standard
DHCP	dynamic host configuration protocol
DHS	Department for Homeland Security
DMZ	demilitarised zone
DNS	domain name system
DoS	denial of service
DPA	Data Protection Act
DR	disaster recovery
EAL	Evaluation Assurance Level
EDGE	Enhanced Data Rates for GSM Evolution
EDI	electronic data interchange
EDS	ETSI documentation service
EFTA	European Free Trade Association
eIDAS	Electronic Identification, Authentication and Trust Services
ENISA	European Union Agency for Network and Information Security
EPC	European Patent Convention
ERP	enterprise resource planning

ETSI	European Telecommunications Standards Institute
EU	European Union
FAIR	Factor Analysis of Information Risk
FBI	Federal Bureau of Investigation
FCA	Financial Conduct Authority
FIPS PUBS	Federal Information Processing Standards Publications
FIRST	Forum for Incident Response and Security Teams
FoIA	Freedom of Information Act
FSA	Financial Services Act
GATT TRIPS	General Agreement on Tariffs and Trades, Trade Related Aspects of Intellectual Property Rights
GCHQ	Government Communications Headquarters
GDPR	General Data Protection Regulation
GFS	Grandfather-Father-Son
GIAC	Global Information Assurance Certification
GLBA	Gramm-Leach-Bliley Act
GPEN	GIAC Penetration Tester (qualification)
GPRS	General Packet Radio Service
GSM	Global System for Mobile Communications standard (2G)
HIDS	host intrusion detection system
HIPAA	Health Insurance Portability and Accountability Act
HRA	Human Rights Act
HSDPA	High-Speed Downlink Packet Access
HTTP(S)	hypertext transfer protocol (secure)
IA	information assurance
IaaS	infrastructure as a service
ICO	Information Commissioner's Office
ICS	industrial control system
ICT	information communications and technology
ID&A	identification and authentication
IDC	inter-domain connector
IDS	intrusion detection system
IEC	International Electrotechnical Commission
IETF	Internet Engineering Task Force
IISP	Institute of Information Security Professionals
IKE	Internet Key Exchange protocol
IM	instant messaging
IoC	indicator of compromise
IoT	Internet of Things
IP	intellectual property

IPR	intellectual property rights
IPS	intrusion prevention system
IPSec	internet protocol security
IRT	incident response team
IS	information systems
ISDN	integrated services digital network
ISF	Information Security Forum
ISMS	information security management system
ISO	International Organization for Standardization
ITIL	Information Technology Infrastructure Library
ITSEC	Information Technology Security Evaluation Criteria
ITT	invitation to tender
ITU	International Telecommunication Union
LAN	local area network
LTE	long term evolution (see also 4G)
LOB	line of business
MFA	multi-factor authentication
MiFID	Markets in Financial Instruments Directive
MPLS	multi-protocol layer switching
NCA	National Crime Agency
NCSC	National Cyber Security Centre (part of GCHQ)
NDA	non-disclosure agreement
NIDS	network intrusion detection system
NIS	Network and Information Systems directive
NIST	National Institute of Standards and Technology
NOC	Network Operations Centre
NPCC	National Police Chiefs' Council
OCTAVE	Operationally Critical Threat, Asset and Vulnerability Evaluation
OES	operator of essential services
OOB	out of band
OSA	Official Secrets Act
OSCP	Offensive Security Certified Professional (qualification)
OSI	Open Source Intelligence
OT	operational technology
OTP	one-time password
PaaS	platform as a service
PABX	private automatic branch exchange
PACE	Police and Criminal Evidence Act
PAS	Publicly Available Specification
PCBCM	Practitioner Certificate in Business Continuity Management

PCI	Payment Card Industry
PCI DSS	Payment Card Industry Data Security Standard
PCIRM	Practitioner Certificate in Information Risk Management
PDCA	Plan–Do–Check–Act
PenTest	penetration test
PGP	Pretty Good Privacy
PII	personally identifiable information
PIN	personal identification number
PKI	public key infrastructure
ProtMon	protective monitoring
RDSP	relevant digital service provider
RFC	Request for Comments
RIPA	Regulation of Investigatory Powers Act
ROI	return on investment
SaaS	software as a service
SABSA	Sherwood Applied Business Security Architecture
SANS	Sysadmin, Audit, Network, Security
SCADA	supervisory control and data acquisition
SIEM	security information and event management
SLA	service level agreement
SOC	security operations centre
SOMA	Security Operations Maturity Architecture
SSL	secure sockets layer
SSO	single sign on
TCSEC	Trusted Computer System Evaluation Criteria
TLS	transport layer security
ToE	target of evaluation
TTPs	tactics, techniques and procedures
UMTS	Universal Mobile Telecommunications Service (3G)
UPS	uninterruptible power supply
US NCSC	United States National Counterintelligence and Security Center
VOIP	voice over internet protocol
VPN	virtual private network
WA	Wassenaar Arrangement
WAN	wide area network
WAP	wireless access point
WEP	wired equivalent privacy
Wi-Fi	wireless fidelity
WPA	Wi-Fi protected access

PREFACE

Data and information have been important for a very wide variety of reasons and for as many centuries as man has been able to pass valuable data to another person. The location of the nearest water hole, herd of wild animals or warm cave was a carefully guarded secret that was only passed on to those with a need to know and who could be trusted not to divulge the information to other, possibly hostile, tribes. The methods of transfer and the storage of such information were perhaps rather more primitive than today, but the basic principles of information security have not changed too much since those days.

Information assurance is now well founded in three major concepts – those of confidentiality, integrity and availability. Managing these concepts is critical and, as information has increasingly become one of the modern currencies of society, it is the retention of assurance in an appropriate and cost-effective manner that has become of keen interest to businesses in all sectors, of all sizes and in all locations. Specific measures taken to ensure that information is held securely is termed 'information security' – the way of achieving information assurance.

As an example, even within living memory, the quantity of numbers we are given and need to enable us to exist and participate in modern society has risen almost exponentially from virtually zero in the early part of the 20th century, to several hundred (and still growing) now: PIN codes; licence numbers; credit card numbers; number plates; telephone numbers; employee number; health, tax and insurance numbers; access codes; customer numbers; train times; tram or bus numbers; and so on. We now need to remember such numbers on a day-to-day basis, and that is before we start work proper and have to deal with all those things that allow us to earn our salary, where even more numbers and other elements of information will occur.

The mechanisms we use to manage information are the areas where we have seen very significant change, notably in the last few decades. The advent of computers in particular has extensively altered the way we manage information and has also meant that we have much more information to worry about than ever before. Information has become the key to success in almost any field and so the assurance of it has gained in significance and, perhaps more importantly, in value to a business or organisation. It may not necessarily be financial value that is the most important factor. Lack of knowledge of some issue or the way things are done, or knowing the currency of specific pieces of information may be more important than any financial valuation. Nevertheless, looking after it properly is still very important.

One other factor that has significantly altered our need for assurance of information is that of mobility. Life was straightforward when the only place we had business information, and where we were able to look after it properly, was the office – to secure information, we closed and locked the office door. Today we expect and need to have information in a wide variety of locations, including wanting it on the move in cars, trains and planes. With open plan offices and the increasing mobility of the office environment, we now have a critical need for improved assurance if we don't want others to gain access to our information inappropriately.

Threats, vulnerabilities and countermeasures have also changed and grown in complexity in some areas, although it is still essential to consider the easiest and often cheapest countermeasures before getting into large or expensive solutions. The increase in capability of those intent on causing harm to companies, public bodies and other organisations means that the role of the information assurance manager and the professional has increased in complexity to such a degree that it is now quite possible to have a full and very satisfying working life entirely within this field of expertise.

Since the late 1980s a new term that has come to prominence is 'cyber security'. The reasons for this are largely down to the significant increase in threats – the complexity of threats, the number of threats and the potential impact of threats – that now arise from the internet and the use of the World Wide Web. Cyber has been used to describe the risks and vulnerabilities that arise primarily from the use of the internet and so cyber security has become the most commonly used term to address these areas. In this edition of this book, we have continued to use the term 'information assurance' where general principles are discussed, have used 'information security' again where it is the most appropriate term, but have also referred to 'cyber security' where the threats are specifically internet based. With the seemingly meteoric rise in what are now known as cyber-attacks, we see more and more attempts to misappropriate information. Criminals and others want to steal information and sell it on or use it for other purposes; to encrypt information and then demand money to release it back to its rightful owner; and to use information gained fraudulently through any means to extract financial gain from seemingly innocent victims, be they businesses or individuals. This is cyber warfare and leads to cyber security.

The legislation that is introduced by governments to address the increasing problems of information assurance in all its guises, is also an area of concern and this book covers the most important principles and the implementation of such laws. Once again, though, it is important that you understand that this book has been written in the UK and is based on English law. Other countries, even Devolved Administrations within the UK, may have further or different legislation with which you should become acquainted. Reference has been made to national and international standards applicable to information assurance, but there is no requirement in the examination for the BCS Certificate in Information Security Management Principles (CISMP), upon which this book is based, for detailed specific knowledge of any of those standards. They are naturally important, but it is recognised that they will change over time and be more applicable in some parts of the world than in others. You should ensure you are familiar with the standards relevant to your country, your area of interest, your organisation and your business sector.

This book accompanies the BCS Certificate in Information Security Management Principles. This qualification, one of a series covering the whole area of information

assurance management, is the first step towards a full understanding of the issues and the comprehensive management of the assurance of information wherever it may be. This book is intended to be a first read for those new to information security and concentrates on the high-level principles. It is not intended to be a comprehensive guide to everything that a practitioner in the area would need to know.

The technical aspects of information security, including the technical details of information systems (IS), computer networks, communication systems, cryptography and related areas, are not part of the syllabus for this qualification despite their importance. However, they appear in higher qualifications, so in this book reference is made to them in passing but they are not covered in any detail. The syllabus and this book have remained technology neutral as far as possible.

While BCS, The Chartered Institute for IT, is clearly mainly concerned with the impact and effective use of computers, it is recognised that it is impossible to divorce the management of information security in computers from the management of information in any other media or from the security of the tools used to process information. Thus, in this book, the boundaries between different forms of information storage, processing, transmission and use are deliberately blurred or indeed removed entirely. It is not significant whether a particular piece of information exists in electronic form, paper form or indeed in someone's head. Its appropriate protection is the main factor, and all aspects of its assurance must be considered from all angles.

The latest version of the examination syllabus can be downloaded from the BCS website[1] and it is the guide for the contents of this edition of this book. As a result of studying this book, you should gain a very clear understanding of the various elements of information assurance and should be able to consider taking the professional examination. It would naturally be useful for an individual to undertake a period of study with an approved training provider to enhance their understanding, and those who deliver such training will inevitably add value to the knowledge given here, probably increasing the chances of success in the examination.

There are some areas where this book does not provide all the detail necessary to answer all the questions in the examination, but there are ample suggestions for additional study and resources for further reading that would help. A simple scenario has been introduced in order to help develop full understanding and to provide a close-to-life example of the real world. Activities based on the scenario are suggested throughout the book, again to help bring reality into the concepts discussed, and it is hoped that you will do these in an appropriate manner – formally or informally as suits you best. The format of the multi-choice questions in the book is broadly in line with the questions in the examination, but naturally there will be different questions in that. A sample examination paper can also be downloaded from the BCS website.

1 https://bcs.org/get-qualified/certifications-for-professionals/information-security-and-ccp-scheme-certifications/bcs-foundation-certificate-in-information-security-management-principles/

After studying this book and the related syllabus, you should be able to demonstrate a good knowledge and basic understanding of the wide range of subject areas that make up information assurance management. The examination tests the knowledge of principles rather than the knowledge of specific technologies, products or techniques. This means that where in the book specific technical examples are used to illustrate particular principles, it is the understanding of the principles that is of prime importance when considering the examples, and not the examples themselves. If more information is required in specific areas, such as risk management, business continuity or project management, other BCS publications are available that provide a much deeper understanding. Full details of appropriate publications can be found on the BCS bookshop.[2]

2 https://shop.bcs.org/

1 INFORMATION SECURITY PRINCIPLES

This chapter covers the basic principles of information assurance (IA). It introduces some specific terminology together with its meaning and definitions and considers the use of such terminology across the field of information assurance management. It also discusses the way in which information assurance management relates to its environment.

CONCEPTS AND DEFINITIONS

As in any area of business, information assurance management has its own language, although, being very closely related to business need, it is limited in scope and complexity to enable the wider business population to appreciate the concepts with little difficulty. Each of the terms listed below will be further discussed and expanded upon later in the book in the appropriate section.

In the following sections the definitions in italics have been taken from the BS ISO/IEC 27000 series of standards (latest editions at the time of writing) where the definition exists, and from other ISO standards where there was no 27000 definition. Where there is no extant definition, it is provided by the authors or from other sources, noting its source where applicable.

LEARNING OUTCOMES

Following study in this area, you should be able to define and explain each of the following terms and to describe their appropriate use as applicable.

Information security

Confidentiality. The property that information is not made available or disclosed to unauthorised individuals, entities or processes (ISO/IEC 27000)

Information will often be applicable only to a limited number of individuals because of its nature, its content or because its wider distribution will result in undesired effects, including legal or financial penalties or embarrassment to one party or another. Restricting access to information to those who have a 'need to know' is good practice

and is based on the principle of confidentiality. Controls to ensure confidentiality form a major part of the wider aspects of information assurance management.

Integrity. The property of accuracy and completeness (ISO/IEC 27000)

Information is only useful if it is complete and accurate, and remains so. Maintaining this aspect of information (its integrity) is often critical and ensuring that only certain people have the appropriate authority to alter, update or delete information is another basic principle of IA.

Availability. The property of being accessible and usable upon demand by an authorised entity (ISO/IEC 27000)

Information that is not available when and as required is not information at all but irrelevant data. Availability is one area where developments in technology have increased the difficulties for the information assurance professional very significantly. In the past, in an ideal world, all important information could be locked up in a very secure safe of some form and never allowed to be accessed – just about perfect assurance but, naturally, totally impractical. There will, therefore, always have to be a compromise between security in its purest sense and the availability of the information. This compromise has to be acknowledged throughout all aspects of IA and has a direct bearing on many of the principles covered in this book.

DATA OR INFORMATION?

Data (sometimes clarified as raw or unprocessed data) are generally accepted as being the basic facts and statistics that can be analysed and subsequently used for many different purposes.

Information is the result of the analysis of the data – the refined information that is useful to operators and managers to understand what is going on; for example, on their IT systems.

Assets and asset types

Asset. Anything that has value to the organisation (ISO/IEC 13335)

Assets come in as great an array of types as the mechanisms for using them. In information assurance, three main types of assets are considered, although the sub-categories that fall within each of these main types can be numerous. The three main types are:

1. pure information (in whatever format);
2. physical assets such as buildings and computer systems;
3. software used to process or otherwise manage information.

When assets are considered in any aspect of IA, the impact on all three of these asset types should be reviewed. The value of an asset is usually estimated on the basis of the cost or value of its loss or unavailability to the business through a business impact assessment. There are, however, other aspects to consider, including, but not limited to, the value to a competitor, the cost of recovery or reconstruction, the damage to other operations and even the impact on such intangibles as reputation, brand awareness and customer loyalty.

Threat, vulnerability, risk and impact

The understanding of these terms is critical to the whole of information assurance.

Threat. A potential cause of an unwanted incident, which may result in harm to a system or organisation (ISO/IEC 27000)

A threat is something that may happen that might cause some unwanted consequence. As a simple example, if we see clouds in the sky that look large and dark, we talk about the threat of rain. Naturally, to some this threat is not unwanted at all, farmers perhaps, and so they would not have the same view of the clouds and their potential for rain – and this is an important point to recognise. Threats to one organisation may well be opportunities to another, it is all very dependent on the viewpoint, the environment and the situation in which they are being considered.

Vulnerability. A weakness of an asset or control that can be exploited by one or more threats (ISO/IEC 27000)

A vulnerability is a weakness; something that, if exploited, could cause some unwanted effect(s). To continue the example above, if someone was to venture out into the cloudy environment without an umbrella, this could be considered a vulnerability. If something (the threat) happens (it rains) then the consequences could be detrimental.

Risk. The effect of uncertainty on objectives (ISO/IEC 27000)

Risk, then, is the combination of these two. If there is a threat (of rain) and a vulnerability (of not carrying an umbrella) then there is a risk that the individual concerned might get wet and ruin their expensive clothes. There may well be other risks associated with this same set of circumstances – ruined hair style, late attendance for an appointment, and so on. It is also important to recognise that sometimes there may be a combination of circumstances that lead to further, more serious risks as well. The lateness of attendance at an appointment combined with a number of other similar occurrences could result in termination of employment. It should be noted, however, that if either the threat or the vulnerability is removed in some way, there is no longer a risk. Both must be present for the risk to exist at all.

Impact. The result of an information security incident, caused by a threat, which affects assets (ISO/IEC 13335)

The impact of the risk actually occurring is perhaps the most important concept of all to grasp. It is the potential impact that has to be considered and managed in IA. If the impact is small and insignificant – a wet coat in the example above – then it may be

entirely appropriate to accept the risk and to take no further action other than to monitor it. On the other hand, if the potential impact could be dismissal from a well-paid job, then more appropriate countermeasures need to be considered – the purchase of an umbrella, hiring a taxi or similar. As far as businesses are concerned, the impact on the organisation and its daily activities is usually the crucial consideration and will often warrant further measures being taken.

Information security policy concepts

Any organisation should have a policy for its management of IA. This is normally a short, punchy statement from the chief executive stating that they acknowledge the risks to the business resulting from poor information assurance and will take appropriate measures to deal with them. It should include statements that make it clear that the organisation regards risk as a serious issue, with it being discussed at all appropriate meetings, with those with the correct authority and responsibility taking an active interest in it. It is common for organisations to form an information assurance or security working group to lead the activities necessary to ensure appropriate levels of assurance within the organisation.

The purpose of controls

Controls in the IA sense are those activities that are taken to manage the risks identified. There are four main types of strategic control, although the actual implementation of each of these types can be very varied.

Eliminate. Risk avoidance – 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)

This means taking a course of action(s) that removes the threat of a certain risk occurring at all. This could entail removing a particular item that is unsafe, choosing to do things in a completely different way or any number of other options. This action is sometimes referred to as 'prevent', 'avoid' or 'terminate'.

Reduce. Risk reduction – Action taken to lessen the probability, negative consequences, or both, associated with risk (ISO 22300:2018)

This means to take one or more actions that will reduce the impact or the likelihood of a risk occurring. It is rare for an action to both reduce the likelihood and reduce the impact of a risk. It is often necessary to use several of these measures in partnership to have the desired overall effect. This could include having contingency measures in place that mitigate the effect if the risk does occur – a backup plan or 'plan B'. This action is sometimes referred to as 'treat' or 'mitigate'.

Transfer. Risk transfer – A form of risk treatment involving the agreed distribution of risk with other parties (ISO Guide 73)

This means to take steps to move the accountability for a risk to another organisation who will take on the responsibility for the future management of the risk. In practice, this might mean taking out some form of indemnity or insurance against the risk occurring or perhaps writing contracts in such a way that the financial impact of a risk

occurring is borne by a third party – liquidated damages. It should be noted though that, for example, taking out an insurance policy to cover the costs of rectifying the results of a risk happening will often not take away the impact. Reputation is the most common example where, although the insurance company may pay out the costs incurred by the client in dealing with an issue, the reputational damage to the organisation may still be very evident. This action is sometimes referred to as 'share'.

Accept. Risk acceptance – The decision to accept a risk (ISO Guide 73)

This means senior management accepting that it is not considered practical or sensible to take any further action other than to monitor the risk. There could be a number of reasons why further actions are considered inappropriate, including but not limited to: the likely impact of a risk is too small; the likelihood of a risk occurring is too small; the cost of appropriate measures is too high in comparison with the financial impact of the risk occurring; the risk is outside the organisation's direct control. The decision is also related to the organisation's risk appetite, which determines the level of risk the organisation is prepared to accept. This is sometimes referred to as 'tolerate' but should not be termed the 'do nothing' option.

Identity, authentication and authorisation

Identity. Information that unambiguously distinguishes one entity from another one in a given domain (ISO/IEC 24760-1)

Frequently there is a need to establish who is accessing information, and the identity of individuals may well be required. This may enable, for example, audit trails to be produced to see who changed a specific item of data and hence to assign an appropriate level of confidence to the change. This concept is equally applicable to assets such as specific pieces of information that need to be identified uniquely.

Authentication. The provision of assurance of the claimed identity of an entity (ISO/IEC 15944-6)

This process ensures that the individual is who they say they are and confirms their identity to an appropriate level of confidence appropriate for the task in hand. This could be simply asking them for their date of birth, at the most basic level, through to completing a complex identity check using, for example, tokens, biometrics and detailed biographical-data checks.

Authorisation. The right or permission that is granted to a system entity to access a system resource (ISO/TR 22100-4)

In order for anyone to use a system of information retrieval, management and so on, it is good practice to have a method of authorisation that makes clear the assets to which someone should have access and the type of access they should have. This authorisation will vary depending on the business requirement, the individual, the type of asset and a range of other aspects. Who has the authority to detail and approve such authorisations will vary according to the type of usage required.

Accountability, audit and compliance

Accountability. The property that ensures that the actions of an entity can be traced uniquely to the entity (ISO/IEC 21827)

When any action is carried out on an information system or as part of the information assurance management system, an individual needs to be accountable for that action. The person who has the accountability may delegate the actual work to someone else, but they still retain the accountability.

Audit. The review of a party's capacity to meet, or continue to meet, the initial and ongoing approval agreements as a service provider (ISO 15638-15)

This is the checking (formal or informal) of the records of a system to ensure that the activities that were anticipated to have taken place have actually happened. The purposes of an audit could include identifying gaps in the system's functionality, noting trends over time to help with problem resolution or identification, or a number of other requirements. It can also help to identify misuse of information or the inappropriate use of an authorisation, for example, and thus identify unauthorised activity.

Compliance. Meeting or exceeding all applicable requirements of a standard or other published set of requirements (ISO/TR 19591)

Ensuring that a system or process complies with the defined or expected operating procedure is compliance. This could cover a major operation, such as a whole organisation being compliant with a recognised national standard for information assurance, or could be much more limited with just certain aspects of the operation, or even individual users of a specific system being compliant. In general, compliance should be independently audited to achieve certification against a standard; for example, a legal or regulatory framework.

Information security professionalism and ethics

General awareness of the work done by information assurance professionals (as distinct from IT security professionals) is gradually growing as organisations become increasingly complex with more and more information being managed and processed. The adage that the staff are the most important asset of an organisation could now be seen to be outmoded since it is often the case that it is the information an organisation holds and uses effectively that has become its most important asset. Therefore, looking after it has also increased in importance and the whole profession has grown to meet the need. Professional bodies, such as the Chartered Institute of Information Security (CIISec) (previously the Institute of Information Security Professionals (IISP) that was set up in 2006 in the UK), have helped to raise the profile very significantly, as have the various qualifications ranging from this introductory level to Masters degrees and beyond.

The UK's National Cyber Security Centre (NCSC) have developed a certification scheme (the Certified Cyber Professional (CCP)) where individuals can demonstrate their competence and experience to independent assessors from a certification body, who

will recommend the award of a certificate in a specialism when the appropriate criteria have been met.

An information assurance professional will, inevitably, become party to some of the most important information an organisation might hold. This could be sensitive for a number of reasons, but in all cases it is critical that the professional deals with it in the appropriate manner. Releasing information to a third party or other organisation, albeit with the best of intentions but without the approval of the owner, is probably one of the easiest ways to be dismissed. Non-disclosure agreements (NDAs) are now commonplace even in seemingly innocuous areas such as publishing and the retail marketplace, as well as the more usual research and development, product innovation and financial areas.

The bottom line of all assurance is trust. Without it, it is impossible to operate in the world as it is today. The degree of trust is where there is room for manoeuvre and it is often the degree to which staff, customers, suppliers, shareholders and the like can be trusted that will determine the measures that have to be put in place. It is crucial though that the trust placed in information assurance professionals is not misplaced in any way. They must be above reproach and never be seen to compromise in this critical area.

The information security management system concepts

Information Security Management System (ISMS). Part of the overall management system, based on a business risk approach, used to establish, implement, operate, monitor, review, maintain and improve information security (ISO 12812-2)

The main principle behind the ISMS is that there should be a 'one-stop shop' for all information pertinent to the assurance of information within an organisation. As soon as there is a need to go looking for documentation, policies, practices or anything else to do with assurance, the chances are that someone will not bother and will do their own thing.

While there may well be good reason for them not to do this in terms of rules, regulations, punishments and the like, human nature being what it is, they will find a reasonable excuse for going down a different route if only because 'I thought it was OK and couldn't be bothered to check if it was the right way to do it.' The result of this approach will inevitably be a reduction in the overall level of assurance. In addition, any system that is too complex or difficult to use will result in users finding ways to get around the security measures put in place, perhaps again resulting in weakened assurance.

It is critical, therefore, that organisations make their information as freely and easily available as is possible, practical and necessary and this equally applies to the security rules controlling it. Naturally there will be elements of policy that have to be more secure, available only to those with a strict need to know, but in general everyone should be able to access easily and quickly the appropriate information and the security measures pertinent to it.

National and international security standards

As a policy, BCS have decided not to relate the syllabus for the BCS Certificate in Information Security Management Principles to any national or international standards or frameworks for information security specifically, although there are many such documents that are applicable. The main reasons for this were two-fold: first, to make the syllabus and the qualification as applicable internationally as possible and, second, to reduce the need to update the syllabus at every change to the standards.

It is clear, however, that IA is the subject of several international and national standards and that these should be considered when studying for the examination. The questions set in the examination will never be specific to any one standard, but will be generic to all best practice where applicable. The knowledge of the appropriate standards required for the examination is therefore limited to a general understanding of the principles involved as they reflect on best practice. In the UK, awareness of, for example, the ISO/IEC 27000 series and related British standards is helpful but not critical to the passing of the examination. It is the broad principles that should be used as a basis for study, as reflected in the examination syllabus.

There is, though, another aspect of this. When an information assurance professional is working in an organisation to deliver a secure and effective information management system, the relevant standards should always be viewed as the achievable goal for the system. Whether it is necessary to gain simple compliance or go the extra step to achieve certification is an arbitrary decision often based on other factors. Nevertheless, it is considered good practice to base an effective information assurance management system on the principles of the relevant standards. The use of an internationally accepted standard such as the ISO/IEC 27000 series makes sense in the global nature of operations today.

THE GROUP FOR THE APPRECIATION OF THE NATTERJACK TOAD SCENARIO

The Group for the Appreciation of the Natterjack Toad (GANT) is a conservation group that is keen to promote and preserve the well-being of the Natterjack toad. It has a significant number of members in many different countries around the world, all of whom are keen to promote the work of the Group, which is a charity registered in the UK. All of GANT's information is either on a web-based application available to members over the internet or on old-fashioned, paper-based documents held by Dr Jane Peabody, the honorary secretary/treasurer.

The Natterjack toad is an endangered species that is gradually being destroyed by the development of areas where it prospers and through pollution affecting the brackish water and sand dunes in which it lives.

The membership of the organisation is growing and the system for managing the records of members is one area where there are some concerns about information assurance. Details of GANT's activities, their meeting places, their website and

other aspects of the Group's work have been compromised in the recent past owing to the server containing them having no significant security in place. The chairperson (Ms Rachel Jackson) believes it is the right time to take information security more seriously. She has heard a bit about information assurance but needs to be clear what it really means and, most importantly, what the benefits and costs would be to the organisation.

The GANT scenario in the box above is a fictitious scenario that will be used throughout the book to provide examples and to be the basis of some questions to aid your understanding of the theory. The main objective of the scenario is to implement an effective IA system, but we will take you through various steps along the way to help with your understanding.

ACTIVITY 1.1

Assume that you have been invited to a committee meeting of GANT by the chairperson, who wants you to 'start the ball rolling' by explaining why it would be a good idea for GANT to think about information assurance.

To make your points most forcefully, she has asked you to define three threats to the organisation, three vulnerabilities and consequently three risks that any information assurance system would need to manage.

We have started above with developing an initial idea of the reasons for considering IA based on three possible problems. We will take that on to a more formal approach in due course – this is simply to get you thinking about some of the terms we have introduced in the first section of the book. Solution pointers for all the activities are at the end of the book.

THE NEED FOR, AND BENEFITS OF, INFORMATION SECURITY

Any business will have information that is critical to its continued effective operation. Looking after this information in an appropriate way does not come free but has a price tag attached that can be, in some circumstances, very considerable. It is therefore essential that information assurance professionals are able to justify their recommendations for appropriate security measures in a sensible yet pragmatic manner, which must take into account the specific environment in which the business is based.

LEARNING OUTCOMES

Following study in this area, you should be able to explain and justify each of the following concepts and to describe their appropriate use as applicable.

The importance of information security as part of a business model

Information security – Preservation of confidentiality, integrity and availability of information; in addition, other properties such as authenticity, accountability, non-repudiation and reliability can also be involved. (ISO 19092)

Neither information nor assurance operate in a vacuum. Both need to take into account the environment in which they are operating and address the issues that this environment brings with it. It is therefore critical that any information assurance system must be grounded firmly in the business world. This means that IA is not an issue only for the IT manager or the security officer but for the whole organisation. As soon as only one part of the organisation is given the task of running assurance, the rest of the organisation will bother less about it. All staff members of any organisation, regardless of its nature, its business, its location or any other factor, should be concerned about IA. It might be from a purely personal viewpoint (what happens to my personal data in this place?) or from a wider view of the effective, continued operation of the organisation, but in either case everyone should be concerned and involved.

Information assurance (IA) – The confidence that information systems will protect the information they carry and will function as they need to, when they need to, under the control of legitimate users. (UK Cabinet Office)

Physical, technical and administrative controls are needed to accomplish these tasks. While focused predominantly on information in digital form, the full range of IA encompasses not only digital but also analogue or physical form. These protections apply to data in transit, both physical and electronic forms, as well as data at rest in various types of physical and electronic storage facilities. Information systems include any means of storing, processing or disseminating information including IT systems, media and paper-based systems.

Assurance should not be viewed as an 'add-on' to be included only if there is the time and the money to do it. It has to be built-in to business processes at all stages if it is to be truly effective. While it might be possible in some areas to add in security measures at the last moment (an extra lock on a door or an additional staff security check, for example), they will usually cost more and be less effective than if they had been added at the appropriate time earlier in the design process.

Different business models and their impact on security

In the last 20 years, the world of business has changed dramatically – perhaps more than in the previous 50 or 100 years. One of the principal reasons for this is the increased use of technology that has enabled business to be transacted remotely rather than in person. One of the consequences of this is that more people are able to make business

transactions themselves rather than expecting others to act as intermediaries. No longer do we need to use travel agents to book our flights, local garages to obtain our cars for us or financial advisors to obtain investment packages for us. All these and many more transactions can be carried out directly with the supplier, often using the internet for communications, or with a trader in another part of the country or the world who can offer a better deal. While the access to such facilities is a huge advantage and can provide very significant financial savings, among other benefits, it has brought with it major issues of security both for the individual and for the organisation wishing to trade in this way.

The other very significant change in business has been the shift in the UK away from manufacturing and related primary industries to service and financial industries where the use of technology has an even bigger impact.

It is clear that the use of technology in manufacturing has changed those industries too but, it might be argued, in a more controlled and manageable manner. However, it would be wrong to assume all is well in the factories; issues with the security of industrial control systems (ICSs) are increasing in number and severity. There will be more about this specific issue later.

In the service industry, the availability of information has increased many times over and has liberated the industry in a manner that is similar to the impact of the introduction of the steam engine or electricity in their day. This in turn has increased the importance and difficulty of keeping the information secure.

Many organisations are now based and/or operate in more than one country. With global organisations now moving very sensitive information or other assets around the world at a moment's notice, the need to ensure it is done securely and with proof of receipt, integrity and authority has grown too. Proving that the authorised person sent the correct document at the appropriate time only to the intended recipients, not to mention ensuring that it arrives in the same state as when it left the originator, are all issues that the information assurance manager now has to deal with to the satisfaction of their management and any ambitious litigant. In addition, organisations that operate within different countries need to understand the differing restrictions that local legislation may place on how their information can/must be handled.

There are many further risks from this change in business model. With an increasing amount of trade being conducted across the internet, organisations must be aware of the dangers of virus infection including ransomware, denial-of-service attacks, unauthorised changes to their information in the public domain (e.g. websites) and the impact of any such issues on their reputation, financial status and other related areas. In addition, organisations are having to deal with people about whom they know very little but with whom they still need to establish an appropriate level of trust. The ability of disillusioned employees, ex-employees or groups of activists to damage an organisation by taking, deleting, altering or otherwise misappropriating critical business information from the employer, and either passing it to a competitor or simply using it for their own ill-gotten gains, is now a very real issue. Companies who have been the victims of such events are not inclined to increase the damage caused by making such acts public knowledge if they can avoid it; however, there are many apocryphal tales of the theft of client databases, deletion or alteration of critical financial data and other similar acts, which suggests that some at least are true.

There are also cautionary tales of laptop PCs containing highly sensitive or confidential information being lost or stolen from parked cars, to the embarrassment of the company or organisation. All mobile devices, such as tablets, smartphones and the like, are seen as easy targets for the attackers and, since many such devices are under the ownership of an individual rather than the organisation whose information may be accessed or held on it (what is usually labelled as bring your own device, or BYOD), the way this attack vector is managed has to be considered very seriously.

The use of the internet for transactions, be it shopping for cars, food or financial services, as well as the storage of client, stock, financial and related information in a secure manner, has further increased the problems to be managed. Often this storage is no longer in a place accessible by the owner of the information, since it is stored in a cloud-based system potentially anywhere in the world. This has given rise to the term 'defence in breadth', which means that all connected systems must now be taken into account when considering how an attack might materialise and the effect it might have. The systems of suppliers and advisors may well be an easier way into the more secure systems of an organisation, since the supplier is a trusted partner and perhaps not subject to the same level of security scrutiny as someone coming in from the outside. This aspect is countered by using defence in depth: layers of security that may start off as relatively low level, but which can increase in complexity, cost and effectiveness as the information and systems being protected get more and more sensitive or important. It should not be a straightforward activity for a criminal to gain access to a low value system or network and to be able to traverse into more complex and sensitive areas without significant additional security measures being encountered.

The ability of the consumer to deal directly with the manufacturer has increased the risks for industry as well as for the consumer, as the problems of unreliable services or products still abound. With the rise of business-to-business transactions, just-in-time operations and other similar services that rely heavily on the timely and accurate movement, storage and retrieval of critical information, the loss of a computer system for a comparatively short while can and has created serious financial losses for the businesses concerned.

The UK's Department for Digital, Culture, Media and Sports (DCMS) estimated in their Cyber Security Breaches Survey 2019[1] that there was an average cost of a single cyber-attack on larger businesses (those with more than 250 employees) of around £22,700, in direct costs. This figure does not represent the whole story, as indirect costs such as reputational damage and loss of productivity were not really included. The incidence of cyber-attacks has also continued to increase according to the survey, with 61 per cent of large businesses reporting a cyber-attack in the previous 12 months.

The effect of the rapidly changing business environment

'It is change, continuing change, inevitable change, that is the dominant factor in society today.' This quotation is from Isaac Asimov,[2] and it is now well understood that for a

1 https://www.gov.uk/government/statistics/cyber-security-breaches-survey-2019

2 'My Own View' in *The Encyclopedia of Science Fiction* (1978) edited by Robert Holdstock; later published in *Asimov on Science Fiction* (1981).

business to survive in the current climate of change, it must adapt and be able to adapt rapidly. This means that what was acceptable as a business practice last week may no longer be acceptable this week; therefore, any assurance system put in place must reflect this changing climate and be flexible enough to cope with it. However, this does not mean that the assurance can be relaxed or reduced in any way. Indeed, if anything, the flexibility should produce a higher level of security and assurance that risks are being managed effectively.

Balancing cost and impact of security with the reduction in risk

Life can never be risk free. In fact, it is often considered that life is all about risk and its effective management. The measures taken in an organisation to reduce risk to an acceptable level can at times become excessively expensive. A careful balance must be struck between the cost or business impact of a risk if it occurs and the cost of the measures taken to reduce its likelihood or impact.

A typical example is insurance. An insurance policy may help to offset the cost of a risk occurring by providing the necessary financial backing to be used to deal with the occurrence of a risk. However, if the cost of the insurance policy is too high, it may simply be cheaper to accept that the risk might occur and pay the smaller amount out to deal with its consequences. It must also be remembered that while it may be possible to transfer to a third party some of the impact of a risk occurring – the financial impact, for example – it is frequently very difficult to transfer the other consequences of a risk, notably the impact on reputation, public opinion or other related results.

It is not uncommon for organisations to put in place extravagant measures to reduce the impact or likelihood of risk occurring when in reality the consequences of the risk occurring are limited, or the actual chance of it happening is so small that the expense is a waste of both money and effort in managing the risk unnecessarily.

A second problem is that of maintaining the currency of risk countermeasures. Once defined and planned, it is critical that they are not simply put on the shelf to await the risk arising. The world around us changes and so the countermeasures may not be valid or may change in their effectiveness or cost as time moves on. Thus, risk management, and the maintenance of the consequential actions taken, is a continual and iterative process that must not be allowed to whither through lack of action or misplaced belief that the situation will not change.

Information security as part of company policy

Assurance or security is not an add-on. It is not possible to deal adequately with assurance by considering it as an additional expense to be avoided if at all possible. The most effective way to deal with it is to include it from the beginning in all areas of the organisation. To this end, the inclusion of assurance as part of the operational policy of the organisation is the only cost-effective way of covering the issues adequately.

There are clear similarities between information assurance and health and safety issues. As soon as health and safety are seen as one person's problem (that of the health and safety officer), the battle for a safe working environment has been lost. Similarly, assurance is not the concern solely of the information security manager, but of the whole organisation. It is essential also that this involvement is from the top of the organisation to the bottom. Just implementing IA at middle management or on the shop floor is meaningless and will inevitably lead to further assurance issues. Senior management have a critical role to play to ensure they engender a working environment where IA is the norm and accepted by all.

The need for comprehensive policy, standards, guidelines and procedures documentation

Just having a policy for information assurance or information security on its own is meaningless. It must be fully supported by a range of other documentation covering the standards expected, the guidelines of how to do things correctly and procedures for what must be done to preserve the assurance of the information in question. This documentation must be comprehensive but digestible, pithy, something that can be read easily and something they will actually read. Not a 1,000 page document that, with all good intentions, the average Joe will not read.

It is good practice to ensure that any procedures to be followed are detailed in an easily digestible format, perhaps as desk cards or prompts for users, or as checklists for operators or support technicians. It must be remembered, however, that this is not only about computers. For example, procedures are also required for the management of physical assets such as filing cabinets, including how they should be cleared before their disposal to avoid the inadvertent inclusion of a confidential file for the second-hand filing cabinet marketplace. Where information critical to the organisation's continued operation is held solely in the heads of its staff, it is almost inevitable that one day this will result in one of the key staff members being ill, having an accident or being otherwise indisposed when a crucial decision or operation is required. Considering the management of the information in staff members' heads is just as important as the effective management of technical systems – some might say more so.

Relationship with corporate governance and related areas of risk management

In recent years the advent of some very-high-profile commercial criminal investigations have resulted in much more stringent and invasive legislation regarding risk taking in companies. Sarbanes–Oxley from the USA, the effects on corporate governance of the Turnbull Report, the Companies Act in the UK and related issues have all had the effect of bringing risk management to the top of the agenda in many boardrooms. It is no longer effective or acceptable (if it ever was) to delegate the responsibility for risk management down to the manager of the IT section.

The proper implementation of effective IA should lie at the heart of all organisations regardless of their sector, size or business. Properly implemented, the secure management of information can provide assurance that risk is being managed effectively in that area at least and can form the firm foundation for further risk management in related areas. If all information is covered by the measures implemented, then the financial, operational, intellectual property rights and a whole range of other risk areas can be managed through the establishment of a single framework.

Information and data life cycles

Information and data have a similar life cycle, and this will be discussed in more detail in Chapter 4.

Security as an enabler delivering value rather than cost

In the information economy in which we all now live, the cost of the loss, corruption, non-availability or unauthorised release of information can be very high. The effective implementation of IA measures can have a very beneficial effect on the potential costs of such events. Thus, it is easy to develop a convincing and compelling business case for the effective management of information through the use of an approved standard and related processes. While it may not be possible to remove the risk entirely, it should be possible to ensure at least that the probability of the risk occurring is significantly reduced or that the effects of the risk materialising are significantly reduced in terms of the business impact.

The use of appropriate countermeasures and contingency plans can also have the very beneficial effect of making the work done by an organisation much more orderly by being based on best working practices. Piles of paper and computer disks left lying around on desks, floors and shelves can be a security disaster waiting to happen. With an IA standard in place, such things should be a thing of the past and the need to spend many hours finding a specific piece of information should be long gone.

With the advent of photocopiers in almost every workplace, the ease with which a sheet of information could be reproduced became very much greater. This in turn meant that where initially there might be only the original and perhaps one handwritten copy of a document to look after, there was now the possibility of many copies to worry about and to try and control. Many a leak from organisations, including governments, has been caused by the proliferation of photocopies,[3] mislaid CDs or inappropriate, perhaps covert, use of USB memory sticks. With improved working practices instigated through effective IA, the need to reproduce information declines, since those who need to see a piece of information can do so easily and in a controlled way through the appropriate use of technology, perhaps without recourse to the production of ever more duplication.

3 An example of which is shown in the film *The Post* (2017) – the true story of how journalists from the *Washington Post* newspaper exposed the American government's ongoing involvement in the Vietnam war, using photocopies of Pentagon papers.

The role of information security in countering hi-tech and other crime

Crime is always advancing and developing, often a little quicker than the enforcement agencies who are established to combat it. The hi-tech industry (covering computers, the internet, digitisation, communications and related areas) over the last 30 years or more has provided criminals with ever-increasing opportunities for more advanced and profitable crime in a wide range of activities. Some crimes are old ones, which have effectively been removed from the criminals' handbook. One example is that of fraud, which had been dealt a severe blow by the introduction of sophisticated security devices in banknotes, passports and the like, but, with the ever-increasing use of the internet, has now returned with increased 'effectiveness'. Emails with 'too good to be true' headings, such as lottery win notifications, have been estimated to be responsible for an overall loss well into the millions of pounds in England and Wales alone.[4] What is commonly referred to as invoice fraud, when a company or organisation is tricked into changing bank account payee details for a sizeable payment, is becoming increasingly common, with ever-increasing sums of money being taken.

Instances such as these are no more than old-fashioned fraud dressed up in new clothes. In addition, the ability to obtain personal information through phishing, key-loggers, screen-scraping or similar tactics has increased the opportunities for criminals to achieve their nefarious purposes. Social engineering helps too; for example, in persuading perhaps more junior members of staff to undertake inappropriate financial activity in a company by apparent pressuring from a supposed senior colleague. Often, simple procedural steps could help to reduce the risk of these crimes – techniques totally separate from the electronic mechanisms through which the crime is committed.

IA can help to address all these issues, at least in the workplace. Good practices at work can also lead to better practices at home, where the proliferation of computers in particular has led to increasing instances of criminal activity targeting the home user. The social duty of companies to help reduce crime overall is well established and setting good work practices with the care of information is an excellent opportunity that should not be missed.

The growth of such crime has increased the importance of forensic investigation, and notably the requirement to preserve evidence based on IT systems. Later in this book this subject will be discussed in more detail, but in recent years it has been ever more evident that the skill of the IT practitioner in the preparation of evidence for trials has needed to develop very considerably from the early days of computing, when IT evidence was rarely used except in the most complex of cases. Now, with internet-crime on the increase and the use of IT becoming the norm for many areas of criminality, the use of investigative techniques based on IT systems has increased enormously. With effectively managed IA high on the priority list for all organisations, these techniques are now a vital piece of the jigsaw of helping to reduce criminality. The IA professional is now a crucial element in the fight against crime, both internal and external to the organisation itself.

4 Crime in England and Wales: year ending March 2019 (ONS.gov.uk).

Ms Jackson, the chairperson of GANT, has asked you to help to develop a sound business case for the implementation of an ISMS. She needs to be able to convince her fellow committee members to authorise the expenditure and so needs to be clear about why this would be a good idea. The key aspect is the balance between the costs of implementing an ISMS against the costs of suffering a serious attack on their information.

Property developers are keen to know where the Natterjack toad can be currently found so they can either avoid buying the land or, if they already have ownership of the land, possibly 'remove' the toad in advance of the planning applications being submitted to 'avoid' any problems with the approvals required. This information is on the website, which has no firewall protecting it.

It would cost GANT many thousands of pounds and several years of effort to reintroduce the toad to a habitat once it has been removed by either natural or man-made effects.

The funding for GANT is through members' fees, grants from other nature conservancy organisations and commercial companies who make donations.

ACTIVITY 1.2

Consider three main areas where the chairperson should gather more detailed information to allow the committee to make reasonable judgements on whether or not it is sensible to carry out the ISMS implementation.

SAMPLE QUESTIONS

1. **If the accuracy of information is a major concern, which of the following would reflect that this is covered effectively?**
   a. Confidentiality.
   b. Integrity.
   c. Availability.
   d. None of these.

2. **When a user logs onto a computer system and is asked for their mother's maiden name, which of the following aspects is the system ensuring?**
   a. Accountability.
   b. Authorisation.
   c. Authentication.
   d. Applicability.

3. **ISO/IEC 27001 is an international standard for information security. Which organisation is responsible for its maintenance?**

   a. The British Standards Institute.
   b. The government of the country in which it has been implemented.
   c. The European Union Standards Committee.
   d. The International Organization for Standardization.

4. **How should the implementation of an information assurance system be seen within an organisation?**

   a. As a problem for the IS department only to sort out.
   b. As a problem on which the senior managers should make a decision but then leave to others to deal with.
   c. As a whole-organisation issue.
   d. As an issue where outside expertise is the best solution.

5. **How should the use of an international standard for information security be viewed by senior managers within an organisation?**

   a. As a good idea if there was the right business environment in which to implement it.
   b. As implementing best practice.
   c. As overkill unless there are very serious problems with assurance.
   d. As the pet idea of the IT director who thinks it will look good to shareholders in the next annual report of the organisation.

2 INFORMATION RISK

Information assurance is almost entirely about the management of risk. The concepts of **confidentiality**, **integrity** and **availability** covered in Chapter 1 are merely areas of risk that must be addressed in an information system's environment. This chapter of the book will examine the component parts of risk – **threats**, **vulnerabilities** and **impact**, and combining threats with the **likelihood** or **probability** that the threat will be carried out, the resulting **risk**. It introduces the basic terminology of risk and discusses the potential threats to, and vulnerabilities of, information systems and the processes for understanding and managing risk relating to information systems.

THREATS TO, AND VULNERABILITIES OF, INFORMATION SYSTEMS

LEARNING OUTCOMES

Following study in this area, you should be able to define and explain each of the key concepts of information risk management and have a thorough understanding of the terminology used.

Threats and threat landscape

A threat is something that may happen that might cause some undesirable consequence. As a simple example, a feasible threat is that an unauthorised person might discover your username and password to a system or service. We won't dwell on the consequences of this just yet – that will be covered under impacts, but it is clear that someone else having knowledge of both your username and password is not a healthy state of affairs.

In order to have any validity, threats must be realistic. They may already have happened to someone else, so there could well be records of such incidents to support the validity of the threat. On the other hand, what might be a threat to one person may well be an opportunity to another. You may care to think about this the next time you try and find an available taxi when it is pouring with rain. To you, there is a very real threat that you will be soaked – to a taxi driver, the combination of the rain and wet pedestrians represents an opportunity!

Threat categorisation

There are a number of types of information-related threats.

Physical threats include deliberate forms of threat, such as theft and vandalism, and also accidental threats, such as trackside communications or signalling cables becoming damaged when a railway train is derailed.

Outages and failures include such things as the absence of vital people resources, which is often overlooked as this is not specifically technology-linked; loss of power supplies, whether due to mains failure or the failure of uninterruptible power supplies (UPS) or backup generators; hardware failures, which are much less common these days, but still possible, especially in rotating disk drives; and software failures – again, less common, but still a valid threat, especially when considering resistance to cyber-attacks. Finally, there will be the threat of human errors, which may result in the loss of confidentiality, integrity and availability.

Hacking and abuse are among the most serious forms of threat. They include social engineering and espionage, which often results in both identity and information theft; malware, such as viruses and ransomware; denial of service (DoS) attacks; and the wider-ranging distributed denial of service (DDoS) attacks. Most (but not all) of these forms of threat originate from outside the organisation. Finally, in this category, are those threats that originate from within the organisation, including eavesdropping, again resulting in identity and information theft; and unauthorised changes both to information and to credentials, such as escalating someone's access privileges.

Legal and contractual threats include the organisation's failure to meet its obligatory requirements in delivering service. While these types of threat may not result in the loss of confidentiality, integrity or availability, there will doubtless be consequences – financial penalties or loss of reputation – that will result. Breaches of legislation such as the Data Protection Act (DPA) or the General Data Protection Regulation (GDPR) will also have potentially serious consequences.

Accidents and disasters may cause information-related problems for organisations. Most of these will be accidental in nature, and will include natural disasters such as floods, landslides, earthquakes and tsunamis, but can also include environmental disasters such as chemical leaks and explosions, such as the events in 2005 at the Buncefield oil storage depot in Hemel Hempstead.

Accidental threats are sometimes referred to as hazards, especially when concerned with external events. The implication is that there has been no deliberate attempt to carry out the threat – it has simply happened. There may be no one to blame for an accidental threat occurring, but there may be a means of dealing with the threat, as will be seen later.

Deliberate threats, on the other hand, occur when someone sets out with every intention of carrying out the threat. This type of threat in the computer world includes hacking, malicious software, sabotage, cyber terrorism, hi-tech crime and so on.

Vulnerabilities

A vulnerability is a weakness; something that, if exploited, could give rise to some unwanted consequence. If you write your password on a Post-it® note stuck underneath your computer's keyboard, this would constitute a vulnerability, as a visitor or other member of staff could easily discover your username, and thereby have complete access to your computer. Many vulnerabilities are not of the user's making. For example, poor software design leaves systems vulnerable to attack – witness Microsoft's® 'patch Tuesday', when patches or fixes for problems, including security vulnerabilities, are released for system administrators to apply.

Whether or not a vulnerability might be expected to be exploited will depend on likelihood or probability, which will be discussed later in this chapter, but often it is the most widely available or widely used software packages and operating systems that are most vulnerable to attack as they present a more easily available or inviting target for malicious-software writers and hackers.

Vulnerability categorisation

Vulnerabilities of IT systems fall into two distinct categories – general vulnerabilities and information-specific vulnerabilities.

General vulnerabilities include basic weaknesses in software (including poor design), hardware, buildings or facilities, people, processes and procedures.

Information-specific vulnerabilities include unsecured computers, including personal computers, hand-held devices and memory sticks, servers, un-patched operating systems and applications, unsecured network boundary devices, unsecured wireless systems, unsecured web servers, unsecured email systems, unlocked filing cabinets and the like.

In recent times, the vulnerability of information leakage from smartphones has become widely known, and many of the applications written for them allow others to access not only the device's store of information, but also its metadata, such as the user's location.

The increasing use of cloud-based services – whether for infrastructure as a service (IaaS), platform as a service (PaaS) or software as a service (SaaS) – means that there exists the further possibility of information leakage due to vulnerabilities in the cloud services themselves. This is particularly important in those situations where the cloud supplier is providing access on a 'multiple tenant' basis.

Use of the so-called Internet of Things (IoT) is widespread, with all manner of devices from fridges and kettles to household alarms and cars being interconnected using the internet as a communications medium. Many of the devices sold as being IoT compatible may have limited (if any) security and are highly vulnerable to interception and attack.

Finally, the increasing use of BYOD, in which organisations allow individual staff members to make use of their own computers and smartphones in order to carry out their daily work should be considered. If not properly controlled and monitored, these can introduce many vulnerabilities.

The point here (for either general or information-specific vulnerabilities) is that the organisation's assets are vulnerable because something has not already been done to secure or protect them, or the 'fix' has been ineffective.

Threats are said to take advantage of, or to exploit vulnerabilities in order to succeed in achieving their goal.

Assets

An information asset can vary considerably in form. An asset does not even have to be tangible, although it could be a system, a database or a building. On the other hand, it could be intellectual property, a business service, an organisation's brand or the reputation of the organisation's chief executive. What is important about assets is that if they are lost, stolen or damaged in any way, the organisation will almost certainly suffer as a result, and if that damage is sufficiently serious, the organisation might never recover.

When impacts are examined, it will be seen that it is always an asset that is impacted by an incident, whether this is tangible or intangible.

Impacts and consequences

The impact (or potential impact) of a threat actually being carried out is perhaps the most important concept of all to grasp. It is usually this potential impact that has to be considered and managed in information assurance. If the impact is small and insignificant, then it may be entirely appropriate to accept the risk and to take no further action other than to monitor it periodically. An example of this might be the failure of 'hole-in-the-wall' cash dispensers – if just one machine in the bank's network fails, the impact would generally be very low.

On the other hand, if the potential impact of a threat was the loss of vital company information, then more appropriate countermeasures would need to be considered. As far as businesses are concerned, the impact on the organisation and its daily activities are usually the crucial considerations and will often warrant further measures being taken.

The business impacts of realised threats include the loss of confidentiality, integrity and availability, and frequently lead to financial loss, inability to trade, brand damage, loss of customer confidence and so on. An example of this is the denigration of his own company's products by Gerald Ratner, the chief executive of Ratners, a UK high-street jewellery chain, when he made a comment to a journalist in 1991. The result was that his comment was displayed in banner headlines across the next day's national newspapers and £500 million was wiped off the organisation's share price within a week. That was actually just the immediate financial impact, but the point here is that nothing tangible was damaged in the exercise; nevertheless the organisation was ruined and eventually went bankrupt just the same.

At the opposite end of the spectrum, immediately following the British Midland air crash in 1989 at Kegworth, England, Sir Michael Bishop, the airline's chairman, commented

that he was not worried about the financial impact of the accident, but that he was more concerned about air safety. This reassured the media and the public and the airline went on to become one of the most successful in the industry.

Likelihood or probability

There are very few certainties in this world, and risk management is no exception. Some things are very likely to happen, while some will be very unlikely to happen. Most others lie somewhere in the grey area between. It is generally accepted that the greater the vulnerability, the more likely an incident is to take place – that is that the threat is actually carried out.

There are two basic ways in which likelihood can be assessed – quantitatively and qualitatively – and these will be discussed in greater detail later in the chapter. In quantitative assessment, there will be clear metrics to calculate the likelihood. These may be derived from previously recorded information including statistical data. In the case of qualitative assessment, the work is more subjective and relies on opinions rather than facts.

For example, companies that produce anti-virus software can point to the large number of viruses that their products can scan for and remove, from which one can conclude that without anti-virus software, the likelihood of infection is high.

On the other hand, one does not need to know the exact number of incidents to be aware that the likelihood of a breach of confidentiality or integrity is high without proper password protection.

Both methods of assessment have their place – the important thing is that likelihood assessments are carried out according to agreed and well understood criteria.

Risk

As mentioned earlier, the result of a vulnerability being exploited by a threat is the occurrence of a risk that produces an impact or consequence. The evaluation of the risk for any particular threat is considered to be a combination of the impact or consequence of the threat being carried out, and the likelihood or probability that it can be carried out.

It is also important to recognise that sometimes there may be a combination of circumstances that lead to further, more serious risks as well. For an unauthorised person to discover your username and password combination is one thing. If your files include a list of other usernames and their passwords, this would lead to further (and potentially more serious) security breaches.

Calculating the overall risk

When assessing the degree of risk for an information asset, all the above factors must be taken into account. As many threats or hazards as possible must be identified and, for each one, the potential impact or consequence of the threat or hazard occurring must be estimated. Finally, any vulnerabilities associated with the asset that will lead

to an assessment of the likelihood or probability that the threat might be carried out must be identified.

The first stage of this is called a business impact analysis (BIA) in which the impact on one or more business assets for each threat can be determined. Once completed, the assessment is made of the likelihood or probability that vulnerabilities might be exploited, allowing the threats to be realised.

From these a risk matrix can be drawn that plots impact against likelihood and gives a formal risk assessment. This will be discussed in greater detail in the next section of this chapter.

Ms Jackson, the chairperson of GANT, has been reminded that all GANT's information is held on a single computer system that was recently compromised by a teenage hacker. GANT has no backup of the information and no suitable paper documentation from which to easily recreate their records.

She has realised that all GANT's information is highly vulnerable and has asked you to assess the consequences of loss or failure of this computer.

ACTIVITY 2.1

Looking at the records and information held by GANT, perform a BIA based on the loss of their main computer system. Some possible impacts to consider are loss of:

- membership details;
- Natterjack toad breeding ground details;
- financial records.

RISK MANAGEMENT

LEARNING OUTCOMES

Following study in this area, you should be able to understand the overall process of risk management, and the appropriate use of controls to enable you to manage risk in a cost-effective and appropriate manner for your organisation.

Risk management standards and frameworks

There are several national and international standards in common use for risk management. These include the national BS 7799-3:2017 and NIST SP800-30 Revision 1 2012, and the international ISO/IEC 27005:2018. While there are subtle differences between them, they all follow the same basic structure regarding risk identification, risk evaluation, risk analysis and risk treatment.

The above standards describe what should be done. In terms of how to go about it, there are a number of risk management methodologies – at a high level, there are Sherwood Applied Business Security Architecture (SABSA) and Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE), while at a lower, less detailed level there are Coras (an open source software tool) and Factor Analysis of Information Risk (FAIR).

Risk management process

Risk management consists of distinct areas: context establishment; risk assessment (the combination of risk identification, risk analysis and risk evaluation); risk treatment; communication and consultation; and ongoing monitoring and review. Risk assessments may take place at a number of levels; for example, across a corporation, a business system or process, or a physical location. While these are somewhat different types of risk assessment, the way in which they are conducted and the way in which the results will be used are essentially the same. This process can be represented as shown in Figure 2.1.

Context establishment

The process commences by understanding what the organisation's information assets are and how they fit into the overall business model. This may also involve developing an understanding of what the business objectives actually are and the organisation's place in the overall scheme of things. This need not be a complex piece of work, but if the following risk assessment is to have value, it is well worth undertaking.

Risk identification

One way of beginning a risk management exercise is to identify the threats. This should be carried out in conjunction with the understanding of any known vulnerabilities. For example, if the assessment is looking at the threat of possible hacking attacks on a web server, the operating system and web server software vulnerabilities should be considered.

Sometimes this will result in the identification of more than one threat, while at other times it will become clear that a number of different vulnerabilities will all be covered by a single threat.

Once each threat has been identified (often more will appear during the process of the work), each one should be considered in the light of its impact on the asset concerned. In the web server example, the threat of a hacker gaining control of the server could potentially result in loss of service – perhaps the failure of an ecommerce facility – loss of customer data, defacement of the web pages and so on, all of which would have a high impact on the company's profitability.

Figure 2.1 The risk management life cycle

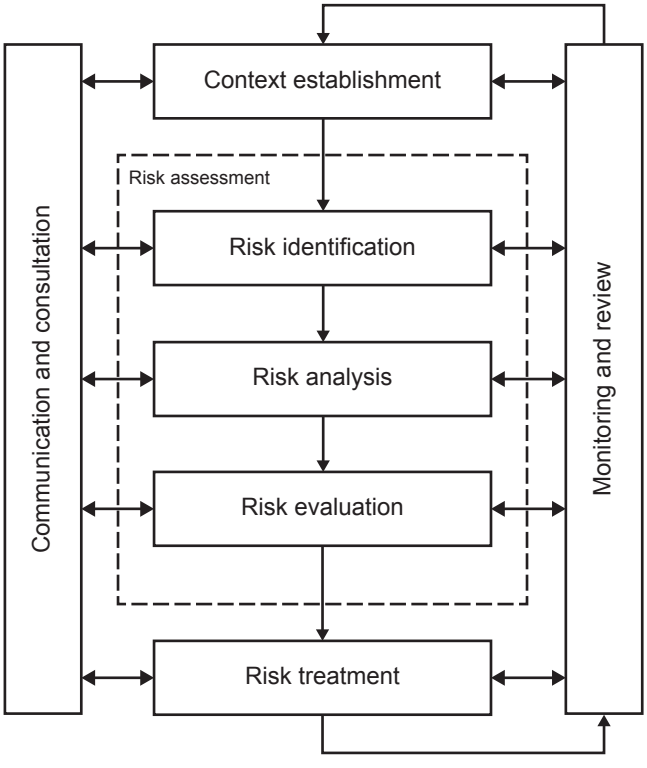

An alternative approach might be to start with a list of the assets that are critical to the organisation, which should have been identified during the BIA, and then determine the potential threats to those assets. In either case the resultant list of assets, their threats and the potential impacts is taken on to the next step of analysis.

Risk analysis

Having identified the impact or impacts for each threat, the next task is to assess the likelihood of each occurring. It is tempting at this point to assume that because the system might be fully up to date with its security patches there is a low likelihood of a threat being realised. However, it must be remembered that this is ongoing work and if the patching falls behind, the likelihood of an attack being successful will increase. It should also be remembered that new vulnerabilities are continually being discovered.

Once this stage has been completed, the risk matrix can be drawn – an example is shown in Figure 2.2.

The matrix itself can be drawn in many ways. It is generally accepted that the simplest form is a three-by-three matrix with High, Medium and Low ratings for both impact and likelihood, as shown in Figure 2.2. An organisation will use the composition of matrix

Figure 2.2 A typical risk matrix

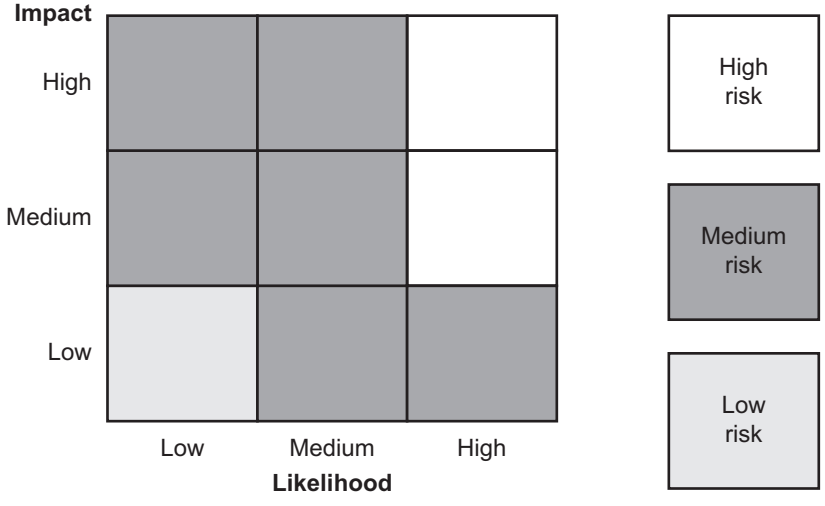

that is most appropriate to its risk programme. Five-by-five matrices are very common and provide a greater level of granularity in the results. However, there is no reason why the matrix cannot contain more ratings for either axis, and there is no reason to have the same number of ratings for each. It is entirely up to the organisation or the individual to decide what size and shape of matrix is appropriate. However, once a particular matrix has been chosen, it should be used throughout the organisation, as this then means that all risk assessments are carried out on the same basis and the matrix is well understood by all involved.

The output from the matrix will be a number of risk levels. Again, these are arbitrary and can be agreed based upon the organisation's 'risk appetite' (this is the degree of risk an organisation is prepared to accept). Organisations that have a low risk appetite include, for example, some aspects of the work of pharmaceutical companies where the introduction of new drugs sometimes requires years of rigorous testing before a product is considered safe enough to launch. Organisations that display a high risk appetite, on the other hand, include some of the work of petrochemical companies, who will spend tens of millions of pounds in searching for scarce oil reserves and will often drill many 'dry' holes before finally finding a rich source. The highest combination of impact and likelihood give the highest level of risk, and these are risks that should be treated as soon as possible. Those lower down the matrix have a lesser degree of urgency, and those risks that carry low impact and likelihood may, if the organisation so decides, be accepted without the need for treatment but should still be monitored over time to ensure that the risk does not increase.

Whatever the risk, the assessment for each threat should be recorded on a risk register, which will include details of the impact and likelihood for each threat, the level of risk calculated, possible treatment options, who is responsible for carrying out the risk treatment and a date by which the work should be completed. It is also considered good practice to note a review date for each risk as ongoing monitoring will show whether

either the impact or likelihood of any threat has changed since the last assessment. It may also be that other factors have had an effect that increases or reduces the threat, likelihood or impact of the risk. It is common practice to do a second assessment of the risk once the mitigating measures identified have been put in place.

Risk treatment

Having decided from the output of the risk matrix the priorities in which to treat the risks identified, a risk treatment plan must be produced. This will be dealt with in greater detail in the following section, and allows for four basic choices utilising what are often known as strategic controls:

- to avoid or terminate the risk completely – often by not doing something that might incur an unacceptable level of risk;
- to reduce or modify either the likelihood or the impact of a risk – usually by some form of risk mitigation;
- to transfer or share some parts or all of the risk – for example, by insuring against the eventuality;
- to accept or tolerate the risk – this is a common option when the assessed level of risk is low.

Communication and consultation

It is essential throughout the entire risk management process that those conducting the work maintain good communications with other parts of the organisation, especially those who are actually responsible for the assets in question and who may eventually own the responsibility of agreeing the form of risk treatment, funding the necessary work and managing the work to completion.

The asset owners will always be the most appropriate people with whom to consult when carrying out the BIA, but, at the same time, they may not have an in-depth understanding of the vulnerabilities, in which case a different group of experts will need to be consulted.

Monitoring and review

The final stage of the risk management process is to monitor the results of the risk treatment work. The frequency of this process may vary according to the type of threat – some threats (and therefore the risks they represent) may change very quickly and will require monitoring at more frequent intervals, while others will change little over long periods of time and will only need occasional monitoring.

The overall risk management process should be repeated over time, as some threats might disappear completely and new threats might emerge. Again, the interval will depend largely upon the risk appetite of the organisation and may well be documented in a risk management strategy or policy document.

Options for treating risks

The output of the risk matrix will determine one of the four courses of action mentioned above in order to treat the risks. These are sometimes referred to as **strategic controls**,

as they provide a high-level approach to risk treatment, leaving the detail to tactical and operational controls, which are discussed later.

Avoid or terminate the risk

Risk avoidance, sometimes called termination of the risk, is usually a fairly clear-cut option. Put simply, it means not doing something that incurs risk. For example, it might be to issue a security policy that states that users of personal computers must not install unauthorised software. This removes the risk of inappropriate software finding its way onto PCs within the organisation, and can be enforced by restricting the administrative capabilities of users. Another example would be not to do, or to stop, a business activity because it put the organisation at too much risk.

Reduce or modify the risk

There are basically three possibilities here: to reduce the threat; reduce the vulnerability (and thereby the likelihood of it occurring); or reduce the impact if it does occur. Actions that take place in reducing the risk are usually referred to as controls.

Reducing the threat can be difficult. For example, although it would be nice to get rid of hackers completely, this would require significant social reform and is therefore an unlikely option.

Reducing the vulnerability or likelihood is a possibility. For example, by applying appropriate security patches to an operating system or tightening the security settings on a firewall, the likelihood of hackers gaining access is reduced, although not removed.

Finally, it is also possible to reduce the business impact if a risk does materialise. For example, if the whole of an organisation's information assets reside on one main system, this would represent a potential single point of failure, and could be mitigated by introducing either a physical or virtual data separation system, or a disaster recovery standby system.

Transfer or share the risk

Risk transfer can be achieved in several ways, but typically an insurance policy is an appropriate method when the impact of the risk can be measured in purely financial terms. Another means of transferring risk is to move it to a third party when the relevant expertise to manage the risk is not available within the organisation. An example of this might be when magnetic media containing sensitive information requires secure disposal and the organisation outsources the work to a specialist company. In cases such as this, however, the organisation itself must still retain overall responsibility and ownership of the risk.

Like information assurance, insurance is also all about the management of risk – it is something in which insurance companies have specialised for many years. You pay the insurance company to do the worrying for you and, in return, they expect you to take reasonable care to ensure that the worst does not happen. The insurance policy does not necessarily alter the threat or the likelihood of the risk occurring, but, if it does occur the impact on the organisation is reduced by the payout charges from the insurance company. The cost to the organisation is that of the insurance premium itself, plus any excess.

It is highly unlikely that transfer can remove the full impact of a risk. It is more likely that the risk is ultimately shared between the two organisations with a formal acknowledgement of the responsibilities. For example, a travel insurance policy will reduce the financial impact on you if your flight is seriously delayed, but it will have little effect on the consequences for your company of not making the critical business sales meeting where you lost a major contract. Indeed, it is usual for these so-called consequential impacts to be explicitly excluded from such policies.

Accept or tolerate the risk

When the level of risk assessed is very low, the organisation may decide that it is willing to live with or tolerate the risk and that it will be accepted. It is important that this is a conscious decision and that one or more individuals are held accountable for it by means of a formal sign-off process and by keeping a record in the risk register as to who has signed as accepting the risk. In some cases, the cost of treating the risk may equal or exceed the potential financial impact if the risk occurs. In this case, the business must decide whether to accept the risk or to implement the appropriate controls anyway, maybe as a means of maintaining customer confidence.

Acceptance of a risk is not the same as ignoring it – this must never be a course of action, as risks that are ignored can cause problems at a later date. It is also important that risks which have been accepted are monitored and reviewed at suitable intervals in case the impact or likelihood has changed since the initial assessment was carried out.

Most aspects of risk treatment involve a cost element of some sort – this must be balanced against the potential losses that might be incurred if the threat should be realised. Where the level of the cost of the mitigation approaches or exceeds the potential losses, the decision to accept the risk may well be the correct decision.

Other types of risk controls

Tactical risk management controls come in four types:

- Detective controls, which are designed to identify information security incidents, such as intrusion detection systems.
- Preventative controls, which are designed to stop an incident from taking place; for example, the configuration of firewall rules that prevent users from accessing banned websites.
- Corrective controls, which, having identified an information security incident, will make appropriate changes to ensure that it does not cause an impact. A typical example of a corrective control is that of anti-virus software, which having identified a virus (for example as an attachment to an email), will block it and perhaps remove it to prevent the virus from promulgating further.
- Directive controls (also sometimes referred to as personnel controls), which are intended to inform users regarding things they may and may not do. An example of this would be a clause written into an employment contract that dictates fair use of the internet and makes clear the possible penalties for abuse.

There are three types of **operational control**:

- Physical controls place some form of device in between the organisation's assets and possible intrusion; for example, securing access to restricted areas such as data centres by means of a card or token-based access control system.
- Procedural controls are intended to guide users in the correct way of undertaking their work. These may include process and procedure documents, standards, guidelines and regulations.
- Technical controls are based on both hardware and software solutions in order to ensure that risks are reduced or avoided. These might include firewalls, intrusion detection systems and activity logging.

Once the organisation has mitigated the risk as much as possible, there may be some residual risk remaining. Generally speaking, this will fall below the level of the organisation's risk appetite, and can safely be accepted, with the proviso that it will be monitored over time. However, some risks may still fall above the risk appetite level, but may be too costly or difficult to address in other ways. These too will have to be accepted, but in this case with a much shorter interval between reviews.

Approaches to risk assessment

Qualitative risk assessment

As mentioned earlier, one of the approaches to carrying out risk assessments is to use a qualitative method. While this is essentially subjective, it may be the best course of action when hard facts relating to impacts and frequency of events are hard to come by.

The fact that the method is largely subjective, however, should not prevent the assessment from being carried out properly. The main thing to agree is what constitutes 'high', 'medium' and 'low' ratings so that any assessment will have a degree of rationality about it, making it easy to understand and straightforward to justify later on.

Alternatively, a 'standard' template can be used. For example, the UK Government has a standard method of risk assessment for use in civil contingency work, as set out in Table 2.1.

Table 2.1 One possible rating framework for risk assessment

Rating	Impact	Likelihood
1	Insignificant	Negligible
2	Minor	Rare
3	Moderate	Unlikely
4	Significant	Possible
5	Catastrophic	Probable

Quantitative risk assessment

Quantitative risk assessments, on the other hand, take a much more factual approach and can use statistical evidence to support both impact and likelihood assessments. For example, when assessing the risk of virus attacks, there will be plenty of numeric information available on the websites of anti-virus vendors to provide the basis for supporting a metric-based assessment. Whether their figures are to be wholly believed is of course another matter entirely.

Semi-quantitative risk assessment

Alternatively, whilst trying to assess the impact of failure of a system or service, and having no hard facts to support the assessment, it might be decided that up to £100,000 constitutes a low impact, between £100,000 and £1,000,000 constitutes a medium impact, and above £1,000,000 constitutes a high impact.

Statistical information to support likelihood assessments is also very likely to be widely available and should also be treated with caution. The expression 'There are three kinds of lies: lies, damned lies, and statistics' has been attributed to many people, including Benjamin Disraeli and Mark Twain, but it remains true to this day. If you wish to explore the subject further, there are two books listed in the reference section of this book that may be of interest (Bernstein, 1996, 1998; Salkind, 2004).

Software tools

Unsurprisingly, there are a number of software tools available that will help in carrying out risk assessments. This book does not aim to offer any specific guidance on the good, the bad or the ugly of these, but merely to state that they exist, and that they should be investigated in greater detail to discover which (if any) are best suited to the need.

Again, a word of caution. It is very easy to become obsessed with choosing the right software tool and working through a complex set of analyses, only to find that the answers are not as required. Very often a simple analysis tool can be created using a spreadsheet and is therefore very much easier to tailor to the needs of the organisation. Try to keep the work as simple as possible and reduce the impression of a 'black art' by making the results understandable to as wide an audience as possible.

Questionnaires

Ultimately, when conducting a risk assessment, it will be necessary to visit various areas of the organisation seeking information from people who understand far more about their particular area of the business. It is worthwhile therefore spending some time in preparing a questionnaire that will guide them through a series of questions designed specifically to discover exactly the information required in order to carry out the risk assessment. Working with a questionnaire also helps to ensure that that there is a level of consistency across the answers provided.

It is usually best to begin with open questions – for example, asking people to describe the processes and procedures by which things happen – as this information will often point to the need for further questions. For example, it may help to begin by asking for an explanation of what the person's department does; what the inputs and outputs are; what processes are involved; who carries out the work; where they do this; what

happens if one or more inputs ceases to work; and so on. This will very often highlight a 'single point of failure' in the process. Closed questions can then follow, drilling down into the detail and uncovering facts and figures that will help to build up a more detailed picture, and that will facilitate the production of a more reliable analysis of what might go wrong and how likely this might be.

While some of this information may seem unimportant at first, it should be remembered that at some stage a business case will have to be presented to the board in order to gain approval for funding to mitigate the most serious risks. This information will almost certainly play a key part in building the business case.

Identifying and accounting for the value of information assets

Before carrying out any form of risk assessment on an organisation's information, it is obvious that each of these 'information assets' must be identified and documented in a BIA. Much of the information to do this will come from the questionnaires referred to above, so it will be useful to list who is responsible for collecting and storing the information, where it is held, how and when it is used and backed up, and so on. On occasion the people themselves could be considered an information asset if, for example, they are the only source of business-critical information or if they have unique knowledge or skills within the organisation.

The value of each of these information assets will depend very much on its function, how long the business can manage without it, how long it would take or how difficult it would be to recover or restore it and how frequently the information changes. One of the key questions to ask when assessing the information value is 'How much will the organisation lose (or not make) if the asset is not available?'

Clearly if this is a human resources database, loss of access to it for a short period of time should not pose a serious threat – the impact would be low – but loss of a database holding online customer orders on an ecommerce website, even for a few minutes, would have a much higher impact.

Information classification policies

The value of information assets links neatly into the subject of information classification. Some information held by an organisation (for example, a product list) will be considered to be public domain information and will be allocated a low classification – often referred to as 'unmarked' or 'unrestricted'. Other information will be more strictly controlled – for example, a list of customer accounts and their annual spend must be kept within the organisation and will therefore have a higher level of classification, such as 'confidential'.

More critical information will have a higher level of classification again – for example, documents relating to a merger or acquisition will not be available to many people within the organisation, perhaps only at board level and a very few senior managers. These might be graded as 'highly confidential' or 'secret'. There is no limit to the number of classification levels that a company can use, although simplicity is again the key here. Three or four levels is generally considered to be about right.

Each information asset should be categorised according to the classification policy, and those assets not graded as 'unmarked' or 'unrestricted' must be protectively marked to indicate this. The classification policy should also identify the procedures for handling, storing and disposing of protectively marked information.

The need to assess the risks to the business in business terms

While it is very straightforward (after some practice) to carry out risk assessments, there will be a great temptation to describe and document these in risk management terminology. This is fine when discussing the assessments with like-minded or similarly experienced people, but when it comes to selling the concept back into the business, this terminology may not be well understood if people are unfamiliar with the jargon. Terminology that is alien to the recipient will diminish the effectiveness of the risk assessment and may make it more difficult to convince the reader that appropriate action must be taken.

It is always advisable for the risk assessor to be able to express the outcome in terms that are readily understood by managers within the business – in other words, to talk their language. This may mean that some risk assessments must be 'translated' for the benefit of different departments in the organisation. For example, different terminology is used in an organisation's production and marketing departments, so the output of the risk assessments must be adjusted so that the language used reflects their own specific terminology. This should not be seen as a patronising approach, but more as a pragmatic method of optimising the results to gain the maximum impact and buy-in. Likewise, the assessments themselves must focus on the areas that the individual departments recognise and to which they can relate, or the exercise will have been wasted time.

Balancing the cost of information security against the potential losses

Once the results of the risk assessments have been made available, there will be recommendations as to how the higher-level risks should be mitigated. While the organisation would not expect a detailed cost estimate to carry out the remedial work at this stage, it would be prudent to have a rough idea at least of the likely order of expenditure, resources required (especially people) to undertake the work and the approximate timescales. In this way, it is possible to present the results of the risk assessments in a more balanced way so that the decision-makers can take a more objective view. For example, if the anticipated losses as the result of a threat being carried out are £50,000, the overall risk may have been assessed as medium. If the costs of reducing this to a lower level will amount to £25,000, the decision might well be to accept the risk rather than reducing it, as the cost of the control is relatively high in comparison to the possible impact.

On the other hand, if the anticipated losses are £1,000,000, the risk has been assessed as high and it will cost £25,000 to reduce this to medium or low, then the decision to reduce the risk is much easier, as the balance is more in favour of risk reduction. Finding the balance largely depends upon the organisation's risk appetite, but in some cases (where the cost–benefit balance is not as clear-cut) the decision will be more difficult to make and may require a more detailed cost breakdown. The experienced risk

manager will recognise cases such as these and will be prepared for them. However, there may be circumstances in which the cost of treating the risk is not the main issue and other factors, such as legal and regulatory requirements, mean that the work has to be carried out regardless.

The role of management in accepting risk

The option to accept risk may sound an easy one to take, but it is not something that should be done lightly. Many organisations are unable to differentiate between accepting risk and ignoring risk (which is never an option). If the recommendation is to accept a risk, then the decision to do this must be a conscious one and should be fully documented. Although it is common practice for a single manager to 'sign off' a risk, when the impact is high it is better practice to have a second manager sign off as well – preferably one who is more remote from the risk itself but nevertheless one who has a good understanding of the potential impact of the risk materialising.

For example, if a production manager signs off the risk of having only one machine of a particular type, a manager from an entirely different discipline (say, finance) should counter-sign the risk in order to provide an objective confirmation that acceptance is in order. This reduces the possibility of individual departments covering up their own mistakes.

Once 'signed off', an accepted risk should still be revisited at regular intervals in order to verify that the threat, the impact or the likelihood have not changed, and that acceptance of the risk continues to sit well with the organisation's risk appetite.

Contribution to risk registers

In 1999 the Turnbull Report established best practice for UK listed companies with regard to internal control, including risk management. This has been updated since then, most recently in 2014 by the Financial Reporting Council's Risk Guidance (FRC, 2014). Risk registers are a vital part of the overall risk management process. They achieve several objectives:

- They permit all risks identified in the risk assessment process to be documented in a formal manner, sometimes a legal requirement.
- They allow an authorised observer (e.g. an auditor) to have visibility of the impact and likelihood of the risk and all the associated details and to assess the suitability of the responses selected.
- They allow ongoing monitoring of the status of the risk and can be used as management reports on the progress of risk mitigation and of any variation in the risks.

A risk register should contain as a minimum:

- the details of the threat;
- its assessed impact and likelihood;

- the overall risk evaluation calculated from these;
- the recommended treatment (accept or tolerate, avoid or terminate, reduce or modify, transfer or share) and the actual action(s) to be taken;
- the person or department responsible for carrying out this work and the date by which it is expected to be completed.

Other fields may also be included, but those listed above form the basic minimum information required of a risk register. It is common practice to review and update the risk register at intervals – typically monthly or quarterly.

You have delivered your impact analysis to Ms Jackson of GANT and, although she now understands the consequences of loss or failure of the computer, she needs to understand the likelihood of the threat actually occurring.

She has asked you to carry out a risk assessment based on the threats you have already identified.

ACTIVITY 2.2

For each of the threats you identified in Activity 2.1, give an assessment of the likelihood of each taking place.

From the impact analysis carried out in Activity 2.1 and the likelihood assessment above, calculate the overall level of risk for each threat.

Draw a simple 3 × 3 risk matrix and illustrate the risks you have assessed.

SAMPLE QUESTIONS

1. **What are the four types of strategic risk treatment that can be used?**
   a. Accept, transfer, ignore, control.
   b. Avoid, ignore, transfer, mitigate.
   c. Accept, avoid, reduce, transfer.
   d. Reduce, transfer, mitigate, control.
2. **A business impact analysis considers which of the following?**
   a. The consequences of a threat being carried out.
   b. The likelihood of a threat occurring.
   c. The likelihood that a vulnerability will be exploited.
   d. The probability that losses might result from an incident.

3. **A risk assessment is designed to achieve which of the following?**
   a. To identify the likely impact if a vulnerability is exploited.
   b. To identify the degree of likelihood that a vulnerability will be exploited.
   c. To identify the likely impact if a threat occurs.
   d. To identify the degree of likelihood that a threat will occur and its likely impact.

4. **Which of the following is *NOT* a threat?**
   a. Failure of the local mains power supply.
   b. An easily guessed password.
   c. A transmission circuit cable break.
   d. Flooding of a data centre.

5. **Once the key risks have been assessed, what action is unacceptable for very low risks?**
   a. They can be ignored.
   b. They can be accepted.
   c. They can be treated.
   d. They can be terminated.

REFERENCES AND FURTHER READING

Publications

Bernstein, P.L. (1996, 1998) *Against the Gods: The Remarkable Story of Risk*. John Wiley & Sons, Inc.

FRC (2014) *Guidance on Risk Management, Internal Control and Related Financial and Business Reporting*, Financial Reporting Council, www.frc.org.uk/getattachment/d672c107-b1fb-4051-84b0-f5b83a1b93f6/Guidance-on-Risk-Management-Internal-Control-and-Related-Reporting.pdf

Hiles, A. and Barnes, P. (1999) *The Definitive Handbook of Business Continuity Management*. John Wiley & Sons, Inc.

Salkind, Neil J. (2004) *Statistics for People Who (Think They) Hate Statistics*. SAGE Publications.

The Business Continuity Institute (2010) *Business Continuity Management: Good Practice Guide*, November.

The Institute of Directors (2000) *Business Continuity*. Director Publications Ltd.

Toigo, J. (1996) *Disaster Recovery Planning for Computers and Communication Resources*. John Wiley & Sons, Inc.

Websites

American Society for Industrial Security (ASIS) Business Continuity Guidelines: https://asisonline.org/guidelines/guidelines.htm

British Standards Institute: www.bsi-global.com

Business Continuity Institute: https://thebci.org

Continuity Central: https://continuitycentral.com

Continuity Forum: https://www.continuityforum.org

Disaster Recovery Institute International: https://drii.org

Global Continuity: www.globalcontinuity.com

Institute of Risk Management: www.theirm.org

Virtual Corporation (BCMM): www.virtual-corp.net/

3 INFORMATION SECURITY FRAMEWORK

The purpose of establishing an information security framework is to ensure that appropriate control mechanisms are in place to manage effectively the information assurance across the enterprise.

This chapter covers the basic principles for establishing such a framework within an organisation and will look at the general area of information security management. In particular it will consider the role and appropriate use of policy, standards and procedures, information assurance governance, security incident management and their appropriate implementation.

ORGANISATION AND RESPONSIBILITIES

LEARNING OUTCOMES

The aim of this section is to provide you with the basic knowledge needed to understand the principles for organising information assurance across the enterprise. Once completed, you should not only be able to define and explain the main concepts but also draft documents to meet the general requirements in the following areas.

The organisation's management of security

Establishing an organisational structure to manage information assurance provides a framework to ensure that the assurance requirements of the enterprise are understood and that responsibilities are allocated appropriately across the enterprise to achieve this. Accountabilities need to be clearly defined, whether at an enterprise level or on a local basis, and assurance activities need to be co-ordinated appropriately across the organisation to ensure that they are being managed effectively. This section covers the necessary organisational arrangements that should be carried out to provide effective control of information assurance.

Information security roles within the enterprise

There should be a nominated resource within the organisation that has responsibility for the day-to-day management of information assurance issues. This is to ensure that good information assurance practice is applied properly and effectively across the

enterprise and for co-ordinating all assurance activities. In larger organisations, this function should be a full-time role and the manager of this function is often referred to as the head of information assurance, the information security manager or the chief information security officer (CISO). In smaller organisations, the role may be combined with other responsibilities.

In this section the role will be referred to as the information security manager. How this role is structured will depend largely on the size and culture of the organisation and is typically supported by other people as part of a dedicated team.

The information security manager needs to understand the information security risks that the enterprise may face, what controls are in place and where the enterprise may be vulnerable. This information must be communicated effectively to senior management (who have ultimate responsibility for information assurance). This is to ensure that they understand the status of assurance within the enterprise so that the appropriate safeguards are put into place. The main activities of the information security manager are:

- co-ordinating information assurance activities across the enterprise, including those delegated outside the team;
- co-ordinating the production of security policy;
- communicating with users so they understand their information assurance responsibilities and are aware of potential threats to the enterprise;
- understanding the enterprise's risk appetite and profile and how it may be evolving;
- monitoring the effectiveness of the enterprise's assurance arrangements;
- reporting on the effectiveness of the assurance arrangements to senior management and suggesting improvements;
- providing expert advice on information assurance matters to the enterprise;
- creating a culture of good information exchange and assurance practices.

There are a number of recognised standards that provide guidance on how to manage assurance arrangements and responsibilities within an enterprise, such as the ISO/IEC 27000 series and the Information Security Forum (ISF) Standard of Good Practice.[1] These standards can be adapted to fit individual enterprise requirements.

Placement in the enterprise structure

Placement of the various assurance roles within an organisation will normally depend on the structure, the particular requirements and the culture of the enterprise. Therefore, there are no definite hard and fast rules as to where the roles should sit specifically, how they should be organised or what their scope should include.

1 The latest ISF document at time of going to print, *The Standard of Good Practice for Information Security 2018*, can be found here: https://www.securityforum.org/tool/the-isf-standard-good-practice-information-security-2018/

In some enterprises, the information assurance function is located within the corporate compliance area. This is common in enterprises or industries that have a strong compliance culture, such as banking or manufacturing.

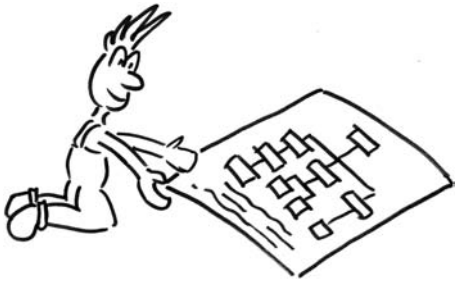

In other enterprises the function is based in the IT group because many (but rarely all) of the controls to protect the enterprise are reliant on computer technology. Sometimes, the function can be placed within a central facilities group since assurance responsibilities often span a number of management areas within an enterprise. To work effectively, reporting structures should include dotted line responsibilities to roles including, but not limited to, the chief risk officer (CRO), the senior responsible owner (SRO), chief information officer (CIO), senior information risk owner (SIRO) and the chief finance officer (CFO).

The scope of the information assurance function may vary. In some instances, the assurance function may include responsibility for setting policy and direction but not for the actual implementation of the security control mechanisms, which then may be carried out by a separate area such as the IT department or local teams. Alternatively, the assurance or audit function can also have responsibility for the implementation of technical security controls and solutions and for conducting investigations and monitoring compliance.

As a regulatory role, the information assurance function should be positioned as part of a formal structure so that it can facilitate the full management and co-ordination of assurance matters across the enterprise.

Board/director responsibility

One senior person within the organisation should be given the overall responsibility for protecting the assurance of the enterprise's information assets and be formally held accountable. This role should be performed by a board member or equivalent to demonstrate the enterprise's management commitment to information assurance. In some organisations, the CISO, or equivalent, is a board member. The main responsibility of this individual is to ensure that appropriate assurance controls are implemented across the enterprise and to:

- provide a single point of accountability for information assurance;
- ensure that assurance goals are identified and meet the enterprise's needs;
- ensure that adequate assurance resources are made available to protect the enterprise to an acceptable and agreed level of risk;
- assign specific assurance roles and responsibilities across the enterprise;
- provide clear direction, commitment and visible support for assurance initiatives, for example by approving and providing sign-off for high-level security policies, strategies and requisite architectures.

The director has the necessary status to ensure that appropriate focus is placed on protecting the enterprise's information assets and to influence and sanction assurance activities. Implementing adequate assurance control mechanisms can in some cases be met with resistance from other parts of the business that are in competition for available resources, and therefore it is important to have a 'security champion' to ensure that priorities are met. Experience has shown that if senior support is not in place, assurance initiatives will probably fail.

There is an increasing quantity of UK and worldwide legislation and regulation that demands this level of accountability and responsibility, for example Sarbanes–Oxley (USA) and the Companies Act (UK). The Turnbull Report in the UK states that a board member for a public limited company has to be responsible for ensuring adequate service continuity requirements are in place to prevent the enterprise going out of operation after experiencing a major problem. The more recent legislative addition has been the European Union's GDPR incorporated into UK law alongside the Data Protection Act of 2018 that tailors how the GDPR applies in the UK. Another piece of European legislation, the Network and Information Systems (NIS) Regulations 2018, has a wide-ranging impact on organisations who are classed as 'operators of essential services' (OESs) and 'relevant digital service providers' (RDSPs).

If these measures are not appropriately implemented, that responsible person could perhaps face a custodial sentence or, in the case of GDPR, huge financial penalties that could be disastrous for an organisation. These potential outcomes may be good for focusing the attention of senior management and for securing appropriate resources. The director should establish and chair an ongoing high-level working group to co-ordinate assurance activities across the organisation to ensure adequate assurance measures are in place to protect the business to an acceptable and agreed level of risk, often referred to as the organisation's risk appetite.

This working group is often called a steering committee or a security forum. The working group should be made up of a cross-section of individuals from the enterprise that are either stakeholders in the requirement of good assurance or have responsibilities for ensuring appropriate assurance arrangements are in place. Membership should include one or more line of business (LOB) managers or departmental heads to ensure that assurance arrangements meet their business or organisational demands. It should also include the information security manager and representatives from vested interest parties such as internal audit, personnel (HR), physical security and the head of IT. In a world with increasing aptitude for outsourcing, it may be appropriate to include in the working group representatives from the key outsourcing partners to be present at meetings or perhaps at those non-commercially sensitive parts of meetings.

The group should meet regularly in order to ensure that the protection of the organisation's information is being managed effectively and that controls are in place to reduce risk to an acceptable level. This includes:

- ensuring that assurance is included in the enterprise's overall planning activities;
- approving and prioritising assurance improvement activities;
- reviewing assurance performance and changes in threats to assess whether the risk profile of the enterprise has altered;

- ensuring that all legislation and regulatory requirements are being met in an appropriate and effective way;
- approving policies, standards and procedures that relate to information assurance;
- acting as evangelists for assurance within the organisation by emphasising its importance to colleagues.

The director will normally delegate authority for the development of information assurance initiatives and responsibilities either to individuals within the working group or to other members of the organisation. However, the director will be ultimately accountable for achievements or failures. The working group is the cornerstone of the information governance structure that will be covered later in this chapter.

Responsibilities across the organisation

Achieving good information assurance requires teamwork and a wide variety of skills ranging from managerial to technical and administrative. It is unlikely that any one person would have all the requisite skill sets or even the time to perform everything that is required; therefore, roles need to be delegated to the appropriate teams or to specific individuals with the necessary skills. For instance, the skill sets required to maintain an enterprise's anti-virus systems are different from those required for administering user identities.

All those involved need to have a proper understanding of accountabilities and be given clear direction and support from senior management to achieve what is required of them. In many cases, individuals may be working together as a 'virtual team' that spans across separate management responsibility areas. For many individuals, their information assurance responsibilities will form just a part of their overall role. Therefore, their activities require co-ordination and monitoring from a central information assurance function to ensure they are successful. It is essential that all individuals have clearly defined responsibilities and that they understand their part in delivering the overall information assurance function within the enterprise. Their job descriptions or terms of reference should include:

- the scope of their responsibilities and their level of authority;
- the processes they should be following to carry out these responsibilities;
- the procedure they should carry out to report and deal with any security breaches that they discover;
- understanding confidentiality/non-disclosure constraints;
- requirements for regular reporting;
- what should take place if they leave the organisation;
- what will happen if they breach the agreed terms and conditions.

These responsibilities should reflect the enterprise information assurance policy and current legislation. They should be reviewed regularly to ensure that they remain current, are relevant and are supplemented with additional guidance as necessary. Where the information assurance function does not constitute a full-time role, it is important that

those engaged in carrying out assurance activities are given a clear mandate by senior management to do so and that the work is included as a formal part of their objectives. These individuals must have the sufficient skills and tools to be able to carry out these tasks and may need training and support to acquire the necessary knowledge and fully appreciate the critical nature of protecting information assets.

Many enterprises have local security co-ordinators that are the 'eyes and ears' at a local level to ensure that security policies are followed and for identifying any security vulnerabilities or breaches. They can offer feedback to the information security manager as to whether existing assurance processes and controls are effective, identify possible risks and help propose new assurance controls. The scope of the local co-ordinators will vary depending on the enterprise's requirements. In a large global operation it may be appropriate to have country or regional co-ordinators, or it may be more relevant to have an individual responsible for a business unit or office location. In the case of GDPR, it is a requirement that an organisation has a designated office or point of contact within the European Union (EU) if they come under the jurisdiction of the legislation. In smaller businesses it may be better to nominate individuals who have responsibility for a specific department or business function.

Anyone who has access to the organisation's information assets will have a level of personal responsibility for its assurance and it is important that these are known and understood. User responsibilities need to be clearly set out in an acceptable information use policy and bolstered by education to that they can help to protect against risk. Guidance on acceptable use policies is described in the next section and user awareness and training is described in more detail in Chapter 5. Individuals may have specific responsibilities for a particular application or system and in this case their responsibilities are best expressed in the system operating procedures. Third-party assurance responsibilities, particularly with outsourcing arrangements, should be included in contractual terms and conditions and should include appropriate auditing and monitoring arrangements.

Any information assets within an organisation should be associated with an owner of that information (i.e. head of department, business manager or process owner) who understands its importance to the enterprise and the resulting negative impact if its confidentiality, integrity or availability is compromised. This will help to ensure that adequate controls and procedures are put into place.

Statutory, regulatory and advisory requirements

External factors can influence how an enterprise's information assurance should be managed and these requirements need to be understood so that the appropriate assurance controls can be adopted to enable the business to fulfil its responsibilities. These requirements can arise from a variety of organisations such as the police, utility companies, government, trade regulatory bodies or telecommunications suppliers. They may be statutory, regulatory or advisory.

Statutory requirements are legal requirements that must be fulfilled. For example, law enforcement agencies must be contacted should certain laws be broken or are suspected of being broken – the download of child pornography would be such a case in many countries. Compliance with these requirements may influence how an enterprise's

incident reporting procedures are organised. For example, how, when and by whom should the authorities be contacted. Privacy legislation such as the GDPR will influence how information is stored and managed within the enterprise and how resources are deployed to ensure that the enterprise complies with this legislation.

Regulatory requirements are often imposed by trade bodies, and these specify how an enterprise should operate to conform to certain standards. Although they are not legal obligations, regulatory bodies have extensive powers and failure to comply could lead to possible fines or, in extreme cases, exclusion from trading in a particular environment. The finance sector is a good example of this as it maintains strict controls to prevent financial malpractices such as fraud or money laundering – official bodies, such as the Financial Conduct Authority (FCA) within the UK, have far-reaching powers. Another example of a regulatory authority in the UK with significant powers is the government agency on health and safety in the workplace. In terms of information, the Information Commissioner's Office (ICO) in the UK is the lead body for setting standards and then enforcing them across all sectors and they also act in an advisory role to other regulatory bodies (some termed 'competent authorities' under the NIS regulations) who have to concern themselves with information assurance. Some of these regulatory requirements support or supplement statutory requirements in certain business operations.

Advisory requirements may arise from government agencies or utility companies and provide advice as to what arrangements should be put into place to help cope with instances such as fires, natural disasters and acts of terrorism. These requirements are not legally binding and are generally issued to help encourage best practice.

Maintaining relationships with relevant external bodies is beneficial to an organisation as it helps the enterprise to better appreciate the requirements placed on them and gain prior warning of any changes. By understanding the requirements of the emergency services, utility companies and government agencies, enterprises can design their contingency plans and incident management processes and procedures more effectively.

Provision of specialist information security advice and expertise

Those involved in the security function should provide specialist security information advice and expertise to the enterprise. A high degree of current knowledge on information assurance matters should be maintained on topics such as awareness of industry trends, changes to organisational threats, new control measures, analysis of risk, legislation and compliance requirements and the latest technological developments. It is an ongoing process. It is not necessary to have all the answers, but it is essential to be in a position to know where to find this information or to have access to someone with this specialist knowledge as and when needed.

One way of achieving this aim is to keep in regular contact with special interest groups and websites or by networking with information assurance peers in other enterprises via professional associations or security forums. Information about new technologies, products, threats and vulnerabilities or how to tackle particular assurance issues can be shared with one another and often a collaborative approach is useful in understanding and addressing these issues before applying them to relevant situations. Bulletin boards, websites and news groups also can provide early warnings of possible alerts,

attacks and vulnerabilities and it is important to identify which ones may relate to the enterprise and which are a reliable source of trustworthy information.

A certain amount of ongoing self-education is needed to maintain this level of competency and to gain knowledge and understanding. Training courses are available to develop this knowledge with courses that range from specific topics on information assurance to training that covers a wider focus such as security management. Masters degrees, some approved by the NCSC, are available from a number of universities and some of them provide part-time or distance learning options. There are also several professional accreditations that can be gained to develop your knowledge and demonstrate the depth of your experience to others.

It is sometimes necessary to seek external security services. Larger information assurance functions are generally able to support larger teams with a wider range of expertise. However, even in large organisations it may be more appropriate and cost-effective to buy in specific expertise that is expensive to maintain and only needed occasionally. Forensic analysis or security penetration testing are typical areas where it is often more appropriate for an organisation to seek specialist external services. When selecting an external specialist, it is important to understand the skills being offered by the supplier and to seek assurance through accreditation, recommendations and references. For example, within the UK, the Council of Registered Ethical Security Testers (CREST) provide assessments of companies providing security testing services and individual testers. The NCSC has established a scheme, the Certified Cyber Security Consultancy, which certifies competent independent consultancy companies who offer information or cyber security advice and guidance.

Creating a culture of good information security practice

Information assurance needs the co-operation and collaboration of everyone with access to the enterprise's information. Involving everyone in the assurance process will help to develop a culture of good information security practice.

As previously mentioned, it is important that information assurance is taken seriously by senior management within the enterprise and that they provide sponsorship and support for assurance initiatives. If so, then their support and commitment will cascade down through the organisation. Line managers will proactively take responsibility for adopting information assurance measures within their teams, and likewise end users will know that they must take their responsibilities seriously. Positive reinforcement of good assurance behaviour by the information assurance function and senior management helps to cement good behaviour, and some organisations even include feedback on assurance behaviour in their performance reviews.

A key factor for success is ensuring that everyone that accesses the enterprise's information knows what is expected of them. Having in place clearly defined assurance roles and responsibilities, and up-to-date security policies and standards and procedures will eliminate any ambiguities. They do need to be clearly communicated and be readily accessible. For example, assurance responsibilities should be included in employee job descriptions and form part of third-parties' contractual conditions. All users need to understand clearly what will happen if they do not follow information assurance

policies and that senior management will be involved should the rules be breached. Policy development is covered in the next section.

Education is also a vital component in creating a culture of good information assurance practice. If everyone understands the value of the enterprise's information assets and how they can be put at risk, then they are far more likely to appreciate why these processes and procedures are in place. Regular awareness campaigns initiated through the information security manager can help to reinforce this message.

ACTIVITY 3.1

Ms Jackson has just returned from a conference and has realised that the enterprise does not have a formal information assurance function. She has asked you to put together a high-level proposal on what should be put in place for GANT.

ORGANISATIONAL POLICY, STANDARDS AND PROCEDURES

Organisations require their staff and third parties to use, manipulate and interpret information and need their co-operation to ensure that the information assets are accessed in a secure and responsible manner. All users need to know what the enterprise expects of them with regard to this. Policies, standards, procedures and guidelines provide this guidance.

LEARNING OUTCOMES

The intention of this section is to provide you with the basic knowledge needed to develop, write and gain user commitment for assurance policies, standards, operating procedures and guidelines. Following study in this area, you should be able to explain and justify not only the main concepts but also to draft documents to meet the general requirements in the following areas.

Developing, writing and getting commitment to security policies

One senior person within the organisation should be given the overall responsibility for protecting the assurance of the organisation's information assets and be formally held accountable to ensure that appropriate security controls are implemented across the business. This director should be supported by a working group to ensure that adequate assurance measures have been put in place to protect the organisation to an acceptable level of risk. Involving senior management will help to endorse the governance process, ensure that adequate resources are made available, ensure that controls are implemented effectively and that any identified security gaps are addressed.

Developing standards, guidelines, operating procedures and so on

There is often confusion concerning the definition of policies, standards, procedures and guidelines, so this should be clarified first of all.

A policy is a high-level statement of an organisation's values, goals and objectives in a specific area, and the general approach to achieving them. Although they should be regularly reviewed, policies should hold good for some time as they are not intended to provide either detailed or specific guidance on how to achieve these goals. For example, a policy might say that each user is responsible for creating and maintaining their system passwords – although it doesn't say exactly how to do this. Policies are mandatory.

A standard is more prescriptive than a policy. It quantifies what needs to be done and provides consistency in controls that can be measured. For instance, passwords must contain a minimum of eight characters, be a mix of numbers, letters and special characters and be changed if compromised or for other similar reasons. Compliance with standards is also mandatory. They should support policy and state what 'must' be done and how it should be achieved. Standards can be either general (e.g. handling sensitive information) or technical (e.g. encryption of data), but they should always relate to a specific subject.

A procedure is a set of detailed working instructions and will describe what, when, how and by whom something should be done. Again, they are obligatory and should support enterprise policies and standards.

Guidelines are not mandatory, but can provide advice, direction and best practice in instances where it is often difficult to regulate how something should be done (e.g. working practices when out of the office).

Whether producing policies, standards, procedures or guidelines, these documents should always be clearly written and to the point. Language should be concise, unambiguous and as free as possible of complex jargon and acronyms. Statements should contain positive rather than negative 'do not' rules, as these tend to make users less responsive. A document should address a clear and well-defined subject area within its scope so that the target audience knows that it is relevant to them (e.g. 'this policy applies to all GANT employees in the UK').

Policies, and any attendant standards, procedures and guidelines, should be endorsed by senior management and have clear ownership (i.e. head of Human Resources, departmental manager, etc.). In addition to senior management, they should also be supported by the main stakeholders and especially those people tasked to enforce them. For example, if an individual is to be disciplined for a breach of policy, then Human Resources will need to support the policy to carry out any disciplinary action in relation to it. To have credibility, policies should be endorsed by all interested parties such as stakeholders and user communities.

Policies, standards and procedures need to be realistic and enforceable. It may be a great aspiration to expect people to keep their laptops with them at all times, but it is far more realistic to state that they must not be left unattended in a public place. Similarly, if a technical control is stipulated (e.g. encryption of all data on laptops) it

must be anticipated whether this is actually feasible. There may be instances where it is not possible to comply with policy and consideration also needs to be given to allowing exceptions to policy in the way of special dispensations.

To be enforceable, policies must be consistent with other corporate policies and compliant with the law. All users need to know what will happen if they do not comply. Policies need to be regularly reviewed to ensure that they remain current, relevant and effective. This will be covered in more detail later in this chapter under 'Information security governance'.

The way in which policies, standards, procedures and guidelines are structured will largely depend on the organisation. Regardless, every organisation should have a (high-level) assurance policy that states the organisation's commitment to information assurance and what it expects to be done to protect its information assets. A security policy is a strategic statement of the organisation's approach to assurance and sets out the formal organisational stance on assurance matters for everyone to see. This security policy should contain statements on:

- how the enterprise will manage information assurance;
- the protection of information assets in accordance with their criticality;
- the compliance with legal and regulatory obligations;
- the means by which users will be made aware of information assurance issues and the process to deal with breaches to policy and suspected assurance weaknesses;
- the fact that this policy has the support of the board and chief executive.

More detailed guidance on what to include in a security policy can be found in recognised standards such as the ISO/IEC 27000 series and the ISF *Standard of Good Practice*. The high-level security policy should be signed off by the director responsible for information assurance. The policy should then be issued or made available to all individuals with access to the organisation's information and systems, both internal and external, in a format that is readily understandable and accessible by the user.

Third parties often require access to an enterprise's information assets in terms of processing information, offering support, providing services or processing facilities. It is important to ensure that there is no misunderstanding between the enterprise and the third party over what controls are to be put in place to protect the enterprise's information assets. Policies, standards and procedures should be extended to third parties where relevant, and specific policies may need to be written to cover third-party arrangements. These should be included within the terms of a contract. Access should not be given to an external entity until the enterprise can be assured that the appropriate controls have been put in place and that the third party has formally confirmed that they understand their obligations and accept their responsibility to comply.

As the relationships with third parties can be quite diverse and extensive (as in the case of outsourcing), any terms associated with policies, standards and procedures may vary according to the type and nature of the relationship. Agreements with third parties should include the enterprise's assurance policy. Again, the ISO/IEC 27000 series and

the ISF *Standard of Good Practice* contain guidance on the type of controls that should be considered for inclusion in third-party agreements, but typically they should include the following arrangements:

- management of changes to the application/facility/service/resource;
- the right to audit and monitor assurance arrangements within the third party;
- notification and investigation of assurance incidents and security breaches;
- the timely sharing of relevant cyber security information and knowledge;
- recruitment of personnel.

Care should be taken to ensure that sensitive information is not disclosed to or by third parties, and policies should reflect demands on the third party for confidentiality and non-disclosure of information. The third party may, in the process of delivering the service, use further subcontractors or service providers. It is important, therefore, to ensure that any policies, standards, procedures and guidelines are applied to them too and this can be controlled in the contract.

Balance between physical, procedural and technical controls

Physical, procedural and technical controls (often termed operational types of control) can provide very effective security mechanisms and do much to reduce the likelihood of incidents occurring. However, they each have their limitations and there are occasions when their use is not appropriate – possibly their deployment would be far too complex or expensive given the perceived value of the information and the associated risk. For example, a £1 million security application to protect a £10,000 information asset does not make much financial sense. In other cases, a physical or technical control may be so intrusive that the users are prevented from carrying out their work efficiently. In many cases there may be no reasonable physical or technical controls that can be deployed to prevent a particular security breach from occurring or it may be that the security controls in place can be circumvented by the user in some way.

Users need to access information systems in order to carry out their tasks and this inevitably introduces a level of risk to the information. They may need to share this data with colleagues or external suppliers and make value judgements as to whether it should be released to them. Reducing this kind of risk is difficult to achieve through technical controls alone. Technical controls introduced by a documental security system, for example, may well provide a good level of security. There will, however, always be exceptions, and these need to be handled in a consistent manner by having a policy and process in place. This might simply involve informing a senior colleague of the issue and the proposed course of action to deal with the issue in the short term.

Formal policies and procedures can be used to make users aware of their responsibilities and the risks relating to the data to which they have access. The policies can empower individuals to make decisions as to whether others should access this data. This can be an effective control measure, but is obviously dependent upon users complying with these policies and associated standards and procedures.

Occasionally, due to time pressures, or perhaps because of expediency, policy rules may be circumvented or ignored. Ignorance or a failure to properly understand the policy will prevent compliance, and in these instances users won't understand the risks to their information assets and are very unlikely to be fully aware of the threats to them. Policies and procedures rely on individuals knowing that the policy exists and understanding what the policy expects of them as well as their agreement to comply with it. So, policy controls have limitations.

There needs to be a sensible balance between using physical, procedural and technical controls to manage the risks associated with information assets. All three elements should be used to complement one another in a layered approach to manage risk to an acceptable level.

Defence in depth and breadth

These principles are a critical instrument that must be fully appreciated and, where appropriate, implemented if security is to be established and maintained effectively.

The principle of defence in depth is that there are layers of security that build on one another. Today it is virtually impossible for any organisation, large or small, to protect all their information assets to the highest degree possible. One reason is that of cost, but there is also a more practical reason: whereas some information assets need to be freely available to staff members in order for them to undertake their routine daily work activities, other information is far more critical and sensitive and so must be afforded a high level of protection, thereby almost certainly making it more difficult to access and use.

Defence in breadth is a more recently coined phrase that has come about due to the need to consider all the connections to any networked system. The complexity of the networks now in place in many organisations is staggering and, when connections to suppliers, customers, different physical locations around the world, homeworkers and many more are considered, they become ever more difficult to manage. The concept of understanding the breadth of the network and its connectivity is critical since, as the old adage says, the weakest link of a chain is where it will break. The part of the very complex network that has the lowest level of security will be where the criminal or other intruder will find their way in. Malware is now very sophisticated software and it has the capacity to find ways of traversing a network to find holes or legitimate portals into other areas that, because it is already inside the network, might, for example, allow the intruder to be seen as a trusted user.

End-user code of practice

The development of a high-level security policy should be bolstered by an end-user code of practice or acceptable use policy that provides a readily accessible way of communicating requirements to end users. An acceptable use policy demonstrates the organisation's commitment to information assurance and must be approved by the director responsible for information assurance. It should be published to all users that need to access the organisation's information management systems and include all employees (permanent and temporary, full- and part-time), contractors and third

parties. The acceptable use policy should detail what is expected from users to protect the organisation's information assets. Elements that may be included in this policy are:

- ensuring that user passwords and PINs are protected appropriately, are not compromised and are changed at appropriate intervals;
- ensuring that users only access information, facilities or equipment for which they have the designated business need and requisite authorisation;
- logging-off from systems when leaving a workstation unattended;
- locking away sensitive documentation and media when not in use (as part of a clear desk policy, for example);
- use of personal devices such as smartphones and tablets;
- ensuring that all security incidents are reported.

An acceptable use policy can also include general statements regarding behaviour in the workplace, such as making it unacceptable to make any sexual, racist, obscene, discriminatory, harassing or other offensive statements regardless of the method used to transmit such statements (email, instant messaging, telephone, text, paper or spoken word). All conditions of employment for permanent or contract employees should contain a statement that compliance with the enterprise information assurance policies is mandatory. To avoid vicarious liability, the policy should also include statements that specify that users must comply with all appropriate legal and regulatory requirements placed on the organisation.

Consequences of policy violation

Anyone accessing the organisation's information assets needs to know what the consequences of a policy violation are, and this should be clearly stated in the policy, standard or procedure. Appropriate processes should be established for reporting and dealing with violations so that they are dealt with in a consistent manner. These processes should be documented and agreed with the relevant stakeholders when the documents are produced.

Violation of a policy may, in severe cases, lead to an employee disciplinary process being instigated, termination of supplier contract or the need to report the behaviour to the appropriate law enforcement agency. Therefore, the rules and processes need to be understood, agreed and put into effect within the organisation before violations may need to be dealt with. Naturally it is essential to involve the HR and legal departments in the development of such policies to ensure the proposed course of action complies fully with all employment legislation as well as other relevant national laws.

However, it is a waste of time having a policy in place unless the organisation is prepared to enforce it. Senior management, and those that have to enforce the rules, need to support the processes to deal with any violations. If violations have not been dealt with appropriately, or have been ignored by line management, then this should also be considered as a violation of policy and treated seriously.

ACTIVITY 3.2

Ms Jackson has asked you to prepare an end-user code of practice for GANT. Identify the main areas that you would include in the policy.

INFORMATION SECURITY GOVERNANCE

There is an increasing amount of legislation and regulation that requires senior management to ensure that adequate controls are in place to protect the enterprise's information assets. To fulfil these obligations, senior management needs to understand the current status of existing assurance controls, where such controls are inadequate and how the risk profile of the organisation is changing. The necessary effort can then be made to improve security mechanisms and manage the risk effectively. This section covers the principles of the governance processes that should be implemented to enable this to happen and to provide senior management with sound and up-to-date information on the state of assurance within the enterprise.

LEARNING OUTCOMES

The intention of this section is to provide you with the basic knowledge needed to understand the principles of information assurance governance. Once completed, you should be able to explain and justify the main concepts and establish procedures and draft documents to meet the general requirements in the following areas.

Review, evaluation and revision of security policy

The production of policies, standards, procedures and guidelines has already been covered in the previous section, but to ensure that they remain current, relevant and effective they should be reviewed regularly. Reviews should take place after any significant changes to either systems or resources or as part of a regular review schedule (e.g. annually).

A management review process should be established to ensure that policy reviews take place in an organised and timely manner. The review schedule should identify all the persons to be involved and a formal record kept of any revisions made – with an explanation as to why content has been incorporated, altered or removed. Senior management should then approve the final version of any amended documentation.

The review should involve all the main stakeholders, including external parties and, where applicable, regulatory authorities. The review should focus on factors that might influence or trigger possible amendments, such as:

- changes to technology, processes, organisation, resource availability or working practices;

- changes to contractual, regulatory or legal requirements;
- changes in threats and vulnerabilities;
- results, actions and recommendations from any assurance reviews or audits;
- findings and recommendations from either incidents or previous assurance breaches, or where there is evidence of non-compliance with the policy.

Once the review has been completed the revised policy should be communicated effectively to the relevant users, both internal and external to the organisation. This process should also be used for the maintenance of all other assurance documents, such as security standards, procedures and guidelines.

Security audits and reviews

Audits and reviews provide a good opportunity to understand how well things are working within the enterprise and provide senior management with valuable information on the assurance of their environment. Regular independent assurance audits and reviews should be carried out across the business to ensure that its information systems are compliant with existing security policies, standards and controls. Possible vulnerabilities to these systems can be checked and the effectiveness of existing controls can be tested. Audits and reviews should be carried out periodically or when a significant change (e.g. a system upgrade, new threat or vulnerability, a change of risk appetite, etc.) has occurred.

To introduce a measure of impartiality into the review, it should be carried out by an independent party, which will also bring to it a fresh set of eyes. Ideally, a member of an audit team or a manager that has no conflict of interest in its outcome could do this. Alternatively, reviews can be carried out by a third party such as an external auditor or a consulting company. In the case of technical reviews, it is often beneficial to engage a company with specialist knowledge in areas such as penetration testing. They can bring with them their experience of having audited similar assurance-based scenarios in other organisations. Reviewers must have sufficient expertise, so it is prudent to verify their abilities before commencement. Technical testing should only be carried out by recognised and approved technicians and engineers. Any organisation providing testers should also be able to verify that the individuals' CVs and background have been checked to ensure that they have a suitable level of personal integrity. Information on technical testing and assessment is covered in more detail in Chapter 6.

A programme of information assurance audits and reviews should be introduced by senior management. The scope of each individual review and their deliverables should be agreed by senior management and the area owner (i.e. head of department, line manager, etc.). A scoping exercise should be completed before the audit or review is started and a checklist should be developed to measure the efficacy of the assurance controls. The outputs should show whether the defined controls (e.g. policy, standard, procedural or technical) have been implemented correctly and are effective enough to reduce risk to an acceptable level.

Access rights to systems or information assets for those performing the audit or review should be restricted to only what is necessary on a need-to-know basis. These should be monitored and logged to create a reference trail of their activities. Wherever possible, auditors and reviewers should be given read-only access to isolated copies

of the system. Audit tools should be restricted to prevent any possible misuse or compromise of data. The legal implications of providing third parties with access to sensitive information should be considered, as should the disposal of any information, reports and scripts that result from the audits or tests. NDAs should be put in place if the information being reviewed is sensitive, as will often be the case. Audits and reviews should be planned well in advance to minimise the risk of disruption to the normal operation of the enterprise. Some audits, such as penetration testing, may produce some unexpected activity on a computer system or network. Change management processes should be followed to ensure that all parties that could be affected are aware of the planned activities and potential change in activity levels.

The results of the audit or review should be recorded within a formal report and presented by the reviewer to senior management and the manager whose area has been reviewed. A plan of corrective action should be agreed with them, including timescales for implementation. The plan should be monitored regularly to ensure that actions are being progressed. Any risks identified during an audit or review should always be added to an information risk register, maintained centrally by the organisation. All documentation produced should be filed securely so that it can be referred to when planning the following cycle of audits and reviews.

Checks for compliance with security policy

Regular checks should be carried out to measure compliance with security policies, standards and procedures. Carrying out compliance checks helps to identify whether controls are still adequate and relevant. Compliance checks also help to gauge the level of user understanding and awareness of their assurance responsibilities and whether or not these are being taken seriously. If regular checks are not carried out, then over time there can be a tendency for users to show less regard for them. Assurance is weakened as users become aware that monitoring does not take place and that they are not likely to be challenged.

If an instance of non-compliance has been identified, then it is necessary to discover why this has happened. This could be due to lack of training, misunderstandings or perhaps simple disregard of procedures. It may have resulted from a change in business processes that has not been recognised in the assurance documentation. The compliance checker must then decide what action should be taken and whether measures need to be put in place to prevent further occurrences. Action taken should reflect the scale of the non-conformity; minor incidents, such as an isolated instance where a procedure has not been followed, could be dealt with in an informal manner, but any major non-conformance, such as widespread password sharing, should be addressed more formally. If corrective action is recommended, subsequent reviews should identify that it has been implemented.

The results of compliance checks should be recorded and serious instances of non-compliance reported to senior management. The findings from the checks can be fed into subsequent policy reviews. Compliance reviews should also be carried out to ensure that the enterprise is using licences, for example for software, in accordance with terms stated in the purchase agreement.

Reporting on compliance status

The finance industry has a long history of regulation and most stock exchanges have their own regulatory controls to prevent financial malpractice, but governance controls have gradually extended to other operating spheres. Many countries have produced their own codes of ethics, often in response to large corporate failures or in response to public pressure. The Sarbanes–Oxley Act was introduced in 2002 following a number of high-profile financial accounting scandals in the USA. The EU's governance legislation was revised in 2004 via the Companies (Audit, Investigations and Community Enterprise) Act, which has been implemented across the member states and has replaced most of their local company legislation.

The type of strategy deployed by the enterprise to meet their information assurance compliance obligations will depend on the risk appetite of the organisation and the external requirements placed on them. The enterprise needs to understand what their specific obligations are so they can implement the necessary controls and reporting mechanisms. In some situations, tightly specified governance requirements do apply, but mostly they ask the enterprise to demonstrate that good information assurance controls have been implemented. Generally, regulators will want assurance that senior management are committed to protecting the enterprise's information assets, understand the enterprise's risk profile and have implemented controls to manage risk to an acceptable level. Regulators will also want assurance that the controls in place are working effectively and that any gaps identified are being addressed.

Senior management and any regulatory or compliance bodies need to have access to sufficient information to be able to demonstrate compliance. To do this, the following types of information need to be made available:

- high-level risk assessments for the enterprise and for critical systems and services;
- a risk register showing how identified risks will be, and are being, managed;
- an up-to-date set of security policies with a review process;
- a register of any dispensations from security polices;
- the results from assurance reviews and security testing and compliance reviews;
- reports from any assurance breaches or incidents, including any actions taken;
- plans to address any compliance weaknesses.

The relevant information has to be gathered, reviewed and presented in a format that is acceptable to the regulator. This activity can be very time-consuming, so it helps to develop a repeatable and efficient process for reporting on compliance issues and to

reuse controls for each of the regulatory groups. This process may form part of an enterprise information assurance policy.

There are various models (comprising a methodology, structure and processes) that can be adopted by an enterprise to provide this level of information. All the models tend to be based upon the principles of implementing a formal control process for:

- understanding risk;
- identifying control requirements to reduce risk to an acceptable level;
- implementing effective security controls;
- monitoring how effective the controls are;
- periodic re-evaluation of risk levels;
- the efficacy of controls to enable continual improvement and to ensure that the risk level is maintained.

Current models include ISO/IEC 27001, Security Operations Maturity Architecture (SOMA) and the Committee of Sponsoring Organizations of the Treadway Commission (COSO), and they all offer levels of accreditation to enable the enterprise to demonstrate their competency to other organisations and regulatory bodies. The ISO/IEC 27001 model provides an approach for establishing, implementing, operating, monitoring, reviewing, maintaining and improving assurance within an organisation. This is known as the 'Plan–Do–Check–Act' (PDCA) or Deming Cycle approach to provide a cycle of continuous assurance review and improvement, although it is no longer part of the ISO/IEC 27001 standard.

SOMA, produced by the Institute for Security and Open Methodologies, provides a framework for measuring the operational security and management process and is structured in maturity levels that can be adapted to work at different levels of assurance maturity within the enterprise as well as being used with other standards. COSO, produced by the Treadway Commission, provides a framework for evaluating effectiveness of assurance by establishing a set of objectives for assurance control and measuring against them. This is often used for testing the effectiveness of accounting controls. There are many more such standards and frameworks, some fairly generic, such as those from the USA in the series of standards issued by the National Institute of Standards and Technology (NIST), and others targeted at specific industries or activities, such as the Payment Card Industry Data Security Standard (PCI DSS) requirements that govern the acceptance of payment cards.

ACTIVITY 3.3

After the recent loss of information, Ms Jackson is concerned that she needs to demonstrate to the regulators and external auditors that good assurance controls are in place within GANT. How would you provide her with evidence to demonstrate that assurance is being managed effectively?

INFORMATION ASSURANCE PROGRAMME IMPLEMENTATION

An information assurance programme provides a high-level view of how the organisation will address its assurance needs. It can help to develop a common understanding of information risk and enable the organisation to prioritise and focus on implementing controls that address the risks that matter most. This section looks at various aspects of assurance planning and how to implement an assurance programme within an enterprise. It is not intended to offer guidance on either project or programme management. There are many publications that address this subject, including the BCS publication *Project Management for IT-related Projects* (Hughes et al. 2019), which provides the textbook material for the BCS Foundation qualification in IS Project Management.

LEARNING OUTCOMES

The intention of this section is to provide you with the basic knowledge needed to understand the principles of how to implement information assurance measures within an enterprise. Once completed, you should be able to define and explain the main concepts and also draft documents to meet the general requirements in the following areas.

Planning – ensuring effective programme implementation

Good planning is the foundation of any successful information assurance programme implementation. It can be used as a powerful tool for gaining support from both senior management and key stakeholders and to demonstrate how the assurance programme is helping to reduce risk within the enterprise. This builds support for further initiatives.

To have credibility, an information assurance implementation programme has to be realistic, be achievable and to address accurately the needs of the enterprise. A programme will need to fulfil agreed objectives within the timescales set, and to demonstrate quality, value for money and overall benefit to the organisation. A programme usually consists of a number of projects, with each addressing either a function, application or information asset.

When planning a programme implementation, it is necessary to understand the current status of assurance within the enterprise and what the programme needs to achieve. Ideally, a prior risk assessment should have been carried out. The outputs will help to define and shape the implementation programme and provide justification as to why it should be delivered. Guidance on carrying out a risk assessment is covered in Chapter 2.

Work should be prioritised to deal with the most pressing issues. A mix of tactical and strategic approaches may be used to address the issues involved. If possible, try to identify where there are some quick wins to be achieved. Those that have a high probability of success using only minimal resources will do much to raise the overall credibility of the implementation plan within the organisation.

When planning an implementation, consideration should be given to how long it will take to implement controls, how easy the implementation will be, what the associated costs are and a measure of the appetite of the organisation for wanting to resolve these issues. The main steps to developing an implementation programme and plan are to identify:

- how the implementation programme will address risks within the enterprise and reduce them to an acceptable level;
- the potential benefits of undertaking the programme;
- the controls or work streams that need to be set up to achieve this;
- the level of effort that will be required and from whom;
- who will be accountable for each part of the programme;
- the costs and timescales associated with implementation;
- how progress will be tracked.

High-level support is essential to the success of the programme and it should always have a senior responsible owner (sponsor). The programme and approach should be agreed with the main stakeholders and signed off by the sponsor. A steering committee should be set up to track the success of the programme and deal with any issues that arise. Resource and budget will need to be secured before the programme starts.

Depending on the scale and size of the programme, it may be necessary to split it into a series of separate projects each with their own planning and tracking mechanisms. During the course of the programme there may be a need for planning revisions due to changes in enterprise priorities, budget cuts, unavailability of key human resources or similar. Such factors might affect what can be achieved; therefore, the overall plan should be kept at a high level with its key deliverables and main milestones clearly expressed, allowing the progress of the programme to be monitored readily. Detailed planning should be carried out near to the time of implementation where there is greater certainty of a settled environment and an appropriate understanding of what needs to be done at that point.

The plan should focus on the key deliverables and work should be divided into manageable amounts that can be measured (milestones). Unless human resource is fully dedicated to the project, it may be necessary to calculate how much of a person's time has been committed to it. Project work can be allocated accordingly and enable resources to meet both project activities and other commitments.

The project plan should be regularly reviewed. The frequency of the reviews will vary depending on the type of implementation (daily, weekly, monthly). The purpose of these reviews is to maintain a vision of its actual progress by understanding what has been achieved thus far, comparing progress against the schedule, handling variations and revising the plan, identifying problems and applying corrective action. Progress should be reported regularly to both the programme sponsor and any stakeholders, and review meetings should be set up to enable them to agree and support programme alterations to meet changing requirements.

How to present information assurance programmes as a positive benefit

An information assurance programme should be seen as delivering positive benefits to the enterprise. Managing down information risk can bring about tangible benefits in terms of greater stability of information systems and improved protection to sensitive information. These benefits can show the rest of the organisation that assurance priorities are aligned with the priorities of the enterprise.

Communicating with senior management, line managers and general users in a manner that relates to their own particular interests can do much to change the view that the assurance function is an inhibitor to a view that it is an enabler by demonstrating how it can add value to the organisation. Similarly, establishing good interpersonal relationships with stakeholders and colleagues in general will always help to present assurance programmes as a worthwhile activity. It is essential to be sensitive to their needs and ensure that they understand what the programme means specifically to them and how it will affect their role. Initially, colleagues may be wary of an implementation programme. Including them within the planning process will help to gain their confidence and appreciation. By aligning assurance objectives with the overall enterprise culture and values it can be seen that you are working collaboratively to achieve the same shared goals.

As with all initiatives, it is essential to have the support and commitment of senior management. Usually, they are not assurance specialists and do not have a full understanding of the information assurance risks and issues facing the organisation. It is important to present the positive benefits of the assurance programme in a manner that is concise and free of jargon. To support the programme, senior management needs to understand:

- the risks facing the organisation;
- the cause and potential impacts of these risks;
- the benefits they will see from their investment;
- where there may need to be changes to ways of working;
- how they can support or sponsor the programme.

The programme should be formally presented to them by way of a sound business case and be accompanied by all the necessary facts to enable an informed decision to be made. Senior management is generally more favourable to persuasion if they can realise a return on their investment. By quantifying impacts and explaining where the enterprise may be vulnerable, it is possible to demonstrate that a reduction in impacts can lead to a reduction in operating costs. Return on investment (ROI) is a mechanism that can be used to justify assurance expenditure and gain budget approval.

By calculating a positive financial return on an assurance investment, it is possible to put forward a persuasive case for its implementation by balancing potential financial rewards against the costs of implementing the controls. This is achieved by calculating how much it would cost to purchase and implement a particular security control and then estimating, in cost terms, what expense the organisation could incur as a

result of assurance incidents. Demonstrating that good assurance controls can be used by the enterprise for competitive advantage can also help to present assurance programmes as a positive activity. For example, many organisations insist that an enterprise has certification to particular security standards, such as ISO/IEC 27001, as a prerequisite before they engage with them. This provides them with an assurance that the enterprise has implemented effective assurance controls. Enterprises that have had assurance breaches can often suffer financially through fines and lost business as their trading partners and customers lose confidence in their abilities to protect their information. For some organisations a security breach can have an enormous negative impact on their brand value. A low instance of security breaches and good assurance controls in place can provide commercial advantage over less conscientious competitors.

Security architecture and strategy

Information security strategy and architecture are two concepts in information assurance implementation that have gained credibility and value in recent years. This section will look at some of the high-level principles regarding these concepts.

An information security strategy is a plan to take the assurance function within an organisation from the reality of where it is now with all its problems and issues, to an improved state in the future. It provides a road map or vision as to how this can be achieved and how it will support the organisation going forward. A strategy should normally cover a period of time where it is possible to implement a significant level of change but short enough to be able to predict changes in technology and organisational objectives. Typically, this is over a three- to five-year period.

An information security strategy has the elements of an implementation programme, but covers a longer period of time and is pitched at a much higher, less detailed level. It should demonstrate how it will enable the enterprise to achieve its objectives and how it will protect it against current and future threats. It should consider:

- the current state of assurance and the strengths and weaknesses of existing controls;
- how the risk profile of the enterprise is likely to change in response to changing business objectives and working practices;
- trends in threats and vulnerabilities to potential types of incidents;
- expected developments in software and hardware;
- legal, compliance and audit requirements and any anticipated changes;
- areas where cost savings can be made.

As it is a vision, this high-level document should be written in concise non-technical language so that the target business audience can clearly comprehend the bigger picture and the vision being presented. The strategy should remain a living document by being regularly reviewed and updated to reflect changes in technology and organisational priorities as they are likely to change over the period of its existence.

Having a strategy in place shows a degree of maturity of the information assurance function within the organisation. It provides an assurance that the organisation is committed to good information assurance governance. Increasingly, there are pressures exerted from external bodies (via legislation and regulation) for organisations to have a security strategy in place.

The second concept is of information security architecture that can be used in conjunction with the information security strategy. The architecture translates organisational requirements for assurance into a set of controls that can be used to protect the enterprise's information assets. The information security architecture should aim to provide a common and consistent framework of global assurance controls and arrangements to be used across the enterprise rather than in a piecemeal fashion. Traditionally, IT-based architectures are focused purely on technology, but, as information assurance behaviour extends beyond technology to include policies, processes, procedures and user behaviour, so it follows that an information security architecture should also encompass these aspects.

An enterprise information security architecture should provide system developers and administrators with a consistent framework of assurance controls that can be used across multiple systems and environments within the enterprise. It can be adapted to fulfil different circumstances, which means that effort is not duplicated every time a new set of controls is implemented. Efficiency and productivity can be increased, while costs can be reduced by providing leverage and economies of scale. This should enable the assurance function to react more quickly to commercial, organisational and technological changes and be able to implement assurance solutions more quickly, efficiently and at lower costs.

It also supports the concept of defence in depth, where layers of security can be implemented such that only the most valuable or sensitive information is afforded the highest protection, a pattern for security sometimes called the onion model. The idea of defence in breadth can also be covered by information security architecture if the bounds of the organisation are extended, in terms of the information security aspects at least, to include key partners and suppliers. This would mean an organisation setting out their terms and conditions for the security of connections of systems and other logical interactions. These are sometimes referred to as code of connection or CoCo.

An information security architecture works on a set of 'principles' that express the type of controls to be implemented. They act as positioning statements that will be adopted within the architecture. An example of this is 'auditing and monitoring controls will ensure that the organisation complies with security policies and legal obligations'. Having identified a set of principles, the architecture can then be modelled through increasing layers of complexity and detail, starting at a high-level conception view of controls though to a detailed specification and design.

Components within the enterprise with similar security requirements can be grouped together into 'domains' so that common sets of security controls can be developed to protect them. For example, all enterprise systems with web-enabled interfaces can use the same domain controls. The term 'services' is used to describe the type of controls that will be used to protect these components.

The need to link with business planning and risk management and audit processes

The aim of an information assurance programme should be to reduce information risk within the enterprise. As we have seen throughout this section, information assurance planning and implementation processes should not work in isolation. To be effective, an implementation programme needs to understand the enterprise's business objectives and goals so that it can identify the appropriate assurance control measures to ensure that the enterprise is sufficiently protected to meet these goals.

Information assurance implementation programmes need to work closely with other organisational and assurance processes to manage risk to an acceptable level. The risk management process should provide awareness and understanding of the risks faced by the enterprise and identify where risks are not being managed effectively. The outputs from risk assessments should determine what controls must be implemented and assess how urgently they need to be addressed.

Similarly, assurance governance processes will identify where existing controls are inadequate and where improvements need to be made. This may be through the reporting of assurance breaches or via auditing or testing of security controls. Governance processes will also determine changes in regulatory or legal requirements that may require additional controls to be put in place.

All implementation programmes should support the enterprise's information assurance policies, security strategy and security architecture. All security controls should support its long-term vision for information assurance and should be seen to add value or business benefit to the organisation.

ACTIVITY 3.4

Ms Jackson is very pleased with the work that you have done so far on information assurance and has given you a budget to implement some additional access controls within GANT. How would you approach this?

SECURITY INCIDENT MANAGEMENT

No matter how good an organisation's risk assessment and the controls that have been implemented, and regardless of how careful the organisation is in conducting its day-to-day business, security incidents will happen. It's not *if*, it's a case of *when* and *how often*. People make mistakes, systems malfunction and there are zero-day attacks that will defeat the best security tools. Organisations that have good information security maturity recognise this and then plan and train for it.

Security incidents don't just affect the confidentiality of data. The impact can equally relate to the integrity or availability of data or any other asset employed by the organisation or provided by a third party (such as cloud services). It is important to have

plans in place to deal with the most likely eventualities before they happen. Trying to think of and implement solutions after the event is much more difficult, will take longer and is more risky in terms of them not working, not to mention bringing additional costs. Planning in advance also means a very hard question from the senior management, 'Why didn't we have a plan to deal with this?', can often be avoided.

LEARNING OUTCOMES

The intention of this section is to provide you with the basic knowledge needed to manage security incidents and to plan and conduct a forensic investigation. Once completed, you should have formed an understanding of the following concepts.

Security incident reporting, recording and management

Having an incident response plan that has been worked out in advance and tested is like having a good insurance policy. Trying to obtain one after the event is too late. Do not think that it is possible to improvise something quickly, because this is a complex subject with a lot of variables. In the ensuing chaos, distractions and pressure to recover to normal operations, the organisation will not have the luxury of time to come up with a good plan, train people and test it thoroughly. Evidence shows that more than half of incident response plans fail when first tested.

A security incident response plan is a set of instructions to help the organisation, and the incident response team in particular, to detect, respond to and recover from information security incidents. These types of plans address issues like:

- DoS attacks;
- discovery of cybercrime, for example loss of data or payment card information;
- malware outbreak;
- incident requiring notification of a regulator, for example ICO for loss of personal data;
- service outages;
- security incident at a third-party service provider upon whom the organisation depends.

The first priority is to ensure that all the people within the organisation know how to recognise an incident and know to whom they should report it. This can be done in a number of ways including awareness training, dedication of a section on the company intranet/portal and by carrying out exercises.

There are normally five phases in the management of an incident:

1. reporting;
2. investigation;
3. assessment;

4. corrective action;
5. review.

In order to ensure that enough information is captured, it is good practice to have a standard form to hand on which to record the information given by the person making the initial report. The form should be easy to find on the intranet and should prompt the user for the following information:

- who they are;
- where they are – geographic location and department;
- contact details – address, telephone (desk and mobile), email, IM;
- brief description of incident;
- whether there is any danger to life, health or company assets;
- other potential impact to business operations;
- description of anything they have done to date in response;
- time first noticed.

Any favourite search engine will help to find plenty of examples of reporting forms and sample incident response plans from which inspiration can be gained to create ones suitable for an organisation.

From the moment of the first report until the incident is closed, a log should be kept of information, decisions made and the consequences of any actions. These records will be invaluable later – both for internal use and also for possible use by external agencies including those for law enforcement. In the event that there are legal or regulatory consequences to the incident, these records may be required to be presented as evidence in a court of law.

Incident response teams/procedures

An incident response team (IRT) must be appointed up front and all members of that team should be properly trained, briefed and prepared in how to use the plan. The members need to come from a cross-section of the organisation to ensure that there is sufficient breadth of knowledge to deal effectively with the situation. They need to be senior and experienced enough to have the authority to make decisions on the spot. The IRT must also be empowered to call upon additional resources, internal and external, as they see fit to use in resolving the incident.

There needs to be a documented escalation process for the team to reach the most senior members of the organisation as and when necessary. The senior risk owner, who is normally a board member, will need to be briefed and ready to provide support and additional resources if required to help respond and recover. It is advisable to give each of the team a laptop with remote access to the organisation. It should have a full set of the incident response plan documentation and be suitably encrypted to protect the contents. Each member of the IRT should also have a company mobile phone so they can be contacted immediately and also reach out to each other when the IRT is activated.

Ideally, one or more persons on the IRT should be designated as note-takers. Their job is not to be involved directly, but to observe and record all details of the incident in a logbook for reference later. It is a good job for someone who is learning about information assurance in general, because they get to see first-hand what does (and does not) work.

The procedures must be quite broad in their content and scope because it is very difficult to predict the nature of the next incident. Some idea of likely events can be gained from looking at the risk register and the output of the BIA, in order to identify high probability and high impact events. There are some events, though, perhaps terrorist activity or environmental damage such as flooding, that are not part of the normal threat profile but can happen. An incident may seem unlikely but in the last 20 years one of the authors has lost three IT systems to terrorist action and was called up as an expert witness after a law enforcement organisation had one of their data centres destroyed by the Buncefield storage depot explosion in 2005.

Need for links to corporate incident management systems

Large organisations often have mature processes in place to support incident response teams. There is often a centralised function with access to resources and expertise to help deal with the incident. The latest procedures and contact lists may well be publicised on the intranet or information portal. In a geographically dispersed organisation it may be that other offices will have plans to send in specialist staff to help with managing and recovering from an incident, to provide cover for the affected location, or to prepare to receive staff relocated from affected premises. Even if an organisation is not one of these, it is worth approaching a large organisation located nearby. In return for the offer of help (perhaps an office with phones and internet access) in the event of a major incident, their IRT may well be prepared to offer reciprocal facilities and advice and guidance based on their experience.

It is worthwhile getting to know these people because the better prepared an organisation is, the better it will be able to handle the incident – whatever it may be. It is also well worth reading advice freely available online, such as the 'Respond' and 'Recover' sections of the NIST Cybersecurity Framework. Partnerships with other organisations can also provide valuable intelligence on attacks and incidents that might give an early warning, allowing an organisation to prepare better.

Processes for involving law enforcement

There are times when it will be necessary to involve law enforcement or other similar organisations in the response to an incident. In some countries there are incidents that require mandatory reporting to law enforcement, and it is important to know who they are in your local jurisdiction and that it may vary if there are offices in more than one country. For example, it is mandatory in the UK to inform the police if there is a suspicion of terrorist activity, a danger to life is suspected or that child pornography has been viewed or processed through the IT systems of an organisation. UK legislation also requires the reporting of suspicious financial activity.

If there is any likelihood of criminal activity or other deliberate action, the appropriate authorities should be notified. It is important that the IRT and senior management have

a good understanding of the legal requirements for reporting certain events and how to capture information to a standard that allows for forensic admissibility. This can be complex and is another reason why prior planning and preparation is essential. One single mistake in the procedure can render everything inadmissible in a court of law. Expert advice and guidance are highly recommended. A good starting point in the UK is to read *Good Practice Guide for Digital Evidence* (V5) (2012) published by the National Police Chiefs' Council (NPCC) (formerly the Association of Chief Police Officers (ACPO)).

Another risk is that of attempted extortion and blackmail where an attacker conducts something like a DDoS or ransomware attack on an organisation. In this case, the National Crime Agency (NCA) is the appropriate body in the UK to contact. Activity of this sort, or malware (malicious software) discovered in UK government departments, should lead to a report being passed to the NCSC. In other countries there will be similar organisations and law enforcement agencies who will be the key point of contact for such incidents. Often these are known as the computer emergency response team (CERT).

One last possibility is that an organisation may be visited by law enforcement officers conducting an enquiry into activities of which management has no knowledge. They may have a warrant to search the premises and remove items, or they may simply be conducting enquiries. The modern UK police force is aware of commercial sensitivities and the nature of intellectual property. While the organisation is best advised to take appropriate legal advice from internal or external sources, an open and co-operative response is almost always the best policy. One example of this is that the organisation must ensure it is abiding by the requirements of the local data protection regulations (GDPR in Europe) before providing any information to a third party, even if they are a law enforcement body. It is normally necessary for a legally-issued warrant to be presented to enable the organisation to provide data without fear of later repercussions from the ICO or other regulatory body.

ACTIVITY 3.5

The last external audit of GANT identified that there was no process in place to deal with any assurance breaches. Ms Jackson has asked you to produce a simple process for managing information security incidents. How would you approach this and who would you involve?

LEGAL FRAMEWORK

This section covers the general principles of the law in relation to information assurance management. This will cover a broad spectrum from the assurance implications on compliance with legal requirements affecting business (e.g. international electronic commerce) to laws that directly affect the way information can be monitored and copied. It will also refer to certain pieces of legislation to explain concepts and highlight the legislative variances between separate countries.

LEARNING OUTCOMES

The intention of this section is to provide you with basic knowledge of some of the general principles of law, legal jurisdiction and associated topics and how they affect information security management. Following study in this area, you should be able to explain and justify each of the following concepts.

Background

The UK's legal system has evolved over a long period of time. Within Britain, it can be traced back as far back as Roman and Anglo-Saxon times, and laws that were in place years ago can still exert an effect on the legislation of today. Legislation is continually evolving to adapt to the changing needs of modern-day life and reflects the cultural development and values of society. Understanding this can help to explain why there is variance in the legislation between different countries.

Countries that have some form of federal government will have multiple levels of law – such as in the USA, Australia, Canada or Switzerland, where there are local state laws that are ultimately subject to national or federal laws. Within the EU, there are European directives (agreements between the member states) that have been produced to harmonise pieces of legislation across member states. Each country has to incorporate the legislation into their own legal system, and this can result in subtle yet significant differences as each country interprets the directives in their own way.

The legal systems across the world do share many similarities and common legal concepts. However, as enterprises increasingly operate or perform transactions in more than one country, it is necessary to understand each particular legislation and how it can affect the enterprise's information as it crosses international boundaries. Organisations should always consult with a qualified lawyer to be sure what legislation is applicable to them.

The requirements from one legislative system may be inconsistent with another, making compliance with all of the relevant laws of multiple jurisdictions difficult. Understanding legislation can be complex, and pieces of the legislation can sometimes conflict with one another. The ISO/IEC 27000 series provides organisations with guidance regarding compliance with legal requirements and covers the following areas:

- intellectual property rights;
- protection of organisational records;
- data protection and privacy of personal information;
- prevention of misuse of information processing facilities;
- regulation of cryptographic controls.

Non-compliance with legislation could prove costly for an enterprise through incurred financial penalties, operating restrictions or, in extreme instances, custodial sentences

for senior executives. Therefore, it is essential that advice is taken from a trained legal specialist before making any important decisions that could put the organisation at risk.

Protection of personal data and restrictions on monitoring, surveillance, communications interception and trans-border data flows

Privacy laws exist to protect the rights of the individual. Most organisations hold and process information about people such as employee or customer information. Organisations need to be aware of the legal restrictions placed on them to protect this information and how it may be used and monitored. The latest EU regulation covering data protection is the GDPR, which came into effect in 2018. In the UK this was supported by a revised Data Protection Act that spelt out how the EU regulation would take effect in the UK and also in the event of the UK leaving the EU. Many countries have legislation to protect the individual and restrict and control the amount of information held and how it should be used and monitored, but the GDPR was arguably the first comprehensive legislation that now affects all the personally identifiable information (PII) about any citizen of the EU regardless of where the information is held.

Although all the various forms of the legislation do share some common principles, there are significant differences in the legislative approaches, and this can cause difficulties when working across different legal jurisdictions. The EU has a legal framework via the GDPR to protect all types of personal information, whereas the USA protects personal information via a number of federal statutes. These tend to target specific areas such as protection of customer information by financial institutions (via the Gramm-Leach-Bliley Act) or preserving privacy of medical information (Health Insurance Portability and Accountability Act (HIPAA)).

The scope of legislation can also vary from country to country. For example, the Canadian Personal Information Protection and Electronic Documentation Act applies to the records of people for up to 20 years after their death, whereas in the EU any protection ceases at the time of death.

The GDPR within the EU protects the individual by ensuring that information is collected and processed lawfully, is accurate and is appropriate. Individuals have the right to have access to information held about them, know who can access it and have any inaccuracies amended. The Act also has provision to ensure that information is only collected (in general) with the explicit approval of the individual and that it is handled and processed in a secure manner, while placing controls on transferring it out of the EU to countries that have less stringent privacy controls. It also covers both electronic and paper-based records. The main points to remember when handling personal information are:

- Personal information must be surrounded by proper robust assurance controls and working practices to protect the data from unlawful processing, accidental loss, corruption or destruction and unauthorised disclosure.
- Processes should be implemented to ensure that information is entered into computer systems correctly and that staff understand that no personal information should be disclosed to any third party without the appropriate written authority being in place.

- Paper records should be kept locked away and computer screens should not be left displaying personal information or able to be overlooked. Information no longer required should be destroyed by shredding or other secure forms of destruction.

Privacy laws often place restrictions on transferring information between countries. For example, the EU's GDPR states that personal information must not be transferred to countries that do not have such similarly strict rules. Certain countries, such as Argentina, Canada, New Zealand and Switzerland have already (at the time of writing) shown that they operate a data protection model which is comparable to the EU GDPR model, and there are no restrictions with these countries. Doubtless, more countries will follow suit over time. For other countries, safeguards need to be considered to enable trans-border data flows to take place legally.

The European Commission and the United States Department of Commerce developed the Privacy Shield framework to enable American organisations to be compliant with the EU's GDPR privacy legislation. American companies that are likely to exchange personal information with EU-based organisations can amend their assurance arrangements so they are compliant with the Privacy Shield framework. For other countries, such as for transfers to the Indian subcontinent, it may be possible to make the data transfer if a number of prior safeguards are put into place. These safeguards would be undertaken via approved contractual terms that are acceptable to the legal body responsible for protecting personal information.

In the UK, the individual has a right to a level of privacy, protected by legislation, that restricts how their personal information can be monitored or intercepted. This legislation may often predate the computerisation of data storage such as the UK's Public Records Acts of 1957 and 1967. Any monitoring and collection techniques employed by an enterprise must comply with these laws. If the monitoring controls are to be used across more than one legal jurisdiction, then there may be differences in the rights of the individuals being monitored. Within the UK, the Regulation of Investigatory Powers Act 2000 (RIPA) was enacted to restrict covert monitoring of an individual's information. It was introduced to take account of new developments in communications technology, the Human Rights Act and the Telecommunications Directive.

Employment issues and employee rights

Depending on the legal jurisdiction, employees have certain rights when using the enterprise's information systems, such as the right to privacy and the right to know what information is held about them by the enterprise. In the UK, for example, under the GDPR, individuals can request a copy of any information that any organisation may hold on them. This is called a Subject Access Request.

Rights may also extend to monitoring controls. Within the EU, employees have the right to know the type and scale of monitoring that is being carried out by the enterprise and why it is being done. For instance, an employer might consider it necessary to protect the enterprise from offensive or pornographic material, or perhaps they need to understand the volume of email traffic being propagated for performance purposes. The enterprise must communicate this information to employees. The easiest way to do this is to include a statement about the extent of monitoring in the enterprise's information assurance policies or employment contracts. If this is not done, it may be necessary

to gain specific consent from individuals to allow their information to be collected and/or monitored. An assessment of the monitoring strategy should be carried out to demonstrate that the monitoring techniques that are being used are justified, not excessive and meet legal requirements.

If monitoring tools detect information that is clearly personal, then care must be taken not to violate the individual's right to privacy. For example, a clearly personal email should not be opened by the employer or an individual's email account should not be accessed unless agreed with them beforehand. These can cause operating issues. For instance, a colleague may need to quickly obtain important company information previously sent to an individual while that person is away on holiday. Employees should be asked to remove any of their personal information from IT resources when they leave the enterprise to avoid it being viewed by others.

Monitoring tools should not be used to target any particular individual and covert monitoring is rarely justified; exceptions might include situations where there are clear grounds that criminal activity or malpractice is taking or has taken place. Internal investigations may lead to an employment tribunal or a court case. It is important that any information that is collected meets the legal requirements, which might include, for example, the UK's Freedom of Information Act or the equivalent in other countries.

Common concepts of computer misuse

Much of the legislation that currently applies to the misuse of computers has not been written specifically to address computer crime. It can be said that crime is crime and criminals simply use whatever means are available to carry it out. Blackmail, fraud, deception, theft and so on have always existed, but developments in technology have enabled criminals now to exploit computing devices in their activities. Similarly, privacy rights can be abused via electronic eavesdropping, hacking or cyber stalking, rather than by an actual physical presence. Therefore, existing laws that predate computers are often used to prosecute computer misuse.

Legislation has been produced to specifically target crimes committed using computers. The USA introduced the Computer Fraud and Abuse Act in 1984 and this legislation has since undergone several amendments. The UK was the first European country to enact a law that specifically addressed computer crime, and this legislation formed the basis of the EU Directive on Computer Misuse. The Computer Misuse Act 1990 introduced three new offences: unauthorised access to a computer; unauthorised access with the intent to commit or facilitate further offences; and the unauthorised modification of computer material. The misuse of computers can include:

- illegal access (hacking) to computer systems;
- illegal interception of information;
- interference with information and systems;
- computer-related fraud and forgery;
- commercial infringement of copyrights;
- download of illegal material such as child pornography;
- trafficking in passwords, digital signatures and encryption keys.

The motives for the misuse of computers can vary. Fraudsters may misuse a computer for financial gain; hackers may try to gain access to a system for the intellectual challenge; a disgruntled employee may sabotage a computer system as an act of revenge.

Computer fraud is the term used to describe stealing money or goods by using or involving a computer. There are various ways that this can be achieved, either by entering incorrect information or by altering the information already held on a computer. It can also be carried out by creating or altering computer code. Misuse of computers in this manner is a major problem as organised criminals continually find new opportunities to use computers to commit fraud. Business fraud, as this is often labelled, is now one of the most common ways of criminals making money from business. In a Bromium Inc. sponsored report researched and written by Dr Mike McGuire in April 2018, titled *Into The Web of Profit*,[2] Dr McGuire estimated that the theft of intellectual property (IP) and trade secrets alone generates $500 billion each year for the criminals.

The increasing use of the internet to trade and shop enables criminals to commit fraud and steal the identities of others to commit fraud. Phishing is the term used where criminals entice individuals to disclose their financial details, for example by sending an email to an individual that purports to have been sent by their bank or perhaps by a senior member of staff in their employer's company. This constitutes obtaining information by deception.

Hacking (despite its benign origins as a term for a general interest in discovering how computers work) is the term given to accessing a computer system without the express or implied permission of the owner of that system. A hacker is the name given to the person that carries out this activity. Hackers often modify information or software programs – which can subsequently cause considerable havoc. Website defacement is an example of where a hacker changes the information displayed on a web page. Sometimes, hackers will change information held within a database. Hackers often gain unauthorised access to a computer system simply for the thrill of being able to circumvent assurance controls and then share their conquests with other like-minded individuals. However, hacking is now increasingly being used as a tactic by criminals to carry out crimes such as fraud or blackmail. Notably one of the most common targets for this type of activity are the systems running out-of-date software or where the latest patches to address known vulnerabilities have not been implemented.

Malicious code (or malware) is the term used to describe programs that have been written to cause security breaches or damage to computer systems by installing unwanted and unauthorised code onto them. Malicious code can cause a number of undesirable impacts including the deletion or corruption of information, the capture of information or the hijacking of computer resources to launch further attacks onto other computers in DDoS attacks. Malicious code comes in a variety of forms such as viruses, Trojan horses and backdoors. Malicious code is being increasingly used by criminals, especially to capture financial information held on a computer that can be used for fraudulent purposes. Ransomware is an example of malware that infects the target computer by encrypting the owner's personal files. The victim is then contacted

2 https://www.bromium.com/resource/into-the-web-of-profit/

and offered the key to decrypt the files in exchange for cash or information. Dr McGuire estimates that ransomware adds a further $1 billion annually to the criminals' gains.

The download of illegal material onto a computer is another form of computer misuse. Many countries have in place legislation that prohibits the download of child pornography and of the 'sexual grooming' of children using the internet. In many countries there is a legal obligation for enterprises (and individuals) to report the discovery of this type of activity to the law enforcement agencies. It is very likely that legislation such as the Obscene Publications Act in the UK would be used to prosecute cases of child pornography, as the penalties are more severe than in computer misuse legislation.

Computers can be misused by a person to harass and stalk another individual (cyber stalking), for instance by sending threatening emails that cause distress. It is essential to ensure that policies are in place to provide clear guidance to all computer users as to what constitutes computer misuse. Cyber stalking of this nature is now recognised in the UK under the Public Order Act, the Malicious Communications Act and the Protection from Harassment Act.

The illegal or unauthorised use of software such as programs, computer games or electronically stored music is known as piracy and is another example of computer misuse. Only legitimate software and material used in line with licence agreements should be installed on an organisation's systems and guidance to computer users that the use of unlicensed material is not allowed should be provided.

Requirements for records retention

Certain documents or records need to be retained by an organisation for legal or regulatory purposes for a period of time. These can include company board minutes, financial reports and accounts or technical specifications. The duration for which documents need to be retained varies by the document type and the legislation of the country in which it is being used. In multinational organisations, records may be passed over to other countries within the same enterprise – meaning that the same data are then subject to different legislation requirements, which might even conflict with one another.

Although most retention requirements state a minimum length of time for keeping data, some legislation conversely states when a record must be destroyed. These usually relate to personal privacy, such as the GDPR (EU) or the Fair and Accurate Credit Transaction Act 2003 (USA).

An organisation may be asked to produce these records (or proof of destruction) either by a government agency or by an opposing party in a legal dispute. Failure to comply with this could result in a legal judgement against the organisation, heavy fines, and closure of business or adverse publicity.

To help in compliance with legislation, it is necessary to have in place a record retention policy and schedule. This should be communicated to staff so that they are aware of their responsibilities. In the case of international organisations, more than one schedule will need to be kept to deal with variances in requirements. A document that needs to be retained should be stored in a format that can ensure its protection (for example on

a secure central repository rather than in a personal file so that it is not deleted or lost inadvertently). For larger enterprises there are document management solutions on the market to do this.

There are a number of externally produced standards available to help enterprises understand how best to cope with legal requirements, such as ISO 15489-1:2016 – Record Management Standards produced by the International Organization for Standardization, or standards produced by the American National Standards Institute (ANSI).

Intellectual property rights, for example copyright, including its general application to software and databases

Individuals and enterprises invest a lot of time, money and effort in creating original works, products, methodologies and ideas. They can be significantly out of pocket if they are unable to realise the benefit of their investment because other parties have used their ideas without compensating them. Intellectual property rights (IPR) is the term given to the legal rights that protect creative works, and most countries have legislation in place to protect such intellectual property.

Copyright law was initially designed to protect original artistic works such as pieces of music, but its use can also be applied to software programs, computer games, documents, books, photographs, video files or other types of work made using a computer or generated by a computer. Copyright is automatically associated with the piece of work deemed as original upon its publication. It usually remains in place for a fixed duration, such as 50 years in the case of music. Copyright gives the creator exclusive rights over certain aspects of the work such as copying, issuing, performing or adapting it. Abuse of these rights by someone else is called infringement. Piracy is the term commonly used to describe the unauthorised use of computer software and is a breach of copyright law. Where software has been developed by an enterprise, the copyright is normally owned by the enterprise rather than the individual(s) involved, unless a special provision has been agreed beforehand.

Copyright legislation is prevalent in most developed countries but there are some countries that take copyright less seriously, such as in Asia and the Far East. There have been a number of initiatives to harmonise copyright protection internationally, such as the General Agreement on Tariffs and Trades, Trade Related Aspects of Intellectual Property Rights 1993 (GATT TRIPS). Within the EU there is a directive to harmonise certain aspects of copyright and associated rights in relation to information systems.

In addition to copyright there are other pieces of legislation that aim to protect intellectual property, and it is useful to have an awareness of some of them. The Common Law of Breach of Confidence (often described as a tort) aims to protect secrets – personal, commercial or governmental. These can only be applied for as long as the data are not in the public domain and covers breaches of confidence made between two or more parties. Trademarks, such as Microsoft® or Apple®, are there to protect brand strength by demonstrating their uniqueness in terms of quality, reputation, reliability, ubiquity, originality, value for money or whatever the brand strives to promote.

'Passing off' is the term used when an object is trying to seem the same as something else in order to cash in on the originator's reputation or ideas. This legislation is intended

to protect the public from deception and to stop misrepresentation. An example of infringement that could apply to information assurance is when someone has set up an internet domain name that uses a very similar name to another, better-known site. The 'impostor' site is branded in much the same way as the original so that people who have mistyped the address believe they have gone to the intended site. This is a typical method of perpetrating banking fraud. Dr McGuire estimates that the trade in illicit and illegal online markets generates $860 billion each year for the criminals.

Patents are used to protect the intellectual property invested in the development of new products or in the creation of inventions, and, like copyright, they are in place to prevent other people from copying or manufacturing the product or invention so that the creator is able to realise their investment (in both time and money) in creating their original work. Increasingly, patents are applied to software processes and such things as 'gestures', for example, on tablet computers. Within information technology, patents tend to be used to protect physical devices such as a new type of computing device. Patents are of a fixed duration, but can often be renewed by the owner. Patents tend to apply only to the particular country in which the application has been made. Extending patents to cover many countries can be expensive as multiple applications may have to be submitted, and so expert advice should be sought. However, there is provision within the EU to protect a patent in 30 countries in a single application, using the European Patent Convention (EPC).

Contractual safeguards, common security requirements in outsourcing contracts, third-party connections, information exchange and so on

When developing contracts with third parties it is important to ensure that controls are put in place to protect the information assets of the enterprise to an acceptable level. In effect, it is necessary to ensure that a third party would take the same level of care in protecting an organisation's information as the organisation would internally. The types of safeguards required will vary depending on the type of service being provided and the sensitivity of the enterprise data.

Contract conditions should include clauses to ensure that proper assurance controls are in place. Security conditions are often handled via a security schedule within the contract. The type of clauses needed to provide adequate protection might include clauses to:

- carry out regular assurance reviews and health checks;
- apply security patches in a timely manner;
- protect information against malicious code;
- provide business continuity arrangements that meet agreed service levels;
- vet new staff to an appropriate level;
- enforce discipline against any security breaches;
- manage security incidents (including reporting any incidents to the parent organisation);
- protect against disclosure of sensitive information;

- allow the enterprise the right to audit and monitor the services being provided;
- prevent further subcontracting without written authorisation.

Many organisations use cloud computing services. It is essential that the organisation understands the services that are being bought and contractually protects its information adequately. This is discussed in more detail in Chapter 6.

Collection of admissible evidence

There are a number of rules and processes that need to be followed when collecting evidence so that it can meet certain criteria when used in a court of law (described as admissible evidence). If legal guidelines are not followed, the evidence may be excluded as being inadmissible. This could result in a court case being lost, adverse publicity, embarrassment and financial penalties to the prosecuting party.

This generally means being able to demonstrate that the evidence is authentic, has not been tampered with in any way and has been gathered in an acceptable manner that meets legislative requirements, which includes being able to retain and document the state and integrity of items at the crime scene. Most countries have produced legal requirements that specify how evidence should be handled. Examples are the Federal Rules of Evidence (in the USA), the Police and Criminal Evidence Act and the Civil Evidence Act (in the UK). The appropriate guidelines should be followed when collecting evidence.

Developing a procedure for dealing with investigations and gathering evidence will help to avoid mistakes when working under pressure. Only trained personnel should carry out the securing of evidence. Some organisations, such as banks, may have an in-house facility for carrying out investigations as they may need to do this regularly. In other organisations, where investigations are rare, it may be better to call in an external specialist organisation to manage the investigation and collect the evidence.

Each person that has handled any evidence may need to testify in court that the evidence is in the same state as when it was processed during the investigation. Therefore, keeping the number of people involved in the investigation to a minimum helps to simplify the presentation of evidence and preserve confidentiality. Evidence needs to be presented in a form that is understandable to the judge, jury or adjudicator.

Evidence can be excluded because it was gathered without the correct authorisation or in a manner that contravenes guidelines. Sometimes a warrant may be required to seize evidence, although this is not necessary if the evidence is in plain view, for example, or consent has been given by the individual.

Collection of digital evidence can be complex and can come in a variety of forms, such as audit trails, application logs, firewall logs and closed-circuit television (CCTV) footage. Some of the information may be deleted, incomplete or partially overwritten. It is more difficult to isolate and preserve evidence from a communications network as it is in a state of constant change and the sources of evidence may reside in different locations. As mentioned previously, The *Good Practice Guide for Digital Evidence* (V5) published by NPCC (previously ACPO) in the UK, provides detailed advice on recovery of computer-based evidence. The guide states four principles when handling digital evidence:

- No action taken by the police or their agents should change data held on a computer or other media that may subsequently be relied on in court.
- In exceptional circumstances when a person finds it necessary to access original data held on a target computer, that person must be competent to do so and to give evidence explaining the relevance and the implications of their actions.
- An audit trail or other record of all processes applied to the computer-based evidence should be created and preserved. An independent third party should be able to examine those procedures and achieve the same results.
- The officer in charge of the case is responsible for ensuring that the law and these principles are adhered to. This applies to the possession of and access to information contained in a computer. They must be satisfied that anyone accessing the computer, or any use of a copying device, complies with these principles.

Use is often made of organisations that specialise in recovering and securing lost data from computer disk drives.

Securing digital signatures

Traditionally, a handwritten signature on an original document proves who signed it and any alterations can be detected. In the electronic world the original is indistinguishable from a copy and therefore there is potential for fraud. Digital signatures are a form of electronic signature that addresses this problem. A digital signature electronically binds the sender of a message to the contents of the actual message to prove that it is genuine. It also proves when it was sent, to whom it was sent, that it has not been tampered with, that it has been kept confidential and that neither party can deny its transmission. Enterprises are increasingly using digital signatures to conduct their business, and legislation has been developed to facilitate and control their use. However, what is acceptable varies across legal jurisdictions, so it is important that legal advice is obtained before adopting the use of digital signatures.

Within the EU, the legal regulation Electronic Identification, Authentication and Trust Services (eIDAS) came into force on 17 September 2014 and states that electronic signatures will not be denied legal effect or admissibility simply on the grounds that they are in electronic form. Electronic signatures will be treated as handwritten signatures if they are backed by qualified certificates, which are provided by a certification service provider and created by a secure signature creation device. Electronic signatures are admissible as evidence in legal proceedings both in relation to the authenticity of the transmission and as to the integrity of the contents of that communication. Applications supporting digital signatures are now available on smartphones and tablet computers as well as desktop and laptop computers.

There is much emphasis on the certification authority (CA) to be deemed trustworthy. CAs generally have their signatures verified by other CAs to build a greater degree of trust. CAs may be liable for any compromise to the integrity of digitally signed documents authorised by them, so, to limit their liability, many certification authorities stipulate a financial cap on transactions. There have been malicious attacks on digital certification companies, though, and this has led to some significant improvements in recent years.

Some legal jurisdictions control the extent to which foreign certificate authorities can issue certificates that meet local laws. In the EU, there has to be a recognised arrangement between the EU and the country that has issued the certificate.

Restrictions on purchase, use and movement of cryptography technology

Cryptography is a powerful tool for protecting privacy that can be used by businesses, governments, criminals and individuals to protect confidential information. Governments argue that it is in the national interest for them to control cryptographic activity in order to protect the individual and to prevent and track criminal or terrorist activity. As such, there are numerous controls in place over its use. Cryptography legislation varies greatly from country to country. In some countries the controls are quite draconian, especially where repressive political regimes are in government. It is important that organisations which operate internationally understand the local operating restrictions as penalties can be extremely harsh (i.e. for treason) and the death penalty is included in some statutes.

In China, foreign organisations and individuals have to gain permission to use cryptography under the China State Council directive 273 of the Regulation of Commercial Encryption Code. In Pakistan, all encryption hardware and software has to be inspected and approved by the Pakistan Telecom Authority. Even within the EU, there is variance on acceptable use of cryptography. France, for example, has a number of very specific requirements as to how cryptography can be used that are in addition to the EU directives.

The export of cryptographic controls is controlled in many countries by the Wassenaar Arrangement (WA) 1996. The purpose of this agreement was to ensure that transfers of conventional firearms and dual-use goods and technologies between countries were carried out responsibly and did not further the development of hostile regimes. There are, at the time of writing, 42 participating countries and, although export controls are implemented by each individual WA participating state, the scope of export controls is determined by Wassenaar directives.

Cryptographic controls should be used in compliance with all relevant agreements, laws and regulations. Local legislation may disallow or restrict the use of strong cryptography. It may restrict its sale and movement to another country or put in place regulatory controls and registration requirements for its use. You should seek legal advice to ensure compliance with any national laws and regulations and when transferring encrypted information or tools from one legal jurisdiction to another.

The ISO/IEC 27000 series advises that the following factors should be considered:

- restrictions on import and export of computer hardware and software for performing cryptographic functions;
- restrictions on import and export of computer hardware and software that is designed to have cryptographic functions added to it;
- restrictions on the use of encryption;

- mandatory or discretionary methods of access by the countries' authorities to information encrypted by computer hardware and software to provide confidentiality of content.

For organisations that are subject to regulatory control, the associated regulatory body may define additional constraints on how cryptography should be used; for example, the finance industry may specify use of particular security standards. These factors also need to be considered when applying cryptographic controls.

You should be aware that legislation is in a continual state of change and improvement, and should verify for yourself that you are using the latest legal instruments.

ACTIVITY 3.6

Following the information leak, Ms Jackson is worried that GANT's controls for protecting personal information may be weak. She has asked you to carry out a review of the privacy legislation affecting GANT to ensure that the organisation is compliant. What would be the main areas that you would look at?

SECURITY STANDARDS AND PROCEDURES

Standards affect many aspects of our daily lives and they have been developed to ensure that products and services are safe, reliable and efficient. They also provide for interchange ability (interoperability between products and systems). Today, there are recognised standards in place for almost all aspects of commerce, industry or government, and information assurance is no exception to this. Standards are produced by recognised standards bodies and they enable organisations to demonstrate a requisite level of technical, operational or administrative competency. There are many standards and technical regulations in existence and many organisations worldwide producing them. This section will look at how externally produced standards affect information assurance management within the enterprise.

LEARNING OUTCOMES

The intention of this section is to provide you with the basic knowledge of some of the main externally produced standards that apply to information assurance management. Following study in this area, you should be able to explain and justify each of the following concepts.

National and international standards

In the area of information assurance there are many standards that apply. These typically define a set of requirements for products, processes or procedures and they are produced by organisations known as standards bodies. They collaborate with

industry experts in different areas, whether representing vendors, scientific research agencies or government departments, to produce good practices that can be applied by others. The jurisdiction of a standards body may extend to a specific industry sector, a particular country or internationally. The standards that will apply to an enterprise will vary depending on a number of factors, which may include the actual country in which the enterprise is based, whether it works internationally, the industry sector in which it operates, or perhaps engagement in government contracts. Most standards are produced by non-profit making organisations and are funded by the various parties that have a vested interest in their existence. Typically, they do not actually regulate the adoption of their standards, although some do provide certification or accreditation to organisations to allow them to demonstrate compliance to the set standards.

Although not mandatory, failure to implement or comply with accepted standards may have a significant adverse effect on an organisation. For instance, if a product or service is not certified as being compatible with other vendors' offerings, then this may dissuade potential customers from placing orders. Alternatively, the organisation may not be able to demonstrate a sufficient level of competency in managing particular processes. Therefore, it is important to understand the standards that can or need to be applied within the enterprise. The remainder of this section will look at some of the common standards related to information assurance as well as the standard-producing organisations themselves.

The International Organization for Standardization (ISO) is the world's largest developer of standards and has published over 22,500 international standards since it was established in 1947. The organisation was founded to facilitate the international co-ordination and unification of industrial standards. ISO standards are mainly technical and cover a wide number of sectors, including agriculture, construction, engineering and information technology. ISO standards are developed collaboratively by committees from more than 160 participating countries. Each standard is reviewed at least every five years to ensure that it remains current and those that are no longer relevant can be withdrawn. Editions of these standards are formally published by the ISO and can be purchased either from ISO directly, or through national standards agencies such as BSI in the UK.

The ISO works in collaboration with two other international standards organisations, the International Electrotechnical Commission (IEC) and the International Telecommunication Union (ITU), to form the World Standards Cooperation. These organisations have been extremely influential in producing standards that affect the information technology industry. The naming convention for standards approved jointly by IEC and ISO is ISO/IEC and there are a number of joint standards that directly affect information assurance management. The most significant is the ISO/IEC 27000 series, which is the current set of standards for information security management. There are, however, many others that cover explicit technologies, techniques or architectures used within IT.

The two main standards in this series are ISO/IEC 27001 and ISO/IEC 27002. Others have been developed within the series to cover specific topic areas or industry sectors. For example, the series includes a standard on risk management (ISO/IEC 27005), network security (ISO/IEC 27033), guidelines for managing information security in the health sector (ISO/IEC 27799) and using cloud resources (ISO/IEC 27017).

ISO/IEC 27001 specifies the information security management system requirements standard. Organisations can be formally certified against it and this will be covered

in more detail later in this chapter. ISO/IEC 27002 provides a code of practice for information security management. It is probably the most influential standard for information assurance management. It describes a high-level set of controls to protect the confidentiality, integrity and availability of an organisation's information assets, and looks at the various aspects of assurance such as security policy, information assurance organisation, asset management, human resources assurance and compliance. ISO/IEC 27002 is a generic advisory document rather than a formal specification like ISO/IEC 27001. In order to use the standard, the organisation will need to assess the risks of their enterprise and apply the recommended control measures from the standard that are applicable to mitigate these risks. ISO/IEC 27007:2017 provides guidelines for auditing information security management systems, helping auditors to assess the compliance with ISO/IEC 27001.

Other international organisations produce standards on security management; for example, the ISF's *Standard of Good Practice* for information assurance, which is reviewed by them every year. This standard focuses on how information assurance can support an organisation's business processes and provides guidance on implementing appropriate protection. It focuses on security governance, security requirements, control framework and security monitoring and improvement. The ISF membership (who are in the main corporate organisations) fund the forum via an annual subscription, and then collaborate with them in developments for best practices in IT security and information risk management. However, the *Standard of Good Practice* is available to other organisations or individuals.

There are many other standards not produced exclusively for information assurance that do affect information assurance management. These cover other related business functions or processes such as retention of records (ISO/IEC 15489), the implementation of business continuity (ISO/IEC 22301:2019 and PAS 77), project development (COBIT), the management of information technology services (ISO/IEC 2000 ITIL) and quality assurance (ISO/IEC 9001). If these standards have been implemented within an organisation, it is necessary to ensure that information assurance management controls are compatible with them and support their requirements. For example, ISO/IEC 20000-1:2018 Information technology – Service management includes requirements for operational security.

It is necessary to be familiar with the standards that apply to the country and industry sector in which the organisation operates. There are a wide range of standards to which the financial and manufacturing industries need to adhere, and an inability to meet these requirements could, for instance, prevent an organisation from actively trading in the financial services market.

Certification of information security management systems to appropriate standards

Gaining information assurance certification is a means of demonstrating that an organisation takes information assurance seriously and that good assurance processes and controls have been implemented. An increasing number of organisations now look for certification in their trading partners and, for some, certification can be a prerequisite for doing business. Certifications can apply enterprise-wide or to a specific

set of processes within the organisation. Certification usually involves the enterprise undergoing an external audit by an accredited third party.

The ISO runs a number of certification schemes against its standards, including ISO/IEC 27001, which enables an organisation to have its information assurance governance and management processes certified against ISO/IEC 27001. To gain accreditation, the organisation's ISMS has to undergo an external audit carried out by an accredited third-party organisation. The auditors use standard processes to check the organisation's ISMS policies, standards and procedures against the ISO/IEC 27001 requirement and then look for evidence that they are being used within the organisation. The findings from the audit are reported back to the organisation and certification will be granted if successful. After the initial certification, periodic follow-ups (reassessments) will take place to ensure that the standards are still being met. There is also an ISO standard (ISO/IEC 27006) that is used to guide the accredited certification bodies on the formal processes for certifying or registering other organisations' information assurance management systems.

In the UK, industries such as financial services require that certification to certain standards are in place, and serious instances of non-compliance can lead to sanctions from the FCA, such as heavy fines or withdrawal of their registration. An example of an applicable standard is the PCI DSS. It was originally implemented in the USA, but has now been extended to include many other countries. It was introduced to reduce credit card compromise and to deal with increasing cases of fraud. It specifies the protection measures that organisations processing card-payment information must put into place. Failure to comply may result in an organisation receiving substantial fines or steps being taken to prohibit the organisation using cards as a payment method. This could seriously jeopardise the organisation. PCI merchants, retailers and service providers must store credit card account data securely as specified by the standard and demonstrate compliance to their acquiring (merchant) bank. All acquiring banks need to have certified proof of PCI compliance from their merchants or they will be liable to fines themselves from the FCA.

Product certification to recognised standards

Many products require independent testing and certification before they can be launched onto the market to ensure that they conform to safety requirements, technical specifications or other compliance regulations. It is useful to have an independent third party to verify that a new product does meet expectations and that it can be trusted. This particularly applies to security products, as it is often difficult for the consumer to be able to test the security of the product for themselves. Certificates provide customers with the assurance that the security features do offer the level of protection that is claimed by the vendor. It is helpful to know that a standards-based approach has been used to do this evaluation, as this will aid understanding as to how rigorous it has been. Test results produced in a standardised format will enable straightforward comparison with other competing products.

Security testing, evaluation and certification have mainly been carried out by either government agencies or organisations serving the defence market. Different countries have developed their own evaluation and certification systems using a variety of classification models and approaches. This has often made life complex when dealing with other internationally recognised certification schemes. It has meant that products have to be recertified each time for use in different countries or industry sectors, which has exacerbated an already time-consuming and expensive process.

Over recent years the numerous certification schemes have been streamlined. In the 1990s the EU harmonised the various schemes of its member states into the Information Technology Security Evaluation Criteria (ITSEC). More recently this scheme was itself harmonised with other models, such as the Canadian Trusted Computer Product Evaluation Criteria (CTCPEC) and the US Federal Trusted Computer System Evaluation Criteria (TCSEC – often referred to as the Orange book) to form the Common Criteria for Information Technology Security Evaluation Criteria (CC ITSEC). This evaluation system has become the internationally accepted approach for security certification, replacing national and regional systems. In 1999 the ISO incorporated its evaluation criteria into a standard that was revised in 2009 to the current ISO/IEC 15408-1:2009 Information technology – Security techniques – Evaluation criteria for IT security – Part 1: Introduction and general model. This standard, together with parts 2 and 3, is now the most widely used certification model and, at the time of writing, it is currently under review for a second time before being reissued.

The ISO/IEC 15408 standard specifies a number of functionality and assurance classes that can be tested. The standard has seven levels of assurance to describe the level of rigour used to carry out the testing, from the entry level of Evaluation Assurance Level (EAL) 1, which tests claimed functionality to the highest classification, to EAL 7, which provides a formally verified design that has been subjected to rigorous testing. The higher the classification, the more complex and rigorous the testing is. Gaining classifications EAL 5–7 is generally less common for commercial organisations, as testing requires very specialist security engineering techniques and is complex. These higher classifications tend to be used by military and government organisations.

The developer of the product or system has to define what is being submitted for evaluation, known as the target of evaluation (ToE), and specify the assurance level for which they are aiming. Certification is carried out by an approved testing agency. Within the UK, security certification is managed by the NCSC, part of the Government Communications Headquarters (GCHQ) at Cheltenham. This is the UK Government's national technical authority for information assurance.

An NCSC Tailored Assurance Service (CTAS) evaluation provides a view of assurance on the IT security attributes of a system, product or service and will be carried out by a company with an NCSC-approved test laboratory with CTAS capability. The scope of the evaluation is specified in a security target and the range of evaluation activities is detailed in an evaluation work programme. The accreditor and NCSC, with other key stakeholders, will agree the scope and technical approach of the evaluation and will review the CTAS activities and results documented in an evaluation report. At the end of the evaluation, NCSC will issue a CTAS assessment statement on the results of the evaluation, making recommendations on the significance of any issues that are discovered. CTAS maintenance provides continued assurance to evaluated

configurations by understanding and assessing changes in an efficient and tailored manner.

To enable the international use of existing certificates, agreements have been put into place to enable security certificates to be recognised by other countries. For example, the Common Criteria Recognition Arrangement (CCRA) enables Common Criteria (CC) certificates up to EAL 4 to be recognised within all participating countries.

NCSC manage a number of certification schemes, including the Independent Evaluation for Assured Services (CAS), Commercial Product Assurance (CPA) and Certified Assisted Products (CAPS) for any commercial product and its developer. They also manage the Cyber Essentials scheme for companies that helps small and medium-sized enterprises to ensure the very basic cyber security controls are in place and effective.

Awareness of the production of key technical standards

There are several technical standards applicable to information assurance management. This section will examine some of the more well-known technical standard producing bodies.

The Internet Engineering Task Force (IETF) is a large, open international community that develops and promotes standards for the internet. Its governing body meets two or three times a year. Standards are developed by working groups of interested parties, such as network designers, operators, vendors and researchers, that each focus on a particular topic. The standards generated are known as Request for Comments (RFCs), and upon production are subsequently issued to the IETF community as draft RFCs for comment and review. Once an RFC has been issued it is not withdrawn, although it may in time be superseded by further RFCs. This, in many ways, can show the development of standards. The published RFC documents have a status of either a proposed standard or an informational statement.

Federal Information Processing Standards Publications (FIPS PUBS) are standards and guidelines developed and issued by the NIST for federal government computer systems within the USA. Where possible, the US federal government uses existing (internationally recognised) published industry standards, but should none be suitable it will ask NIST to help develop them. NIST collaborates with national and international standards committees such as IETF and other interested parties (such as vendors and industry bodies) to produce FIPS PUBS.

Within Europe, the European Telecommunications Standards Institute (ETSI), based in France, has official responsibility for standardisation of information and communications technology (ICT). It is recognised by the European Commission and the European Free Trade Association (EFTA) secretariat. Its main purpose is to provide technical specifications (or standards) that may be used in European directives and regulations or by manufacturers to show that their products are compliant with these directives and regulations. Products demonstrate conformance by attaching the 'CE' mark on their goods. ETSI members represent areas that have a vested interest in the process and include manufacturers, network operators, administrations, service providers, research bodies and users. They come from a wide selection of countries both inside and outside Europe. The members determine the Institute's work programme, allocate resources

and approve its deliverables. Documents can be downloaded from the ETSI website via their documentation service (EDS).

The European Union Agency for Network and Information Security (ENISA) also provides guidance and advice, based sometimes on its own scientific research work. Most recently at the time of writing, it has been looking at the human aspects of information security and considering the psychological, societal, ethnography, anthropology, human biology, behavioural economics and any other subject that takes humans as its main focal point.

SAMPLE QUESTIONS

1. **Which of the following activities should *NOT* be handled by the information assurance function?**
   a. Monitoring the effectiveness of the enterprise's assurance arrangements.
   b. Providing advice on information assurance.
   c. Effectively delivering a secure environment across the enterprise.
   d. Reporting on the effectiveness of the enterprise's assurance arrangements to senior management.

2. **Where should the information assurance function be placed within the enterprise so that it can facilitate full management co-ordination of assurance across the enterprise?**
   a. Within the compliance function.
   b. At board level.
   c. It will depend on the structure of the enterprise.
   d. Within the IT group.

3. **What is the main role of the board director with responsibility for information assurance?**
   a. To ensure that appropriate security controls are implemented across the enterprise.
   b. To have a detailed understanding of the threats facing the enterprise.
   c. To implement information assurance solutions across the enterprise.
   d. To provide day-to-day management of the information assurance function.

4. **Clearly defined responsibilities for information assurance should include which of the following?**
   a. Operating procedures and reporting requirements.
   b. The scope of the responsibilities and level of authority granted.
   c. Disciplinary procedure.
   d. None of these three.

5. **Which would be the best way to hear about and plan for any regulatory changes to your industry that may affect information assurance?**
   a. Permanently employing consultants.
   b. Scanning bulletin boards and websites for snippets of information.
   c. Waiting until the changes are announced in the press.
   d. Maintaining a relationship with regulatory bodies for the industry.

6. **Which of the following groups of people should have access to the high-level security policy for the enterprise?**
   a. Senior management and all line management.
   b. All staff within the enterprise.
   c. Third parties that have access to the enterprise's information systems.
   d. All of the above.

7. **Which of these security documents is *NOT* mandatory?**
   a. A policy.
   b. A standard.
   c. A guideline.
   d. A procedure.

8. **Which of the following statements best describes an information security architecture?**
   a. A technical overview of assurance controls applied within the enterprise.
   b. A framework of assurance controls that can be applied across the enterprise to protect its information assets.
   c. The physical security controls applied within security locations.
   d. A blueprint for future security controls.

9. **Which of the following is the security standard that applies to the accreditation of security controls within products?**
   a. ISO/IEC 27001.
   b. ISO/IEC 15408.
   c. ISO/IEC 9000.
   d. ISO/IEC 13335.

10. **Privacy legislation is in place to protect the rights of:**
    a. Criminals.
    b. Companies.
    c. The individual.
    d. Data protection officers.

11. Which of the following is *NOT* a phase in incident management?

a. Assessment.
b. Investigation.
c. Reporting.
d. Elimination.

REFERENCES

Hughes, B., Ireland, R., West, B., Smith, N and Shepherd, D. I. (2019) *Project Management for IT-related Projects*, BCS.

NIST Cyber Security Framework: www.nist.gov/cyberframework

NPCC (2012) *Good Practice Guide for Digital Evidence*. National Police Chiefs' Council.

4 SECURITY LIFE CYCLES

In this chapter we discuss the life cycle of information that, in turn, drives the security issues that arise from the development, testing and implementation of new software. The ongoing life cycle of software is also a concern and is addressed here too.

You should gain an understanding of the importance and appropriateness of audit and review processes, of effective change control and of configuration management. You will learn about the differences in security between open source and propriety solutions, commercial off-the-shelf software and bespoke systems, and certified and non-certified systems. You will also learn about some of the techniques involved in reducing the security risks in the development of code.

THE INFORMATION LIFE CYCLE

LEARNING OUTCOMES

The intention of this section is to provide you with an understanding of the life cycle of data and information and its impact on system and software development.

The life cycle of information (Figure 4.1) must be managed in a way that supports the assurance or security of the information in the life cycle. The life cycle consists of three main stages, starting from the generation, creation or acquisition of the information, through to its ultimate archiving or disposal. Each must be considered appropriately, with the necessary controls and procedures put in place to support the confidentiality, integrity and availability of the information.

The first stage is how the information comes into the possession of the custodian. There are any number of ways this can happen, but basically, regardless of whether the information has been created by someone else and sent to the custodian by email, letter, telephone, data transfer or any other method, or it has been generated by the custodian themselves, the information comes into the organisation. This should result in some form of classification being attached to that information, either by the creator or someone acting in that role. The classification system will be discussed in more detail in Chapter 5, but in essence, it labels the information by importance and value so that the appropriate cost-effective security measures are put in place to look after it. Factors to consider at the point of acquisition include the planning of the system for the

Figure 4.1 The data and information life cycle

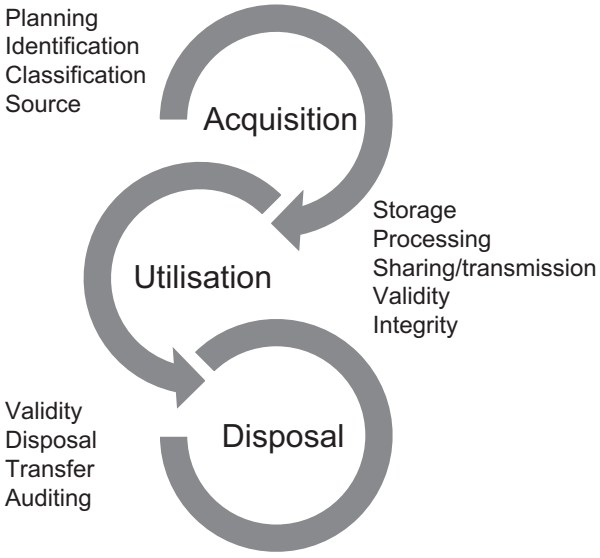

management of the information, the unique identification of types of information, the source and the classification of the information.

The second stage is the one that will often last the longest. The information will be utilised in some way. This could be to educate the organisation or to be published for others to understand or learn from it. The publication could be through paper format, as in a book, letter or other physical document, or it could be published, as is more common today, using electronic means to the public through the internet or internally on an intranet. This use of the information will perhaps happen just once or multiple times during its life. The factors to consider during this stage will include the secure storage, the processing, the sharing and transmission, the integrity and the validity of the information.

The third stage is to dispose of the information once it has served its intended and useful purpose. Disposal could mean deletion, or it could mean archiving out of the normal daily business so that it could be retrieved should the need arise. The factors to consider during this stage include validity dates, disposal methods, transfer methods for disposal and auditing the process.

Data are created at an ever-increasing rate. It has been estimated by Social Media Today (Ahmad 2018) that by 2020, 1.7MB of data will be created for every person on Earth each second.

One of the problems with information and data is that it is now so easy to store and share that it is rarely really deleted or destroyed. In the days of paper-based information

storage, filing cabinets became too full and so were 'weeded' to get rid of old information no longer of any use or value, often by shredding or burning. Today, we are much less inclined to do that and so the mountain of information simply grows. Storage has become cheaper as well, which simply encourages poor behaviours of proper information management.

In each of the three stages, it is vital that the most appropriate ways of looking after the information are considered and then implemented. This could be both software and hardware, in addition to physical and procedural controls, and will often be driven by the value of the information, the need for its security and its availability.

TESTING, AUDIT AND REVIEW

LEARNING OUTCOMES

The intention of this section is to provide you with an understanding of the issues surrounding security of the IT infrastructure and the content of the associated documentation.

Methods and strategies for security testing systems

Having built what is believed to be a secure system that meets the needs of the business, there is almost always value in proving that the end result is secure. This provides confidence to senior management in both the systems and the organisation's abilities to design and implement them effectively and securely.

A single test after completion is not sufficient, however, as threats and business requirements are constantly changing. Tests and reviews should be repeated at periodic intervals to look for any new issues of technology, threat or process that need to be addressed. Some of this requires expert testing by a professional penetration test team (as described in more detail in Chapter 6) and some of it requires review by a combination of business and security analysts. From time to time the advice of an independent external consultant can help to identify areas that may have been overlooked or about which the internal team have limited knowledge. This should form part of the ongoing risk management process, which exists to manage all risks, including these.

The continued review and analysis of vulnerabilities in systems is also a critical element of the ongoing drive to maintain the security of systems. Vulnerabilities are uncovered through a number of mechanisms, including the results of penetration tests, the analysis of viruses, system manufacturers' further development and specialist vulnerability analysts.

The choice of what to do about these vulnerabilities is up to the organisation to decide and a policy should be implemented to direct the actions of those staff members responsible for dealing with them. One of the major sources of a lack of security in systems is a poor patching policy and so this should form part of any organisation's security system. There is further discussion on a patching policy later in this chapter.

Correct reporting of testing and reviews

The test and review process requires accurate and comprehensive reporting if it is to serve any value. The report must be an open and honest 'warts and all' report that highlights any shortcomings in the security architecture. Any attempt to hide or downplay problems may lead to vulnerabilities being left in place that can be exploited successfully.

As ever, the report must contain detailed technical content, and an executive summary for those who don't have the time or knowledge to digest the entire report. This summary must contain the 'take-away' messages and important conclusions, along with a brief justification for further action and expenditure. Since it may contain details of vulnerabilities within the organisation, it may be necessary to give this report some level of protective marking to prevent unauthorised access.

Findings should be prioritised so that emphasis is given to the most serious shortcomings of the system. Sometimes these are categorised with a level of impact if exploited (e.g. high/medium/low) accompanied by a level of difficulty (easy/medium/hard) to exploit. Therefore, the overall rating will combine these two scores to focus on the worst, that is high impact that is easy to exploit.

Verifying the links between IT and clerical processes

The importance of aligning the information security architecture, policy and procedures with the needs of the business and its primary operational processes has already been emphasised in this book on more than one occasion. In the best traditions of the Deming 'Plan–Do–Check–Act' Cycle, this is where it is important to check that the task has been done properly and that the basis for the original design has not changed since the last review. This check will show if people are following the procedures and that those procedures are correct for the current circumstances. If it is found that the procedures are being widely disregarded or side-stepped, it is often a good indication that the design, the procedures or both are wrong, and changes should be considered.

Principles of monitoring system and network access or usage

As mentioned earlier in the book, there is a need to collect event log data from a whole range of systems, appliances and devices, and to monitor the traffic passing over the network and any external data links, such as the internet. There are commercial devices and software applications that can be used to perform this role, automatically processing anything up to hundreds of thousands of events per hour and capturing data for further use later. Many organisations will keep six months' worth of event log data for just such eventualities.

The data that have been collected can be analysed to detect unusual patterns of behaviour, malware and signatures of known attacks. They can also be reviewed

forensically to gather evidence of wrongdoing and abuse that can be used in an internal disciplinary case or provided to criminal justice organisations as part of their enquiries. The analysis needs to be performed by well-trained and skilled individuals. The training must not only be in the technical side of recognising unusual activity, but also in how to collect and preserve data in such a way that it is legally admissible in court. This kind of work can be outsourced to specialist third parties by smaller organisations without resources of their own.

Security operations centres (SOCs) often provide this function, and within their armoury they will use security information and event management (SIEM) tools to provide real-time analysis of the collected audit events. These tools combine detailed log management with a powerful analytics engine to enable them to detect patterns of behaviour that could not be detected by boundary or single end-user devices. They are able to compare results across multiple devices, enabling them to indicate an emerging problem better and more rapidly.

New systems are up and running, supported by a third party. GANT now needs to be able to audit not only the activities of the users, but also the services provided by the third party. This audit is for both the technical services and the quality of service against defined service level agreements (SLAs). It is also important to check that the work done does meet the requirements of the organisation, so some internal auditing must be done.

ACTIVITY 4.1

How will you define the content and standard of reporting that you require from your third-party suppliers?

ACTIVITY 4.2

How would you check for alignment between the actual business processes and the information security management system?

ACTIVITY 4.3

As part of an in-house exercise, your consultant has recommended that part of it should be the simulation of recovering data from a PC for use in an investigation. How would you plan to do this?

SYSTEMS DEVELOPMENT AND SUPPORT

LEARNING OUTCOMES

This section outlines the principles behind developing and supporting systems with an appropriate level of assurance.

Security requirement specification

The design of any application, system or network must meet the operational requirements of the users and also be aligned with the information security architecture of the organisation as a whole. The security requirement must be part of the overall statement of requirements document from which the design is generated. It is most important that the assurance requirements are captured at the start of any project in order to ensure that they are effective and that there is no adverse impact on the project or product from trying to reverse-engineer the security requirements later on. Adding them later will almost always add complexity and cost to any project.

Another issue can be the attempts by the project team to reduce the security requirements to save time and money on the project if there have been cost overruns or slippage of timescales. The security manager must be ready to defend their requirements but not be totally inflexible to urgent operational requirements. It is important to have the project and senior management sign-off acceptance of the increased risk that results from any changes.

Security should not be thought of as just defending against improper access and misuse, it also means:

- defensive coding to make sure that only valid and accurate data are processed by the system;
- proper functional testing of the system to ensure it behaves as expected and within the design criteria;
- methods to backup and secure data against loss or damage;
- adequate assurance of availability;
- compliance with any legal and regulatory requirements;
- security of communications;
- effective auditing of activity for, for example, regulatory and legal reasons.

Security involvement in system and product assessment

All new systems and products should have to go through some form of appropriate acceptance testing before being used in production. It does not matter if they were developed in-house or purchased, they should be assessed for acceptable and

appropriate levels of security. For example, a product bought from a reputable supplier should be afforded more trust than a piece of freeware written by someone you have never heard of, but should still be tested for appropriate operation.

Every product should be considered for its potential effects on confidentiality, integrity and availability, both directly and indirectly in conjunction with other assets, as part of the risk assessment process. Many 'best practice' organisations maintain a separate test environment that replicates the live systems to allow assessments to be conducted without risk of adverse impact. Another approach is to examine the source code (not always practical) by eye or with automated tools. Use of a malware scanner is always recommended for new code.

Security issues associated with commercial off-the-shelf products

The most obvious threat with commercial off-the-shelf (COTS) products is of rogue code hidden within an application that performs an activity against the best interests of the organisation. It could also be that there are 'bugs' that, while not intentionally malicious, introduce vulnerabilities that could result in a serious adverse impact. Mention has already been made of a separate test environment, and this is why it is important – to help find any such code by identifying its behaviour before it affects production assets. When a new product is installed, it is vital to ensure that all security updates have been applied to it.

Sometimes unscrupulous people will advertise cheap copies of applications because they have altered the code to include malware. The reduced price means it is more likely to be purchased and their malware installed. Security issues don't just mean checking for rogue code. It is also important to check that the product is a legal copy and not pirated. Make sure that the supplier is reputable, not some dubious market stall or website selling cheap copies. Failure to buy genuine copies can leave the organisation open to prosecution under the (UK) Copyright, Designs and Patents Act of 1988. That can mean financial penalties, impact on operations and loss of reputation. Make sure that all assets are purchased through trusted and authorised channels.

Links with all business areas

The development process is another area that benefits from contact with all the business functions that will be impacted by the new deliverable. All too often it happens that end users are given what the designers thought they wanted, but about which they had never bothered to ask. Consultation from day one has all sorts of benefits. The end users get the deliverable they need with a form of security built in that they can not only live with but also see a positive benefit from it being included. Project managers call this stakeholder engagement; it is a powerful tool that should be used by everybody.

A good security manager is in close contact with all their fellow managers throughout the company to ensure open communications and good feedback. It may be that the security team learns something new during the development process that could come in useful in the future. Keeping current with new system technologies and software tools is as important for security architects and managers as maintaining current knowledge of legislation and threats. Developing good relationships across the business with key stakeholders and earning their trust will greatly assist the credibility and trust of the security manager.

Separation of development and live systems

This has already been mentioned briefly. The main reason for keeping the live and development systems separate is to protect the live data from any unintended actions that might compromise them. Work to develop new systems and applications almost always contains mistakes in coding or design, and sometimes both. That is why functionality and acceptance testing is required. Any attempt to run unproven and incomplete code against a live database could have a major impact on the ability of the organisation to function. The international standards' ISO/IEC 27000 series contains a requirement to keep these two systems separate, as do some of the regulatory frameworks that exist around the world, especially in the banking and finance industries. Indeed, it is often considered best practice to have three separate systems – one each for development, test and live.

The last issue to consider is that the users may well need additional training before they can use the new systems properly, and the development or test systems can be a good place to allow them to make their mistakes, away from the live data. Accidental errors introduced by users is a regular source of issues and during training this can be a more frequent occurrence because they are less familiar with the system. Training on the test or development system removes the concerns about errors being introduced and also allows trainees to make mistakes in a safe environment. A script can be used to reset the data to known values for another attempt at the procedure, further testing or the next group of trainees.

Security of acceptance processes and authorisation for use

Once a deliverable, be it hardware, software or both, has completed development and is ready for deployment, it must be tested to make sure it does exactly what the requirements specify as documented in the functional test plan. If the product is an update of an existing product, there must also be regression testing to make sure that no unexpected changes to existing functionality have happened during the update process. This includes testing the security aspects of the product and also ensuring that the testing is conducted securely.

The deliverable(s) must not only work, but do so securely and not have any unintentional adverse impact upon the business processes or other business areas. A risk assessment should have been conducted as part of the design and development life cycle and the forecasts should be checked against actual outcomes of testing. Security testing needs to consider:

- effectiveness of defensive coding;
- protection against malware and code injection through interfaces;
- backup and recovery of data;
- access control;
- auditing and behavioural analysis;
- communications security;
- resilience.

Final acceptance testing should be performed by representatives from:

- the project team;
- end users of the deliverable;
- business management;
- the assurance team.

The final authorisation to go live should require sign-off from all these representatives before it can proceed.

The role of accreditation for new and modified systems

Some organisations have an accreditor, or the equivalent, who is responsible for ensuring that any changes or additions to their information systems and networks are of a required standard from a security standpoint specifically, although from all other aspects to some degree as well. This person has to approve the information security architecture, policy and procedures before the product(s) can be deployed and used. Normally this process is supported by formal documentation to standards defined in organisational, regulatory or legal documents.

Accreditation can apply particularly in the business world, especially in finance and aviation systems, where systems must be accredited by a regulatory body as being fit for purpose before they can be used. An alternative approach is where a new system needs to be capable of accreditation to a standard such as the ISO/IEC 27000 series. It may be that the organisation already has the accreditation, or is working towards it, and wants to ensure that the new system is capable of meeting the required standards for controls and countermeasures so that they will pass audit without remedial action.

The same principle applies to existing systems that are modified or updated. All changes should go through the same review process to make sure that the standards defined when the system was new are being maintained in the latest work. Many organisations also require periodic review and re-accreditation even if there does not appear to have been any change. Sometimes users will make changes in design or working practices without permission, or the environment changes (e.g. new threats and technology). Periodic review will help to identify these, and formal processes can then be used to take remedial action.

Many commercial organisations, and certainly smaller organisations, do not have a formal accreditation function. In this case there needs to be a formal sign-off for implementation of a new system so that any residual risks and vulnerabilities are documented, understood and agreed. Very often this is the business owner, but the sign-off must be aligned to the role or line of responsibility so that when the business owner moves on, the new owner inherits the decision. It is not the place of the security manager to perform this function.

Change control for software integrity

Any change to a software application, while designed to enhance its functionality, can introduce unintended problems. Every organisation should implement and enforce an

effective formal change control process to manage the risks to their information assets and reputation.

The start of the process is the submission of an outline of the proposed changes to a review board, who will assess the benefits against the risks and the work required to achieve the change. One of the members of this change board should be a representative from the assurance team, who will determine the risks and any changes to threats and vulnerabilities it may bring about. If the board approves the request, they may specify certain conditions and approaches to be used in order to manage the risks. Once the development work is complete, the new version must undergo regression and functionality testing, as described later in this book.

The process must apply not only to the software or hardware, but also to update documentation that describes its use, function and design. A copy of the new code and accompanying documentation must be lodged in a secure place for business continuity purposes. In some cases, this may be a formal escrow process, which protects the organisation against suppliers' businesses failing or suffering a business-impacting incident of their own.

Security issues arising from outsourcing software development

The practice of outsourcing has become more widespread. It often drives down costs, but it can also introduce new risks to the process. Some of these risks carry security implications, such as the introduction of malicious code, deliberately or accidentally (both have been known to happen), into the deliverable or customer systems during installation.

There is also the risk that there will be a loss of intellectual property or trade secrets through the information that has to be given to the third party, which may find its way into the possession of a competitor. A similar risk applies to any data sent to the third party. The laws and regulations on the protection of data (see Chapter 3) apply to anything sent to a third party as part of the development process.

A further concern is of a legal dispute developing between the customer and supplier. This risk can be managed by having appropriate terms and conditions in the contract, including agreed terms for dispute resolution, possibly by mediation, and understanding of the legal system or country in which disputes will be resolved. However, the business must realise that they are likely to be the biggest losers if a contractual dispute has to be sorted out through the courts. They may not have the system they need to operate effectively. It is always advisable to manage the risks by selecting a supplier that has reached level 5 on the Capability Maturity Model (CMM) for managed organisational processes.

The introduction of commercial offering through the 'cloud' has brought a significant increase in the number and diversity of issues for security to manage. This is discussed in more detail in Chapter 6. It is essential to ensure that there is a clear understanding of who owns the data on the platform, what format they will be returned in if the agreement is terminated, and that there are adequate controls to protect the confidentiality from other organisations hosted on the same platform. Again, with any third-party development it is good to have escrow protection to ensure access to the code if the organisation goes bankrupt or suffers severe business disruption. This is discussed later in this chapter. However, this is not easily implemented on cloud platforms.

Preventing covert channels, trojans and rogue code

Mention has already been made of the risks of unwanted code ending up in a product that is being developed or updated. Some kind of methodology should be used to inspect the code in order to identify any such malware. Code should also be developed to a clearly defined set of standards.

For short pieces of code, the code walk-through process is a simple yet effective way of checking for any extraneous lines of software, but for many modern products that are much larger, this is not practical. The testing process will require use of a system that is separated from the live network and replicates it as far as is possible, as previously described. Some kind of automated testing tool could be used, in conjunction with the examination of the resulting data, audit logs and outputs from network analysers, to look for unexpected and abnormal behaviour as part of the testing and acceptance stage of the development life cycle.

This work can be complex and lengthy if the application is large. It should be noted that this is one of the most complex and difficult tasks to perform fully and soon involves very complex mathematics if taken to its fullest extent. There are experts who understand the process, the various tools and their outputs and it may be advisable to involve one of these in this process if the risk of malware and its potential impact is deemed sufficiently high. There are also organisations who run Commercial Licensed Evaluation Facilities (CLEFs) throughout Europe with the skills and toolsets to do this kind of task on behalf of an organisation. They are generally measuring against CC, which is discussed elsewhere in this chapter.

Security patching

It is a fact of life that every software application and operating system contains bugs. The complexity and length of the code makes it impossible to test completely every single execution path through it. These bugs can have different impacts ranging from incorrect values being stored in a database to allowing unauthorised access to the system or network. One way or another, they will have some form of adverse impact on the confidentiality, integrity or availability of the information assets of the organisation.

When bugs are found, the supplier will normally issue a patch that can be installed in order to remove the vulnerability. These patches need to be tested and installed at the earliest opportunity. Hackers will also download the patches and attempt to reverse-engineer them in order to exploit the vulnerability if they can. The elapsed time from release of patch to the release of a usable exploit is now often measured in days.

Some people argue that patches should not be installed on certified products (see next section) as this changes the code away from the evaluated target. The official advice is that installing a patch to fix a known vulnerability is a much lower risk than that of accidentally introducing another vulnerability at the same time. Patches should always be applied. Having said that, patches should usually be tested, before they are rolled out, in an environment that is not connected to the live system to make sure they don't have an adverse impact on business functionality.

Some platforms are easier to patch than others. Older legacy systems may struggle with patching and become unstable. Unfortunately, these almost always tend to be mission critical! To protect them and the rest of the estate, additional protection measures need to be implemented so that the patches can be applied safely. Generally, these platforms should be managed out of the environment over as short a time frame as possible.

Use of certified products and systems

There are some industry sectors and circumstances under which it is advisable or even mandatory to use software products (e.g. firewalls), hardware devices (e.g. network switches) and operating systems that have been formally certified as providing a minimum standard of security, safety, reliability or a combination of these. Examples of this might include the nuclear industry, air-traffic control systems, finance, government and defence organisations. There may be industry or government requirements for the use of this kind of software or it may just be a requirement defined by the management of the organisation as part of a drive towards higher standards.

Probably the best-known system in use today is the CC assessment scheme that is recognised internationally, as discussed in Chapter 3. This provides a scale of product assurance, ranging from EAL 1 to 7; the higher the number, the greater the level of assurance. The concept is that an assured product can help to formally reduce risk in a quantifiable way when designing security architectures. The key issue to note is that each product will have a 'security target' or 'target of evaluation' of the features and functions that have been assessed. It is very important to make sure that the features you plan to use are included within that target, otherwise the certification is of no value.

Unfortunately, the testing process tends to be lengthy and expensive, so many manufacturers don't bother to have their products tested through this scheme. There are now other initiatives. The Commercial Product Assurance (CPA) system (run by the NCSC) is designed to provide a minimum standard of assurance, quickly and at minimum cost, that the product has the functionality it claims. This is sufficient for many organisations and helps to meet one of the controls in the ISO/IEC 27000 series for use of assured products. Another form of certification applies to encryption products, but that is discussed later in Chapter 9.

Use of escrow to reduce risks of loss of source code

If source code has been written or provided by a third-party organisation, the customer is dependent on that supplier for support, updates and changes to their software. There have been cases in the past where a supplier has gone out of business or been sold to a competitor and the end user has been forced to spend considerable sums of money to resolve the subsequent problems that pose a threat to their business, especially in getting support if something goes wrong.

One solution to this is escrow. The supplier and customer agree on a neutral third party (often a firm of lawyers or a bank) who will hold a copy of the source code and development materials. There is a legally binding agreement that specifies the circumstances under which the third party will release the material to the customer and ownership passes to them, along with all the relevant rights to use and develop the

application further as required. It will often include the circumstances described above and there may be other specific conditions that must be met.

The work that GANT is doing now requires some fairly complex bespoke code to be written, as there aren't any COTS packages dealing with amphibians. As the information assurance representative, you will need to be involved in the design and testing stages of the project to develop and implement the software.

ACTIVITY 4.4

You have been asked to suggest selection criteria for a third-party application developer. What would you put forward as mandatory and desirable factors in the selection process from a security perspective?

ACTIVITY 4.5

You hear two of the main users of the new application discussing a change they plan to make to the way they input data. What do you tell them they need to do before they go any further with their plans?

ACTIVITY 4.6

GANT has been offered a grant from local government towards the costs of the new application in return for helping to meet some of their environmental data reporting requirements to central government. There are some conditions attached, one of which is to use security approved products to protect the data. The management asks you to explain what this means.

SAMPLE QUESTIONS

1. **In the life cycle of information, which of the following is *NOT* one of the main stages?**
   a. Disposal.
   b. Creation.
   c. Acquisition.
   d. Utilisation.

2. **What technique should be used on a newly developed system just prior to its release into a live environment?**
    a. Penetration testing.
    b. Multi-factor authentication.
    c. Protective monitoring.
    d. PCI DSS.

3. **What is a COTS product?**
    a. Commercially operated temporary storage.
    b. Confidential organisational tested software.
    c. Certified off-the-shelf.
    d. Commercial off-the-shelf.

4. **The management of all alterations to an information system is best achieved by what service management process?**
    a. Configuration management.
    b. Requests for change.
    c. Change control board.
    d. All of the options above.

REFERENCE

Ahmad, I. (2018) 'How much data is generated every minute?', *Social Media Today*, 15 June. https://socialmediatoday.com/news/how-much-data-is-generated-every-minute-infographic-1/525692/

5 PROCEDURAL AND PEOPLE SECURITY CONTROLS

In this and subsequent chapters, ways of addressing the risks to information security are covered based on the three main categories of operational risk controls. This chapter discusses the controls involving procedures and people and how to manage them by use of the appropriate measures.

There are three main types of operational control:

- Physical – for example, locks on doors and secure cabinets.
- Product/technical – for example, passwords or encryption.
- Procedural – for example, checking references for job applicants.

At the time of writing, the latest version of the ISO/IEC 27001 Annex A (Reference control objectives and controls) contains 114 controls within 13 functional groups, and this does not cover everything. This clearly indicates that the subject of controls is an almost bottomless pit. All that is described here are the principles of the generic use of the major operational controls within information security dealing with procedures and people. More detailed information about specific controls is outside the scope of this publication. The key topic of training and awareness, which helps to reinforce the security derived from people, is also discussed.

GENERAL CONTROLS

LEARNING OUTCOMES

In this section you will be given an insight into the three main operational controls. Subsequent sections and chapters will look at each of these in much more detail.

There are three principal types of operational control that are available to the security manager and each will be covered in this and the following two chapters. Each has its place and role to play and, when used in conjunction with one another, can supplement and enhance the overall assurance of an organisation significantly and effectively. However, if they are used inappropriately or without due consideration, they can actually end up reducing overall security by providing a loophole, weak link or back door into a secure environment. The old adage that a chain is only as strong as its weakest link is still very true.

Physical security

Physical security relies on the presence or otherwise of physical limitations to the activities that a criminal or other unauthorised person might wish to carry out. The origins of such security are almost as old as man himself – ancient earthworks such as Maiden Castle in Dorset are evidence of an era when man had to protect himself from fellow man and indeed probably from the animal world too to some extent. The high ramparts of earth were later replaced by stone ramparts of castles, which were then supplemented by moats and the like.

Today, these are replaced by walls, fences and other obstacles that can prevent, or at least make it very difficult for, an intruder to gain access. Locks on doors of varying degrees of sophistication do the same job – provided the door itself is strong enough. There have been instances where very expensive digital or combination locks have been put on doors that themselves could be lifted off the hinges or simply forced open by battering the hinges off their mountings.

Technical security

Technical security is the general term used for any security measure that employs technology in some way. This is usually related to computers and software techniques that can be employed, but it could equally apply to technical locks using tokens or fingerprints, to hardware through the 'locking' or disabling of ports or to some other technological solution for a specific application.

Procedural security

Procedural security covers the rules, regulations and policies that an organisation puts in place to help reduce the risk of issues arising. They could, for example, include clauses in employment contracts that legally bind employees to obeying the security policy, the appropriate use policy or other necessary rules and regulations. While in itself this doesn't prevent problems happening, it can make them less attractive to staff if they know they could be disciplined or dismissed for contravening the rules.

Procedural measures also cover the correct vetting of staff before they are employed to ensure they don't have any convictions or other incidents in their background that might mean they are unsuitable for employment in a specific area. The induction training or probation period of employment might be another way of ensuring all staff are fully aware of their responsibilities as soon as they join an organisation. Setting standard ways of doing particular tasks associated with information might also be applicable. If, for example, there must always be two people present when the safe is opened or two people have to confirm the destruction of highly classified material, this too would be a procedural measure.

Overall, it is a layered approach to assurance, often called the onion model, that provides the best solution. This means that all three types of operational security control are used in varying degrees to protect the organisation's information assets. A set of controls, implemented effectively, might include:

- controls getting into the site (procedural and physical);
- further checks for specific high-risk buildings (procedural, physical and technical);

- specific logons to computer systems containing or processing classified data combined with limitations on what users can do on the systems to which they have access (technical and procedural);
- a set of well-drafted and effectively policed policies of which staff are well aware and have signed contracts to that effect (procedural).

This combination will provide a high level of assurance and should prevent most incidents happening.

PEOPLE SECURITY

LEARNING OUTCOMES

In this section you will be given information about the sort of procedural and people controls that can be used and their appropriate application.

Security culture within organisations

The most sophisticated information assurance system on the planet is worthless if the people, whose data it is designed to protect, are not security conscious. They need to be made aware of the threats and dangers, how relevant they are to them and their data, and how to use the systems to make sure that the information assets are protected. This is very, very important and the lead must come from the top of the organisation. As mentioned in Chapter 1, there should be an information security policy document, signed by the chief executive or equivalent, which says words to the effect of:

> We take information security and assurance very seriously and it is a high priority to us. It is the responsibility of everyone within the organisation to be security conscious and to abide by our information security policy and procedures when dealing with colleagues, suppliers and customers. If you are uncertain as to the correct course of action, or are suspicious about a set of circumstances, your duty is to consult the information security manager for advice.

A positive security culture needs to become part of the culture of the organisation. The organisation's security culture should align with the business' needs and strategic direction and with the security strategy. It will only happen if those in senior management lead by example. Do not expect people to obey the organisation's rules if senior management in particular do not stick to the same rules. A large percentage of assurance incidents are caused by the inappropriate actions of an organisation's staff, with far more problems being caused by accident rather than maliciously. A positive security culture will reduce the number of incidents. Without it, people will either not use the countermeasures or find ways round them. Be proactive in promulgating this culture throughout the organisation and keeping the topic fresh in people's minds.

Security awareness

A large part of creating and maintaining the culture is by developing an ongoing security awareness programme. People need to be taken on a journey to develop a positive attitude. It is no use telling people that 'there is a risk to our organisation and its assets'. That means nothing to the majority of the staff; it has to be made interesting and something to which they can relate. There should be a programme of training for all employees and it should be included as part of their induction training programme too. Make sure that everyone is aware of the particular risks to the organisation and why they are risks. This may include training in specific laws and industry regulations that apply specifically to the organisation. Give them examples of organisations that have suffered as a result of security incidents. If it can be shown that the organisation will lose money or be prosecuted, that is likely to mean something to them – it could affect their job security. Use humour and stories and other strategies to make the messages real and relevant.

It is also important to make people aware that this isn't just about confidentiality. Introduce them to integrity and availability too. In today's fast-moving business world, it is often these two aspects where greater problems can occur and have a much more significant effect. Another factor to consider is that of non-repudiation – evidence that the information provided has actually come from the source indicated and therefore can be believed. This issue is now causing much more significant fraudulent activity than has been the case. As an example, invoices that seem to be genuine but that are actually copies created and sent by a criminal are a serious problem being seen increasingly now in the UK and elsewhere.

If possible, keep a record of assurance awareness training – who, when and what – it has several uses. First, it allows the identification of when refresher training is required, for example, due to the passage of time or changes to risks or the law. Second, it could prove very useful in the event of legal action by helping to prove due diligence.

Contracts of employment

A contract of employment is a very important document because it has legal standing. It will define the terms and conditions of employment, including the responsibilities of the employee towards the organisation, and anyone else with whom they interact on behalf of the organisation. The document also defines the obligations that the organisation has towards the employee. It will include all the standard information about pay, leave, illness, training, health cover and so on.

In many cases the document, once signed, will not see the light of day again, but in some unfortunate circumstances, such as disciplinary cases, it will need to be referred to and quoted. The contents may end up being cited at an employment tribunal or produced as evidence in a court of law. This is why it is important to get the contract right, including the parts that define:

- acceptable standards of behaviour and conduct;
- ownership of intellectual property;
- acceptable use of company assets;

- grounds for disciplinary proceedings and the disciplinary process;
- adherence to all applicable laws and regulations;
- duty of care to the organisation and other staff;
- non-disclosure/confidentiality of information;
- privacy responsibilities;
- responsibilities concerning access to and use of PII.

Service contracts and security undertakings

The nature of the modern world means that almost every organisation now has a service contract of some sort with a supplier and customer defined within it. Both parties give undertakings about the meanings of certain standards and naming conventions, such as protective markings, and the identification of policies and procedures to be observed. This document can be amended by mutual consent without having to resort to contract renegotiation.

The security details for such contracts are normally contained in a security aspects letter, which is often included as an appendix to the contract issued by the client to the service provider. This forms a binding part of the contract and is to be observed by all parties to the contract. It will also often contain the possible consequences of not complying with the assurance requirements, set out with any consequential damages specified.

Codes of conduct

The most obvious elements of a code of conduct in the context of this book are the obligations placed upon employees regarding information assurance – confidentiality, integrity and availability, which are described in Chapter 1.

A code of conduct can cover a lot more than that. It is also a means of expressing the ethics and standards of the organisation. It will contain examples of the kind of behaviour expected of employees in their dealings with each other and other people, be they customers, suppliers or anyone else.

This can include rules on accepting (or not) hospitality, and guidance on accepting and declaring the receipt of gifts and inducements from third parties. It is not unusual for the rules to say something along the lines of: 'Gifts should be accepted when it would cause offence to refuse, but they must be declared to the organisation at the first opportunity.' Sometimes staff are given the option to buy the gift at a fair price, it might be used by the company or it may become a prize in an annual raffle or prize-giving.

Additionally, the code of conduct may state that staff must not offer gifts, inducements or unreasonable levels of hospitality to others, or that alcohol is not allowed on company premises without approval of senior management, and usually only for entertainment of clients.

The ethos of customer relations can also be included in this heading; for example, 'always be helpful, polite and approachable to everyone, employees or otherwise'. Many supermarket chains have a programme along the lines of 'Every Customer Offered Help', which is why they will ask if one needs help with packing nothing more than two packets of chewing gum and a new toothbrush.

While some of these areas may seem distant from information assurance, it must be remembered that one of the most important areas of security is that of social engineering. Members of staff may be provided with extravagant gifts or put into embarrassing situations by others who want some information or access to other assets of the organisation; this is a type of social engineering. There are many other situations where those intending to damage an organisation or its assets use the social behaviour of staff to achieve their aims.

The passing of amusing emails is another type of social engineering. Providing an email with an attachment of an amusing photograph that is readily passed around the staff may seem innocuous enough. If that email's attachment contains malicious code, it is a very efficient way of compromising the whole IT system.

This sort of attack is not limited simply to IT systems, though. In the more complex, perhaps higher-value world, gaining access to keys to buildings, access codes or passwords for security systems is also often achieved by variations of social engineering. It is in this light, therefore, that banning gifts, for example, helps to reduce the threat to some degree, although it would be very naive to think this would stop social engineering in its entirety.

Acceptable use policies

The acceptable use policy, sometimes known as an end-user code of practice, is the document that defines the standards for the use of organisational information and communications systems by employees. This serves as an adjunct to the contract of employment to protect both the organisation and the individual from the actions of others. In law, an organisation can be held accountable for the actions of employees under what is known as 'vicarious liability'. The only defence an organisation has is to show they have carried out due diligence in telling its staff that they are not to break the law or any relevant regulations. This document can also help to protect staff from harassment or malpractice by employers and other employees.

In this document the management must make clear the level of infringement (e.g. misconduct, gross misconduct, etc.) for each offence and the disciplinary steps that will be taken against those who are considered to have broken the rules. It is advisable to consult staff before introducing a document of this kind, to ensure that they understand the reasons why it has to exist and support it as being fair and reasonable. It is common practice to include a brief on the implications of this document in an induction training course, and this is an effective way of dealing with a lot of related issues while also ensuring full understanding by the staff members. This topic is also discussed within 'Organisational policy, standards and procedures' in Chapter 3.

Segregation of duties and avoiding dependence

The segregation of duties is the concept that one person may not perform the duties for more than one role where there could be a conflict of interests. The requirement for segregation of duties has two functions:

1. To limit the scope that any one individual has to attack and compromise the information security of the organisation. A commonly used term for this kind of activity is 'system misuse'.

   If one individual has all the passwords, access rights and privileges for the entire organisation, they have the ability to systematically alter, extract or destroy any or all data within the control of the organisation with little or no risk of discovery. Many risk management systems consider one person having full access to be a serious risk. There should be separation between the roles of system administrator, system user and system auditor in order to manage the risk of collusion and fraud. Some legal and regulatory bodies require high levels of segregation and, in addition, mandatory third-party audits. Examples of this include the (UK) authorities regulating financial services and the (US) Sarbanes–Oxley legislation.

2. To limit the dependence that an organisation has upon any one individual.

   This is about the obvious fact that if the organisation relies on the knowledge or skills of one person, it is vulnerable. If that person falls under the proverbial bus or resigns, taking their knowledge with them, the organisation has a real problem. This is a very difficult problem if staff numbers are low, especially in IT. Good documentation can help, providing that it is kept up to date, as can cross-training and succession planning.

Obligations on third-party suppliers of goods and services

It is sometimes overlooked that the obligations upon an organisation need to be taken into account when dealing with other businesses. There may be times when contracts for goods, services or both are outsourced. The contracts to cover this will need to include legally binding clauses that cover the information assurance aspects of the data and services concerned. The information owner has a legally binding duty of care to ensure that the external body is competent to process the data securely and will observe the same high standards as the organisation on behalf of which it is performing the work.

This stipulation (and the enforcement) of obligations has two benefits for the data owner:

1. It manages the risk of loss of goodwill, and punitive or corrective action as a result of an information security breach.
2. It manages the risk of leaving the organisation exposed to the impact of a business continuity risk event happening at a supplier. Setting obligations upon suppliers to use good practice to manage their own risks improves the level of confidence in ones' own business processes not being affected, or at least being minimised.

It is important to note that the obligation should include the right of the organisation to audit suppliers (either directly or via a specialist third-party auditor) to ensure that they are complying with the requirements. This may be in the form of planned or no-notice inspections.

The profile that GANT has among the public has started to rise, thanks to an inspired publicity campaign. The workload is such that volunteers and one or two full-time specialists are now working for the organisation from an office. There are the beginnings of what may well become a full-blown IT infrastructure as time passes. Owing to the limited resources, some work is to be outsourced to service providers, who will be much cheaper than staff doing the work in-house for the time being.

ACTIVITY 5.1

Now that the number of officers and volunteers working for GANT has started to rise, what do you think needs to be done in order to promote the need for assurance awareness among staff and service providers to the organisation?

ACTIVITY 5.2

How would you advise GANT to manage the risks that the use of third-party suppliers can bring?

USER ACCESS CONTROLS

LEARNING OUTCOMES

The measures taken to provide information assurance within an organisation are often referred to as controls. This section describes the basic building blocks upon which many other controls are predicated – the access controls. Without these, most of the other possible strategies would count for nothing. This section provides an overview that is essential knowledge for every information assurance practitioner.

The intention of this section is to provide you with the basic knowledge to understand how people and organisations should be managed within a culture of assurance.

Authentication and authorisation mechanisms

The process of authentication and authorisation is often referred to by the acronym 'ID&A', which stands for 'identification and authentication'. First in this process, the user has to tell the system who they claim to be (identification) by entering a unique username. The system will then challenge them to prove that identity by providing some form of knowledge that can only be known, or possessed, by the individual that they

have claimed to be (authentication). The system compares the data it receives against a known value it holds and, if they match, it provides access to the system.

Traditionally, the second (authentication) value has been a password, which the user is supposed to remember and not tell anyone else. The reality is that, although users are becoming more sophisticated in protecting their passwords, people still do write them down or they choose something that is easy to guess, like a date of birth, name of spouse/child/pet, car registration, sports team and so on. Perhaps the most significant issue with passwords is the tendency to use the same one across multiple systems such that if the password is discovered for one system, all the others then also become vulnerable. Many information assurance professionals believe that too much faith is placed in the ability of passwords to control effectively the access to systems.

The traditional defence against password guessing has been to allow the user three tries and to lock them out if they fail to enter the correct password. Unfortunately, this provides a form of DoS attack – allowing an attacker to disrupt the availability of a system by deliberately locking out users. In addition, there are well-known techniques to capture passwords travelling across a network or to grab copies of the file on the authentication server that holds all the values for comparison. Copies of programs that will attack and 'crack' these passwords are easy to find on the internet, meaning that this form of attack is relatively easy to carry out. People are also very easily fooled into giving out their passwords through social engineering attacks, mentioned earlier, for example where they believe they are talking on the telephone to someone in IT support. Passwords are not strong security; they are actually fairly weak. It is very unwise in general to put a lot of faith in passwords as a security measure.

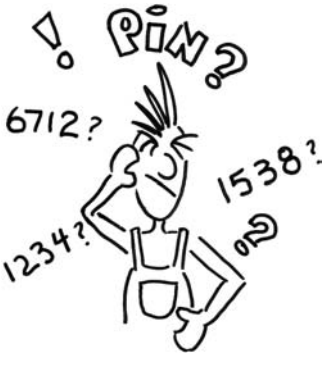

In order to provide a more effective level of assurance, many organisations are using two-factor authentication (2FA). This is where the user has to enter a password and something else as well before the system accepts their claimed identity. Quite often this involves use of a token, such as the RSA™ SecurID™ device. The traditional type is the size of a keyfob and has a liquid crystal display. This displays a six-digit number that changes every 60 seconds. The values displayed are based on an algorithm and secret key value that is known only to the organisation that owns the system. The sequence of numbers displayed is not predictable and it has resisted attempts to break it for many years. The sequence for each user is different, so they cannot be interchanged with other users. The user is asked to enter a secret personal identification number (PIN) supposedly known only to them and then the value showing on the token. This is compared to the value calculated by the authentication server. If the values match, access is granted.

The PIN provides protection against the token being stolen, providing time for the loss of the token to be detected and that particular unit to be disabled. Some banks and

online trading organisations have issued their customers with such tokens. There is no doubt that they provide a greater degree of security. The downside is the cost of buying and managing the system and tokens, which normally have to be replaced periodically, usually because of the battery life limitation. Other systems use one-time passwords (OTPs) sent to mobile phones, which are obviously cheaper to support. Methods are constantly evolving.

Another approach increasingly gaining acceptance is biometrics. This is the use of a characteristic of an individual that is unique to that person, either anatomical (e.g. fingerprints or facial recognition), behavioural (e.g. signature) or a combination of both aspects (e.g. voice). Although the concept has been known for a long time, its introduction has been delayed due partly to the technical challenges in producing a system that is appropriately reliable, and partly to issues of public perception and acceptability. Many of the early systems were regarded as being too intrusive or unreliable, but technology is rapidly evolving and improving and hence the increasing take-up. Biometrics have distinct advantages over many other forms of identification and authentication methods:

- They are free with every user and very difficult to steal or lose, they even self-repair, although in certain trades and professions (for example the manual trades), fingerprints can wear away or be damaged for considerable lengths of time.
- The person to be identified can be required to be physically present at the point of identification.
- Identification based on biometric techniques reduces the need to remember a password.
- You can't write down a biometric on a piece of paper for someone else to find.

They do, however, require the use of sensors that can reliably read a biometric and detect attempts to defeat the system. These require capital outlay and integration into the security management system.

Effective use of controls

Now that the control of access through the perimeter of information systems has been discussed, it is necessary to look at how to limit the access a user has once they are granted entry. Users should only be granted the minimum level of privilege to perform the role assigned to them. For example, a person working in a warehouse may need access to the stock records, but they should not have access to the detailed financial or personnel records of an organisation. They simply do not need to know. The concept of need to know is very important. It is estimated that over 50 per cent of attacks originate from within the organisation; this is another line of defence against that type of attack to help in securing information. There is also the issue of privacy of data and legislation such as the Data Protection Act to be taken into consideration.

In order to make this limitation possible, it is necessary to assign attributes to users and to data, describing their profile to the remainder of the systems. This information can then be used to control access to data and systems, providing another level of protection for data and the users.

The first concept is that of user groups to grant role-based access. Those users who perform a similar function are grouped together (e.g. 'Accounts', 'Sales', 'IT', etc.). This can be used to control access to applications, or functionality within large integrated applications such as enterprise resource planning (ERP) systems. An example of this could be that access to the payroll system is only permitted for finance and human resources group users. A user can be denoted as being a member of more than one group if their role(s) requires it. This type of designation may be based on some other attribute, such as geographical location. All users in country A will have access to the data relevant to their own country, but no access to the data of country B, and vice versa.

One other kind of user account that must be mentioned is that of the system, or application, administrator. This role has access not just to the data (quite often all of them), but also to the software and operating system itself. The 'sysadmins' can add and delete users, groups or levels of privilege, rebuild the system, erase data, grant or deny access to applications, change passwords and even alter or destroy event logging or auditing data. These accounts have great power and wide-ranging capabilities and their use must be limited, tightly controlled and safeguarded. Their potential to disrupt operations, accidental or otherwise, is enormous and these accounts must be locked down as far as is practical.

This degree of protection can be extended to the data as well. The standard approach is that there are three levels of privilege:

- each file has a designated owner, who has full control of the file;
- other members of the same user group as the owner who may have some degree of access, as described below;
- the 'rest of the world', that is other users in other groups, who may also have some limited access, as described below.

There is another level of granularity that can be provided, which is to say whether the user can do the following:

- Read – the user can see the contents of the file or database, but not change them.
- Write – the user can change the contents of the file or database. This includes deletion privilege, since the user could overwrite the data to destroy it.
- Execute – the user can run this if it is an executable or a command script file. This implies full ownership, including write privilege.
- A final option is not to grant any level of permission to one or more users, so that they cannot even open the file.

The actual names for these functions vary from one operating system to another, but the concept is the same.

The attributes on a file may then be as follows:

- User/Owner – execute permission, allowing the file owner to read, write and run the script or application, if that is what it is.

- Group – write permission, allowing other members of the same user group to update the file.
- Other – no permission, so that the data or application cannot be read, amended or run by anyone who is not a member of the same group as the owner.

Administration of controls

The administration of access controls is another of the important jigsaw pieces that make up the whole picture. Their use at the appropriate time is essential. The levels of privilege that each user or administrator has should be reviewed and updated regularly. This takes into account the fact that people change role (e.g. get promoted or transfer to another team). New privileges that are needed should be granted, but ones no longer required (i.e. old group memberships) should be removed.

The role of the administrator should include the following:

- Enrolling new users in the system after appropriate validation of identity. Without access they cannot do any work.
- Removing user access rights when users leave the organisation to prevent any further access.
- Modifying user access rights if users change role within the organisation. They may not be able to perform their new role without a change in rights and they may no longer need access to data they previously accessed.

People who leave the organisation should have their accounts deleted and all rights removed on the day they leave. Ownership of any data assets should be transferred to another user.

Role-based access is an issue that often gets overlooked. The roles of system or database administrator (very high system privilege required) and those of standard users must not be combined into a single account. The lack of high-level privileges is a good security countermeasure. The separation of these roles provides protection against accidental and deliberate abuse of the system. It also ensures a higher-quality audit trail. It is a legal requirement in some industry sectors to separate out not just the 'sysadmin' functions, but whole departments, especially in finance and banking.

If the size of the organisation is such that the same individual is required to perform more than one duty, often within the IT department, then the segregation of duties must still be enforced. This can be achieved through the requirement to log in using a different account with the appropriate privileges when performing the different duties. While this may seem bureaucratic, it provides an effective audit trail and also helps to reinforce the serious nature of information security. There must always be someone else (another individual either internal or, as a last resort, external) who is used to audit the work of the main 'sysadmin' function.

There can be occasions when access is needed to other areas of the system in order to meet an unusual or temporary operational requirement (e.g. to cover for a colleague

on leave). There should be a process for users to apply for additional privileges to be granted temporarily so they can do the work. The approval should specify a start and end date for these, and those dates should be rigidly enforced.

These procedures must also be applied to any temporary staff who work for the organisation. A regular check should be made for obsolete roles, user accounts and privileges to guard against any lapses in the process.

Access points

An access point is any location from which the internal systems of the organisation can be accessed. This can be via one of three main types:

1. Direct connection from a hardwired terminal.
2. Wireless network access within the perimeter.
3. Remote access over a third-party network, such as via a broadband link at home or from a hotel, coffee shop, train or pretty much anywhere.

There are two main security concerns with giving access:

1. To ensure that the user completes the ID&A process successfully.
2. To protect the data in transit being used to complete the authentication process and then the session itself.

If a user is connecting from a hardwired terminal, then they are probably located within the premises of the organisation. This provides a degree of physical security to manage the risk that the person sitting at the terminal is not an authorised user. The main risk is that someone is watching the user, either visually or electronically with a 'network sniffer' or key-logger to capture the ID&A data. It is best to make sure that the login session is encrypted as it passes across the network in order to stop the data from being reused to impersonate the user later.

Remote access presents additional challenges of physical and network security. It is quite likely that the network connection used is not controlled by the same organisation. This means it is even more important to protect not just the ID&A traffic, but all the data. This is usually achieved by setting higher requirements for ID&A, such as two-factor authentication using tokens described earlier, and protecting the data connection using encryption, such as by using a virtual private network (VPN) tunnel or secure sockets layer (SSL) via the secure 'https' protocol.

A considerable challenge has been created by the rise in popularity of wireless networks. It is almost considered a right by many to have access to a wireless network wherever they go. They have a lot of advantages, primarily large savings in cabling costs, but they also have disadvantages, mainly to do with security. The radio waves do not stop at the physical perimeter of the organisation or house. They travel outside and allow unauthorised access to attackers if not properly configured and located in the network architecture. Readily available aerials mean that a determined attacker can connect from a range of several miles rather than sitting outside in a car, which would be fairly easy to spot. If they are in a nearby building (e.g. at a neighbour's house or another hotel

room) they can be almost impossible to spot. Wireless access points must be strictly controlled in terms of installation, configuration and their physical links to the network. Many organisations connect them into their 'demilitarised zone' (DMZ) on the network, so that users have to authenticate to a higher standard and their traffic is screened by a firewall before having access to the internal organisational network.

Another significant risk is that wireless users in an organisation may accidentally connect to an unsecured wireless network belonging to someone else in a nearby location, exposing the organisation's data to them or possibly even leading to accusations of hacking.

Protection of data

The protection of data by means of an information classification system (also referred to by some organisations as protective marking) is one of the oldest, yet most effective, security countermeasures yet invented. The point to remember is that the definition of the word data has changed, and continues to change.

Originally, the data were only in the form of paper files and knowledge inside the heads of the employees. Now it includes all forms of media, whether in storage or in transit from one place to another, including:

- magnetic – such as external disks, USB sticks, magnetic tape, tablets, mobile phones, digital cameras;
- optical media – such as CD, DVD and even still microfiche;
- paper – such as handwritten notes, printed files, punched tape, blueprints and plans;
- data on Wi-Fi and radio frequency networks;
- physical – some devices may have a protective marking because of their design or content;
- email, texts and the myriad of social media platforms such as Facebook, LinkedIn, Instagram and Twitter.

All of these are considered to be information assets and thought must be given to their value to the organisation. This is not restricted to their immediate commercial value, for example the design for a new product, but also the impact they could have on the organisation or other people if their contents were to become known to a competitor, foreign country or the public. The impacts can range from 'negligible' through to a 'grave impact on national security' for a government, or 'major loss of goodwill with the public' for a commercial organisation.

All such assets need to be identified and valued against an agreed impact system; it is another form of risk assessment. Government and commercial organisations tend to use different classification terms. The UK Government at the time of writing has three levels of classification:

- top secret;
- secret;
- official.

Official is the lowest level and covers the majority of information that is created or processed by the public sector and covers business operations and services, some of which could have damaging consequences if lost, stolen or published in the media, but not to heightened levels. Other governments, such as the USA, may include classifications such as:

- confidential;
- restricted;
- protected;
- unclassified.

Details of what these levels mean and how information should be handled can be found on the various government websites.

Commercial organisations may have a system such as:

- highly confidential;
- confidential;
- internal only;
- public or open.

Whatever the system used, once values have been assigned, a set of rules for handling and distribution of each classification of information must be drawn up to define their use. The most fundamental guideline is universally referred to as the need-to-know principle – information should not be made available to people who do not need to know it. The fewer people that are aware of the knowledge, the easier it is to protect, yet that also presents a challenge in that enough people need to know to make best use of the information. There are even regulations in place to control the distribution of knowledge in some organisations. For example, in finance the concept of the 'Chinese' or ethical wall is used to guard against conflicts of commercial interest and insider dealing of shares. The trick is understanding where the ideal balance point lies for each piece or type of data.

Each level of classification requires protection appropriate to the value it has or, more accurately, the impact its inappropriate release or knowledge would have on the organisation or individual affected. The controls will not be just physical, but procedural and people-related too. Data of the 'internal only' variety are probably sufficiently well protected by the normal ID&A mechanism for the system, the standard business rules and the locks on the doors of the building. Data that are 'top secret' is often required to be kept in very strong safes inside heavily guarded buildings, handled in strictly defined ways, and can only be accessed by people who have been through an extensive security screening process.

In addition to these protective markings, data can also be given 'caveats'. These are additional markings that define a finer layer of protection and discretion. Some examples are:

- 'Human resources only' – personnel files containing sensitive personal data.
- 'Board member eyes only' – not to be shown to anyone who is not a board member.
- 'Commercial in confidence' – not to be shown to any competitor organisations.
- 'Confidential until …' – for information such as a product launch or a campaign that needs to be kept confidential until a certain point.
- 'Intellectual property' – subject to non-disclosure rules and possibly pending a patent application.

The key point overall with using any system of this type is that once a piece of information has been given a security classification, it automatically imposes certain constraints on the methods that can be used to process, store, transmit, dispose of or otherwise deal with it. These conditions are imposed on anyone who may come into contact with that information.

The growing numbers of members and their records mean that it is high time for GANT to register with the Information Commissioner. This means complying with the requirements of the UK Data Protection Act for the proper protection and processing of personal data. In order to achieve this, it has been recognised that GANT must have some policies in place to control user access to their IT systems and data. This requires a survey of the IT equipment in use and the methods used to access it, together with the places from which the internal systems can be accessed.

ACTIVITY 5.3

What do you think needs to be included in the policy documents for access control?

ACTIVITY 5.4

You have been asked to conduct a survey of the IT systems in use and the means used to access them. How would you set about conducting a survey and what would you look for?

TRAINING AND AWARENESS

Protecting the organisation's information is not usually at the top of most managers' priorities. They are more likely to be concentrating on immediate pressures like hitting sales targets or meeting deadlines. They may not have considered how reliant they are on their information systems to help them achieve these goals or whether these systems are vulnerable. Too often, the need to secure data properly is only brought to one's attention after it has been lost or becomes corrupted. Therefore, ensuring that users understand their assurance responsibilities, and are aware of the risks to their information systems, is a key security control.

LEARNING OUTCOMES

The intention of this section is to provide you with the basic knowledge needed to develop assurance training initiatives within an enterprise. Following study in this area, you should be able not only to explain and justify the main concepts but also to develop a high-level approach and draft documents to meet the general requirements in the areas of:

- purpose and role of assurance training;
- approaches to training and promoting awareness;
- available training materials;
- sources of information for training material.

Purpose and role of security training

Organisations need both their staff and any third parties accessing their information to comply with their information assurance policies and procedures in order to reduce the likelihood of assurance issues. Providing appropriate security training will help individuals to understand their assurance responsibilities, how the enterprise's information assets can be put at risk and how this can be avoided. Enterprises that do not implement awareness and training initiatives are more likely to experience security-related issues. Security training is, in relative terms, a low-cost assurance control that can create a positive and lasting change in users' behaviour.

By understanding the risks, the users are more likely to remember what they need to do to protect the organisation's information and the systems containing it. For example, most people would realise that leaving their wallet on display in an unlocked car could attract an opportunist thief. They would certainly have a personal appreciation of the loss and inconvenience caused by its theft. However, they might not equate the loss of information assets in quite the same way. They also might not be aware that disclosure of sensitive information could lead to a breach of current data protection legislation or that not following a set procedure, such as a data backup, could result in a severe financial loss to the organisation.

Anyone with access to the enterprise's information systems should receive some form of information security education and training. The level of training that they may need can vary with their role, but it should be sufficient to ensure that they can carry out essential assurance procedures and have sufficient understanding of the correct use of their information systems. It should always include awareness of the acceptable use policy, no matter who they are.

The key messages, tone and approach of a security training or awareness programme must be relevant to the intended audience and consistent with the values and goals of the enterprise. Messages may contain common themes, but the language and delivery should be tailored to suit the audience. Therefore, when developing security training, thought should be given to the messages to be conveyed and what needs to be achieved. You should consider the following questions:

- What does this group of people need to know?
- Why do they need to know it?
- What is their current understanding?
- What is an appropriate method to gain and maintain their attention?
- What should they think and do after the messages have been delivered?

Do use language that they will understand and avoid jargon with which they are not familiar. Examples and case studies will need to convey 'real-life scenarios' and be appropriate to them. Security incidents that may have occurred previously within the enterprise or within other similar organisations are always useful to get the message home.

Security awareness and training should be seen as a mandatory, continuous process rather than as a once-only exercise. Its overall objective is to reduce information assurance risk by developing a positive security culture. This is achieved by increasing the level of understanding about information assurance and explaining to users what is expected of them to protect the organisation's information assets.

Approaches to training and promoting awareness

There are two broad approaches to improving levels of knowledge: first, through specific information security training; and, second, through raising awareness of information security. Training tends to be focused and addresses specific issues. Its primary aim is to achieve within the user a certain level of competence in a given area. Awareness is more general and aims to create a change in user behaviour and influence the perception of risk.

Individual campaigns should be developed to target particular areas for improvement, cater for the various types of audiences, utilise different mechanisms to engage the intended audience, or to cover some specific security matters. Effective campaigns need to be meaningful to their audience to result in a long-lasting change in user behaviour, and should concentrate on what an individual can do to improve security.

As with other security activities, it is important to gain sponsorship from the senior management of the enterprise. If senior management are seen to value and support positive security behaviour, then line management and general staff members are more likely to adopt similar behaviour themselves. Without senior sponsorship, line managers may be reluctant to release staff members to take part in the campaign or communicate the need for security to them. In turn, staff members may not take the campaign seriously as they may not have been given sufficient time and support to be involved, nor may they appreciate the importance of security within their roles as this may not have been communicated to them.

An awareness campaign or security training programme should be developed as a formal project with agreed objectives so that it can be delivered efficiently and measured for success. It is important to concentrate on the security issues that are relevant to the enterprise and not just around what the hot topics are within the industry.

The issues that will be addressed by the project and the training messages should be identified. Each training message should explain the security issue and what can be done to address it. If users do not understand what the problem is and what is expected of them to help address it, then they are less likely to adopt the desired behaviour. Constant repetition of the same messages or information presented in a dull manner will create user disinterest. Some issues, like password sharing, are perennial problems, but it is important to try and deliver them each time in a fresh way.

Getting the audience's attention is crucial. Using techniques such as humour, storytelling or relating messages to how they can protect their home devices or their children online, all help to gain their attention and become involved with your campaign. There are new mechanisms now easily available to engage the younger audience in particular, including games, scenario-based exercises and escape rooms. Some individuals are competitive and an element of gamification and the possibility of winning a prize can work effectively, so think creatively how best to engage each audience.

The type of approach adopted will be constrained by a number of factors, including the size of the enterprise, its culture, available funds and either the scale or the scope of the campaign. If the organisation has a press office, communications specialist or perhaps a training department, then their involvement is always beneficial as they should be able to provide guidance and advice on a suitable approach.

Timing is everything. To get the best attention from the target audience, schedule campaigns to fit in with working schedules and enterprise priorities. Avoid busy times such as year-end accounting, month ends, peak sales periods or peak holiday periods.

Available training materials

There is a wide variety of methods and materials that can be used to support awareness and training campaigns. Choice is usually constrained by budgets and the size and culture of the organisation.

Face-to-face sessions are effective as the participant is able to interact directly with the trainer, but they can be resource intensive, especially if many people need to be trained, and can become routine and boring. External training courses can be used to cover specialist topics. If there are sufficient numbers, it is sometimes more cost-effective to get an external trainer to carry out the training course on the premises. As an alternative, courses or workshops could be developed in-house. Messages can then be tailored to each audience and the sessions repeated many times over at little or no extra cost. Face-to-face training will take staff members away from their normal activities for periods of time, so this approach could meet with resistance from line managers unless support is gained from senior management.

Training videos can deliver a message that is consistent throughout the organisation and have the benefit of being easily transportable. The formats can be adapted to enable people to watch them on their phones or tablets when convenient to them, or when travelling or at home. Long gone are the days of everyone crowding round a television to see the company training video! This could work out as a reasonably low-cost option if it is needed to be viewed by users across many office locations, branches, retail outlets, homes or even countries. However, videos are less personal than face-to-face training and it can be less easy to track who has viewed them. Bespoke videos can sometimes be

costly to produce, and they are not always easily adapted if circumstances or technology change. There are companies that provide off-the-shelf videos covering subjects that may be appropriate to the enterprise or which may be adapted. Short video clips tend to work best.

Computer-based training (CBT), sometimes known as computer aided instruction (CAI), provides a similar, more interactive solution. Most CBT packages are able to offer a tracking system to record attendance and any scores from tests. Again, external companies may be able to provide an off-the-shelf solution at a reasonable cost. In some instances, a CBT module can be used for an initial information security campaign and then included as part of an induction course to ensure new personnel understand the security culture and requirements of the enterprise when they start. They can also incorporate a level of gamification that often encourages the more competitive.

Electronic formats such as workstation screen savers and emails can also be a useful way to deliver important or timely security messages straight to the desktop. Many organisations have intranets, and these can also be used effectively and at low cost to convey training and awareness messages. There are a number of external companies that produce awareness material, again at low cost, that can be adapted for use on intranet sites. The distribution may be quick and easy, but electronic methods rely on the end user choosing to access and read the information. They are particularly good for updating people about topical events such as a breach that may affect them either at work or home.

Escape rooms have gained popularity in recent years. These are essentially games where a small team is 'locked in a room' where there are numerous clues to provide the means of escape. This can be very cost-effective, taking perhaps only 30 minutes to run for a group, and so can be repeated for the whole workforce in a comparatively short period of time. If the time to escape is recorded, the competitive nature of some can be utilised with a reward of a prize of some sort and the key security messages can be incorporated into the clues provided – perhaps a password under a keyboard or a Post-it® note on a message board. If the target audience is more technical, then they could be given the role of trying to defend (or attack) the participants, thereby increasing the overall value for the organisation.

There are numerous visual aids that can be used to convey security messages. These include paper-based media, such as posters, booklets and brochures, or objects such as lanyards and coffee cups to display slogans. In fact, the more creative the better! They are relatively low cost to produce and can be very effective. However, there is no guarantee they will actually reach your intended audience, so they need to be supported by other methods. For example, posters can reinforce the messages delivered within a CBT module or a face-to-face training session. Finally, most people enjoy being given a small gift, so personalised items such as pens, puzzles, screen cleaners, sweets or stress balls can play a part in delivering key information security messages.

Measurement of success in training is a difficult area. A test or quiz at the end of a training session can gauge to some degree how much the participant has learnt. It can also provide a record that the participant has successfully completed the course for

either regulatory, administrative or due diligence purposes. Effective security training should attempt to deliver a positive change in user behaviour, which should lead to a reduction in losses from security incidents and reduce risk, but it is often difficult to properly quantify and measure how much value has been gained by the organisation.

Some organisations carry out exercises to see if their employees are becoming more aware of risks after providing training. An example of this could be to carry out a phishing exercise to see if people spot a rogue email that has the signs expected in an email containing malicious code to see how many people spot it and report it to their support function. It is essential that an exercise like this is managed carefully and done in collaboration with the relevant stakeholders, such as the support function and HR. It would be very upsetting to send an attractive email that is offering benefits when the organisation is about to implement a series of cutbacks. There are professional organisations that can do this for you, and it is best handled by someone who has the correct experience. There is, though, a danger that such exercises have unintended consequences, such as when staff refuse to acknowledge receipt of important information by clicking on a link. It is always important not to vilify the staff as a security risk or they might react in unexpected ways.

However, when compared with other control methods, awareness and training is a relatively low-cost control, and even a minor change in behaviour will far outweigh the costs of any investment. Going beyond the more formal sessions and campaigns, just being available to talk to people about their security issues at the coffee machine, for example, helps to make the security team accessible, reinforce important security messages and increase awareness. Having local 'security champions' in work areas can be a very effective and low-cost option.

Sources of information for training material

As the approach and content of a training or awareness programme needs to be tailored to the requirements of the enterprise, it is necessary to do a certain amount of research in selecting appropriate material. Specialist training organisations can help to source information and help to tailor an approach, but there are many other sources of information, many of which are online and can be accessed at no or very little cost.

Government departments and regulators give advice to individuals and organisations on how to protect their information, provide warnings of potential threats and offer news about information security problems. Some examples are: the Information Commissioner's Office within the UK, which has used YouTube to deliver short videos; similarly, America's National Counterintelligence and Security Center (US NCSC) provides links to awareness videos. Get Safe On Line is a UK website that provides advice to individuals and smaller businesses on protecting information, and much of the advice they provide for individuals can be adapted to help in awareness campaigns. The ENISA is an initiative that has produced user guides on how to raise information security awareness, which are also available in French, German and Spanish; the NIST also has a wide range of resources.

Learning from others via industry conferences and seminars can be useful sources of information as the issues discussed tend to be current and topical. They also provide an opportunity to network with industry peers who are generally facing similar challenges.

Some conferences are run by vendors and may be free to attend. Trade bodies are also able to provide industry specific content. Organisations such as the SANS Institute (Sysadmin, Audit, Network, Security) share good security practice and provide a wealth of information. There are numerous online newsgroups and bulletin boards that can provide relevant information as well.

Books and industry-based magazines and publications can provide many useful articles and features on a wide range of information security matters.

Last month GANT's Dr Peabody left her briefcase on the train. She realised that it contained a report containing details of some of the members. Fortunately, it was handed in and returned to her a couple of days later. Soon after there was a break-in to the office, and a small amount of money was taken, which caused a lot of disruption because membership papers that were on the desks had been strewn over the floor by the thieves as they looked for cash. Both events alerted Ms Jackson that the level of assurance training within GANT is not high and needs to be improved as a priority, especially concerning the protection of members' information.

ACTIVITY 5.5

1. What would you include in an initial awareness campaign and why?
2. What methods would you use to get your message across?

SAMPLE QUESTIONS

1. **What are the three main types of operational controls that can be used to protect information?**
   a. Confidentiality, integrity and availability.
   b. Vulnerability, risk and threat.
   c. Detective, reactive and preventative.
   d. Physical, technical and procedural.
2. **What are the different ways in which controls can be used?**
   a. Confidentiality, integrity and availability.
   b. Vulnerability, risk and threat.
   c. Detective, reactive and preventative.
   d. Physical, technical and procedural.

3. **What is considered the most effective approach to security?**
   a. Have as much security as the organisation can afford.
   b. Only use those security measures that are absolutely essential.
   c. Only use security measures that provide the best possible security available in that field.
   d. A layered approach with a combination of different measures for different risks.

4. **What is the main purpose of assurance training?**
   a. To prevent incidents from occurring.
   b. To ensure the organisation complies with legislation.
   c. To make people aware of their assurance responsibilities.
   d. To ensure that management objectives are achieved.

5. **Assurance training should focus on which of the following?**
   a. Topical assurance issues.
   b. Assurance issues that are relevant to the organisation.
   c. Organisational structures and management structures.
   d. All security issues.

6. **Which of the user groups below should receive security training?**
   a. All users of the organisation's information systems.
   b. Senior management.
   c. End users within the organisation.
   d. All technical and administrative staff.

6 TECHNICAL SECURITY CONTROLS

In this chapter the technical controls that are implemented to provide protection against security incidents are discussed in more detail. This includes the detection, prevention and mitigation of such incidents.

As discussed in the previous chapter, there are three main types of operational control:

- Procedural – for example checking references for job applicants.
- Product/technical – for example passwords or encryption.
- Physical – for example locks on doors and secure cabinets.

Of these, the product and technical operational controls are perhaps the most important in terms of information security since they are often the last barrier to illegal or unauthorised activity. As mentioned previously, this book deals with mainly generic controls because the more detailed information about specific controls is outside its scope.

TECHNICAL SECURITY

One of the main concerns about technical measures is the ease with which they can be overcome. However, some are undoubtedly very difficult to circumvent and, in the case of encryption, for example, and as discussed in Chapter 9, the best encryption techniques are essentially uncrackable in any real sense of the word.

There have been instances, though, where electronic locks have been used that were left in a 'safe state' of unlocked when there was a power failure, hence providing no protection at all. It has been acknowledged that some of the early attempts at technological security measures using tokens were less than successful when it was found that any credit card with a security strip could be used to operate the lock rather than just the authorised one. Fortunately, technical controls have improved over time and there is now a wide array of more robust controls that can be used. There are many sources that provide information on the types of approaches that are available. Organisations such as the NIST in America and the Centre for the Protection of National Infrastructure (CPNI) in the UK are two examples.

PROTECTION FROM MALICIOUS SOFTWARE

LEARNING OUTCOMES

The intention of this section is to provide you with the basic knowledge needed to put in place effective controls to manage the risks from malicious software. Once completed, you should have an understanding of each of the following concepts and be able to implement appropriate controls in straightforward situations.

Types of malicious software

The topic of malicious software is very large and could easily fill a book of its own. In this section the basics are described, and enough information is given to allow you to continue your studies elsewhere if you so wish. Malware (coined from malicious software), as it is often known, is one of the largest threats to the users and managers of information systems. An understanding of the capabilities of malware and those who write it, along with the controls that are needed to counter its threat, are essential for most information assurance practitioners.

A simple definition of malware is something like:

> An unauthorised piece of code that installs and runs itself on a computer without the knowledge or permission of the owner. It then conducts data processing and other operations that benefit the originator, usually at the expense of the system users or the recipient of the output from the malware.

The traditional idea of malware is the virus that infects a computer, attempts to spread itself to others, then trashes the contents of the hard disk or displays a message to show that it was successful in infecting the machine. A lot of the early malware did just this. Things have moved on, however, and the main emphasis now is not on 'spreading chaos while gaining kudos', it is about money. The Federal Bureau of Investigation (FBI) announced that, for the first time ever, organised crime gangs in America made more money from cybercrime in 2006 than they did from dealing in drugs. It is big business in many parts of Eastern Europe and the Far East too. The chances of being caught are much lower than for drugs operations and the sentences, even if the perpetrators are identified and convicted, tend to be much shorter.

Earlier malware authors simply wanted users to know that they had succeeded in infecting their machine, but now it has changed completely. The vast majority of modern malware authors know that if users realise they have an infected system the perpetrators have failed, because the user will attempt to disinfect it.

Modern malware can be split into the following major categories, depending on their payload.

> **Viruses**. These cannot spread on their own. They need to be attached to another piece of data or a program in order to reach and infect another computer. They are

often triggered by opening an email attachment or executable received by email or through removable media such as a CD or USB stick.

Worms. The difference between a worm and a virus is that worms contain the code needed to spread themselves without any user action. They will seek out other computers on any networks they can find. These can spread very quickly. It is estimated that in 2003 the Slammer worm infected 90 per cent of the world's vulnerable computers within 10 minutes of being released.[1]

Ransomware. This relies on viruses and worms being distributed to vulnerable systems. The immediate effect seen by the user is usually a message informing them that their data (or indeed the entire hard drive) have been encrypted, and that a ransom must be paid in order to decrypt it. One of the most visible of this type of attack was the so-called 'Wannacry' worm attack, which took place in 2017, affecting computers running Microsoft Windows®.

Rootkits. These are complex software packages that hijack the operating system and attempt to make themselves invisible both to the user and to the software designed to find and remove malware. They are insidious in that they still perform all tasks that the user requests, but they often make copies of sensitive data such as passwords, account details and logins and then send them to another computer, often to enable financial fraud such as identity theft.

Backdoors. The idea of the backdoor is to do just as it says. It provides a means for a third party to access the computer and use it for their own purposes without having to carry out the normal authentication checks. These can be used to turn the computer into a 'bot' (short for robot) that is effectively under the remote control of the attacker. It can then be used to distribute spam or act as part of a DDoS attack on a third party that cannot easily or quickly be traced back to the attacker.

Spyware. A common example of this is the use of malevolent cookies by websites. Some are designed to be permanent and to track and report the web usage back to a third party without the knowledge of the user. They can also log keystrokes and look for specific information such as bank account or ecommerce site login credentials. These can also be installed by software that performs a legitimate service, and freeware or bogus prizes are often offered as a means of getting a user to install spyware.

Trojans. The Trojan is the hackers' 'weapon of choice' today. Far more successful attacks use Trojans than any other attack vector. These are often disguised as another piece of software or are hidden inside compromised copies of other programs that users are lured into downloading and running. They often successfully avoid security countermeasures because poorly configured user accounts tend to have administrator privileges that allow the Trojan to run.

Another very successful infection route is through compromised websites. Trojans can download themselves without the user having to click on any buttons or links on the page. Simply going to an infected web page can be enough. More and more groups, criminal and otherwise, are writing increasingly sophisticated Trojans to attack computers in order to extract data, particularly via web protocols,

1 https://www.caida.org/publications/papers/2003/sapphire2/

where the malware scanning technology is often much weaker than the email countermeasures.

Active content. This is the means by which a Trojan is often downloaded to a computer running the viewing browser. Modern web applications use active code such as Flash, Java, Active-X and even mine headers to perform complex tasks within the web page to 'enhance the user experience'. There is no question that they are good at this, but they are also good at installing malware on the target computer. If the right level of security is not set in the browser policies, the compromised code will install and run itself on the target without the user having any knowledge of it happening. A typical attack is where a banner advert runs on a well-respected and heavily used website, with the code for the banner being supplied by a third-party advertiser. The attacker subverts the third party and adds the Trojan into the banner code. People view the website, thinking it trustworthy because of the reputation of the organisation, little realising that the advertising hosted there is busy trying to infect their computer. The payload of an active content/Trojan can be any of the forms of malware described in this section.

Whatever the type, detecting a piece of malware on a computer is a cause for concern and should be investigated without delay. It should also be noted that malware is actively and very widely spread; it is not a case of if you receive some malware, but when and how often. It is almost inevitable.

Zero-day exploits

No matter how good and comprehensive the defences in place, there is always a possibility that a new form of attack can get through them.

Hackers talk about 'zero-day exploits'. These are ones that have yet to come to the attention of the companies selling anti-virus and firewall products, so they have not issued an update to detect and remove them. In theory, these exploits can bypass the scanning engines because they are not on the 'stop' list that the updates contain. Some products are better than others in spotting types of behaviour and their analytical tools can identify many new versions of malware because they exhibit behaviour that is known to be unacceptable or has similar code to that found in other known malware. There is even a trade in zero-day exploits, with hackers selling the knowledge to others. Additionally, they are highly effective in cases where they have already been discovered, but have either not been blocked, or the user has not updated their system(s).

Routes of infection

Most of the routes of infection have already been mentioned in passing, but a more comprehensive description is provided here.

Infected media. Any piece of media that has been out of your control or supervision should be considered suspect – CD, DVD, USB stick and so on. It should be scanned for

malware, ideally on a stand-alone computer before being allowed into an operational computer. It may have been infected by any system with which it has interacted before it reaches your system. Even CDs that come with a magazine or as part of a special promotion should not be trusted. Do not assume they have been properly checked before mass-production. These have been issued containing malware on more than one occasion in the past, causing much embarrassment for the organisation giving them away. USB sticks are another source of infection. Malware can use them to travel from one system to another.

The most common routes today are via **email**, as an attachment or a macro in a document or even disguised as another file type, and through **websites**, as described above. Worms can propagate across **networks**, wide or local area, and may spread through unprotected systems.

It is also possible for malware to infect your system through a Bluetooth or infrared port. These should not be enabled unless they are required at the time and there should be a malware scanning application that protects those ports as well as the standard ones. If these functions are never used, the device drivers should not be installed if it can be avoided.

Smartphones and the increasingly complex software available for these types of devices, be they phones, MP3 players, tablets, iPads or similar, all have the capacity to be infected, some more easily than others. The idea that any one operating system is secure has also been shown to be false in recent years. The attractiveness of infecting one operating system or manufacturer's goods over another is often simply a matter of price – is it worthwhile to put in the effort to infect this type of device when there is fairly limited use of it by the wider general public?

An increase in staff being allowed to 'bring your own device', where staff may use their own technology to undertake their work, simply increases the risk to corporate IT infrastructures. The detail of providing security for these systems is beyond the scope of this book, but it can be very demanding and expensive. Depending on the level of security required and the risk appetite of the organisation (how safe your company's information needs to be), there may be a decision to be made whether or not to allow these devices to be used at all for any official business purpose.

Malware countermeasures

The countermeasures required to detect and defeat malware depend upon the configuration of the systems and networks to be defended, and they continually need to be updated to deal with the latest threats. A single computer, connected to a broadband connection at home, is very different from a global corporate network or a small organisation.

Even for the single user, because of the different possible routes of infection, a basic anti-virus package is not enough. The user requires a personal firewall package too. This will provide a defence against worms and web Trojans. Good-quality products also contain a profiling and access control tool. Once installed, they scan for existing malware and remove it, then build a profile of all the existing executables, putting them on a 'whitelist' of allowed products. Any new, unknown executable or active content

can be blocked from running unless manually approved by the user as the result of a prompt on the screen.

In an ideal world, large organisations that have separate systems to receive email and perform web browsing will need products or services for each system, for example:

- content scanning for web traffic and some means of controlling web access to stop prohibited sites from being accessed;
- email content and source checking software;
- firewalls that block ports and check content;
- network intrusion detection or prevention systems;
- 'sheepdip' malware scanners for untrusted media;
- personal firewall or application control software on individual systems, including checking files when they are accessed;
- use of managed service providers to scan mail and web traffic – inbound and outbound.

It is not the place of this book to recommend specific manufacturers' products, but it can list functionality that users should check for when acquiring such countermeasures:

- high degree of effectiveness in detecting and removing malware – read independent reviews;
- frequent and easy-to-deploy updates to signatures and scanning engines;
- ability to create and maintain a whitelist of accepted executables, active code and open network ports;
- support from a reputable company that can provide prompt updates and support to major threats;
- minimal impact upon operation of the systems.

Taking regular, secure backups is also a good way of countering malware. If something does get in and compromises the integrity or availability of data, it is possible to restore from the last good backup to minimise the impact upon the organisation. Use of the Grandfather-Father-Son system (GFS) (maintaining at least three generations of the backed-up data) is highly recommended to provide defence in depth and allow rollback to dates further back in time if necessary.

It is important to remember that there is a never-ending 'arms race' between malware writers and the developers of the countermeasures. The hackers are continually developing new ways to infect systems – new types of code and new routes of infection. Some malware is quite sophisticated and can even defend itself to some degree against countermeasures and other malware.

Methods of control

There are several approaches to controlling malware that need to be implemented at the same time if an organisation is to manage the associated risks successfully. The

first one is not always obvious and doesn't relate to any form of specialist malware application. This approach is patching. The operating system or application that does not contain any bugs or vulnerabilities has not yet been written. Patches and upgrades are released quite frequently, and every organisation should test and install patches at the earliest opportunity. Hackers keep a close eye on patch releases and the more capable ones will reverse-engineer the patch to identify the vulnerability it resolves. They then write or modify malware to take advantage of that weakness. The Slammer worm took advantage of a weakness for which a patch had been issued more than eight months previously. The worm was so successful because a large number of organisations had not applied the patch. The time from a patch being released or a vulnerability being described to an exploit appearing 'in the wild' is now down to as little as three days. Organisations must not only apply patches, but also do it promptly to provide adequate protection from new malware. User awareness is important too. Users that have been educated about the threats are less likely to click on a suspect link or fall for a social engineering attack that tries to trick them into loading malware.

Another approach is to 'harden' the operating system by not installing unnecessary features or applications and to ensure that default passwords and open configurations are not used. This is not the place to discuss the detail of how to perform these tasks, which is best left to experts. Suffice to say that an operating system installed using all the default settings recommended by the manufacturer is often very easy to compromise either manually or by malware.

The further approach has already been mentioned – use of anti-virus and personal firewall software. Some operating systems come with versions of firewall and malware-removal bundled in as part of the product. Experience and much independent testing have shown that these are often not necessarily the best products to use. Larger organisations need to investigate and select specialist products to protect high-bandwidth routes in and out of the organisation, such as email and web interfaces. Good firewall products also contain malware checking applications, and specialist appliances are available to monitor activity on internal networks.

The last, but equally important, approach is to harden the settings in the web browser in use. By default these often have much too low a level of security, allowing active code to run and accepting cookies from any source. Users should change the settings to accept cookies only from the original source and either disable active code completely or, at the very least, prompt the user to authorise a piece of code to run each time it tries to do so in the browser.

None of these products is of much use unless they are kept up to date. Many thousands of new items of malware are identified every day. The application and product providers issue regular updates to the signature files and sometimes to the scanning engines themselves. The same approach as for patching is required: download the updates and install them promptly to benefit from the protection they offer against new threats. Good products are capable of automatically distributing updates across the network to all clients, saving time and resources.

The officers of GANT have decided that they need to establish a better means of communicating among themselves and with the members of the society. Some members report that they have been targeted by persons sending them malware in emails or attempting to extract data about toad populations. The officers have no knowledge of this area of computing and need advice on how to protect their systems, at home and in the GANT office, against malware.

The loss or unauthorised disclosure of sensitive membership or toad population data would be embarrassing and potentially harmful to human and amphibian alike.

ACTIVITY 6.1

What advice would you give to the society with regard to the countermeasures they need in order to provide an adequate level of protection from malware?

NETWORKS AND COMMUNICATIONS

LEARNING OUTCOMES

The intention of this section is to provide you with the basic knowledge to understand the issues that organisations should take into consideration when identifying and managing the security risks to their networks and communication links.

Entry points in networks and principles of authentication techniques

There is an old joke that 'if it wasn't for the users, we wouldn't need security'. That can equally apply to the network and any connections to it. Not having a network would reduce the security requirement by a factor of 10. The network and communication links exist to make the systems connected to them available to authorised users. Unfortunately, it also makes them available to all the unauthorised ones. If there is an internet connection to the network, it means there is an entire planet full of potential unauthorised users.

The increasing use of services such as cloud computing and online applications such as Microsoft Office 365® means that an increasing percentage of work requires the use of a network that extends beyond the enterprise domain boundary. One of the factors that makes this possible is the increase in network bandwidth and the computing power to run the real-time encryption algorithms, such as internet protocol security (IPSec) and transport layer security (TLS) required for VPNs. Add to this the much greater capabilities of networking hardware to support and run a large number of sessions,

and the result is the increasingly connected world that now features ecommerce, the IoT and network television subscription services.

The challenge with increasingly large infrastructures and services is that they bring complexity and ever-higher levels of associated risk. The threats are also greater in size and potential impact; an attacker has the ability to disrupt large parts of the internet and an organisation's operations without attacking an asset owned by the organisation. How is it possible to even think about managing that kind of risk?

Any location, logical or physical, from which a user or device can gain access to a network is considered an entry point. Where the whole system is hardwired, these are fairly easy to define. These include but are not limited to:

- a laptop or PC in an office or home office;
- a console on a server;
- a broadband connection;
- a router for a connection from another network – internal or external;
- a firewall protecting a connection from another network – internal or external.

The problem is that most networks no longer have that kind of structure. In the years since the first edition of this book appeared, the technology behind networks has changed radically. Wireless networks are everywhere; Wi-Fi is used in offices because it's so much cheaper to install than running miles of ethernet cables. Many people have 4G data on their mobile devices and, at the time of writing, the 5G mobile networks are already being rolled out. Many new cars now come with an internet connection.

Almost everyone now has a tablet or smartphone, so even the nature of the endpoint has changed. Many companies are now embracing a BYOD policy where employees are allowed to use their own devices to connect to the network and process data. In the security world, BYOD also unofficially stands for 'bring your own disaster'. The management of the risks to confidentiality, integrity and availability and the governance of a device over which the organisation has no control and may not even be able to mandate the installation of an endpoint security product is, to say the least, challenging. The increasing proliferation of personal data and regulations such as GDPR that mandate their protection mean that network security is an ever-increasing challenge that cannot be underestimated.

Wireless networking

When any aspect of wireless networking is involved, the perimeter becomes much harder to define because of the ability of an attacker to use sensitive, freely available Wi-Fi antennas to greatly increase the effective range[2] from which they can access the network. An attacker can use the organisation's Wi-Fi network to view, download or upload unacceptable content or to conduct other criminal activities. This could lead to a visit from the police with a search warrant for activities that were not conducted by

2 At time of writing the Wi-Fi world record stands at 237 miles, from a mountain in Venezuela.

an employee and of which the organisation has no knowledge, but which used the organisation's connection to the internet.

The existence of a wireless access point (WAP) within a network will add enormously to the challenges of securing the network against unauthorised access. The fact that the hardware is relatively cheap and installing a WAP has been made so easy presents two more challenges:

- Users can buy and install their own hardware without the knowledge of the IT department.
- Wi-Fi security was designed by engineers, not security experts, and any version can be broken with freely available tools. Exploitable vulnerabilities were found in WPA3 and publicised within a few months of the code being released in 2018.

Many organisations now provide a 'guest' Wi-Fi network for use by visitors and contractors while on site. This needs to be logically separated from the internal Wi-Fi to avoid attempts to 'jump' on to the internal networks. The correct security architecture is very important in order to defend against this and detect any attempts to subvert the security controls.

The other insidious threat is that other organisations in close proximity may also be using wireless networking and users may accidentally, or intentionally, connect to the wrong network. There is a real risk of sensitive data being compromised by this kind of activity.

Identification and authentication

The principle and practice of authenticating to a network is very similar to that described in the section on user access controls for identifying and connecting to a computer. It may even be that a single sign on (SSO) system is in use that authenticates the identity of the user to the network and then grants appropriate privileges and access rights for all the systems for which that user has authority.

There are protocols designed specifically for centralised access control (e.g. Radius and Kerberos) and cloud-based services (e.g. OKTA) that work well for networks. These provide authentication of the user and software on a dedicated server. This may be just username and password, or it may involve some kind of token (e.g. RSA token/dongle) and code input.

It is also possible to use a challenge–response mechanism using a different communication system (e.g. a previously registered mobile phone) to receive a validation code for input as part of the authentication process. This is known as out of band (OOB) authentication. Another option is the use of a biometric such as a fingerprint, retina or iris scan to confirm the identity of the user.

An alternative way to authenticate to a network is the issuing of a digital certificate to a user, which is installed on their device or in their web browser and which acts as identification in the form of 'something you possess'. This is done as part of a public key infrastructure (PKI) and can be used as a way of identifying the device or user, depending

on how it is implemented. This kind of system is used in the Chip and PIN payment card to help secure a link to the payment centre and identify the user, in conjunction with the PIN.

All these are forms of multi-factor authentication (MFA) and can be used to authenticate at the network domain boundary or to an online service or website.

Partitioning networks

Partitioning a network is another way of protecting essential systems. This is sometimes referred to as subnetting or segmentation. The internet can be considered as the largest collection of subnets in existence, composed of all the public and private network domains that are connected to it. It is the same principle as physical access control, where doors with card readers are used to restrict access to sensitive areas of the office, or the need to know, where only a defined group of people are allowed to see certain types of information in order to manage the risks to it.

The rules on governance and the separation of roles within some business sectors, especially finance, require complete data separation to defend against insider trading and accusations of market manipulation. Network partitions can provide this function for electronic data.

Networks consist of cabling and devices such as firewalls, switches and routers that are used to connect sub-networks and enforce the rules on what data network protocols and users are allowed to pass through them from one network segment to the next. Part of the role of information security is to help design and define the rule sets (e.g. access control lists (ACLs)) used in these devices to protect data and ensure compliance with organisational, national and international legislation and regulation, e.g. GDPR, PCI DSS and HIPAA.

By using a network 'sniffer', an attacker can potentially record all of the traffic passing across any part of the network to which they have access. The sniffer may be a hardware module or some software (e.g. Wireshark) installed on a workstation or server as a Trojan to capture data and send them to the attacker for later use. The destination for this traffic is often outside the domain boundary so that it cannot be deleted by the victim if the monitoring is discovered. The network traffic will probably include useful sensitive data such as parts of the ID&A process, including usernames and passwords that might be in the clear or encrypted. The routing data in packet headers is also useful in helping to 'enumerate' the network; to build up a picture of its structure from the addressing data. In the same way that the dial code of a telephone number provides the geographic area where it is located, network addressing data in packet headers can help to build a diagram of the connections inside a network and how it connects to the internet and other networks belonging to partner organisations, suppliers and other offices used by a global company.

If the attacker can see the whole network, they can 'sniff' the whole network. Nothing is safe. Partitioning a network limits the amount of data that can be seen and makes the job of an attacker much harder. It is also true that partitioning can limit the damage done by malware. The chances are that any infection may be restricted to one network partition, reducing the effort needed to clean up the system to restore normal operations.

Without network partitions, an external attacker who defeats the perimeter security can access any area of the network with little impediment. An internal attacker doesn't even have to beat the defences, because they are already on the inside. The decision about how much protection to offer should be made through a risk assessment process, but there are certain common safeguards that should be considered by most organisations.

There are various approaches to partitioning networks, from physical cabling separation, the use of VPNs configured in network hardware using the IPSec, or even protocols such as multi-protocol layer switching (MPLS). A department or site may have an individual local area network (LAN) linked to others via routers to form a wide area network (WAN). Each has its good and bad points, ranging from strength of security to cost and network bandwidth overheads. The appropriate solution will depend on the outcome of operational requirement, risk assessment, risk appetite and budget.

The DMZ

Any connection to the outside, such as the internet, should be protected by at least one firewall, ideally a pair with a logical gap between them, known as the DMZ. It is where any traffic from the internet or other locations outside the enterprise domain boundary should be terminated, inspected and authenticated if necessary. Traffic that doesn't need to go deeper into the network should have a server here to act as a protocol break. Examples include email, external web servers, encryption and remote access terminal servers. In addition, technology such as Wi-Fi and IoT devices should have their connections to the enterprise network routed through the DMZ. Any compromise of the radio frequency Wi-Fi network, as discussed earlier in this section, requires the attackers' traffic to pass through the same protocol break and undergo inspection by the security controls before it can either enter the domain through the inner firewall or leave the domain to access cloud services and so on through the outer firewall of the DMZ.

Implementing a DMZ and additional controls within it means that any successful attack on the protocol termination point through the external connection does not immediately grant access to the whole internal network and the data it contains.

It is important to check traffic both entering and leaving the DMZ to identify incoming attacks and unauthorised data that are leaving the domain (data loss prevention, also known as data leakage). A good example of this is the use of a cloud access security broker (CASB) to enforce rules on access to data and services in the cloud both from inside the domain by an organisation's own users and also by any third parties to whom access as a service provider or a customer may have been given. There is a lot of good advice and theory available on the services that should be located here and how to inspect them. Modern heuristic analysis systems exist that use Bayesian statistical analysis to detect unusual network traffic before alerting the SOC and blocking that traffic. Their use is highly recommended.

Cryptography in networking

Cryptography is described in more detail in Chapter 9, but some basic concepts need to be understood now. There are two common mistakes many people make when they think of cryptography. The first one is that they think it stops people from being

able to see their data. This is not the case. Attackers can still see the data, but if the cryptography is correctly used it means they can't understand them. The second one is that they think cryptography is only used to provide confidentiality. Once again, this is wrong. The four main uses of cryptography are:

- secrecy – nobody else can see the plaintext;
- data integrity – the data have not been changed, deleted or inserted;
- user verification – this is the person they claim to be;
- non-repudiation – the sender cannot later deny sending the message or its content.

Different forms of cryptographic algorithm and technology can be used and combined into protocols to perform different tasks. For example, digital signatures are a form of cryptography and do not normally provide confidentiality; their main function is to provide non-repudiation and data authentication.

Data travelling across a network are obviously in transit, but it should not be forgotten that a network provides access to data that is at rest, on a hard drive, in the cloud or stored on other media. The security architecture must protect both transiting and stored data. Good operating systems also use encryption across networks, especially when sending passwords. This feature may not be activated by default; it is always worth checking. This defends against the capture of passwords by attackers with network access 'sniffing' the traffic. This is increasingly important as more data are stored in a remote location somewhere in the cloud. They need to be protected against being read while at rest or in transit to and from the cloud as they are processed.

The most obvious form of cryptography that most people see and use is TLS, which provides encryption for websites, especially ecommerce to protect financial data such as credit card numbers, and in many smart-metering networks, to protect the energy usage and billing data passing to the energy provider over the network. The browser user does not normally need to do anything other than to check the TLS certificate to make sure that it is valid and belongs to the organisation with whom they want to do business. All the configuration is done in advance by the operator of the website. When a user connects to the site, their browser and the website set up an TLS session using a unique session encryption key to protect the data from being read by a third party as they travel across the internet.

In business, the increase in mobile working has meant that there has been a steady rise in the need for VPNs. These are another way of encrypting (protecting) traffic that travels over a public connection, which could be the internet via a fixed or wireless broadband connection. The common risk with all of these connections is that the data are travelling across a system that is owned and administered by people unknown to the user and therefore not fully trusted. It is also possible for a third party to compromise the channel and eavesdrop traffic in transit. Attacks such as domain name system (DNS) hijacking and rebinding can be used to reroute traffic so that an attacker can copy it in transit without the knowledge of either end – a form of man-in-the-middle attack. That is why cryptography is used to create a VPN. The data part of the traffic is encrypted before leaving the sender until after it arrives at the receiver, leaving the address part in the clear so that it can be read and routed by the public network. It's like putting the data

in a letter that is sealed in an envelope instead of sending a postcard. The address can be read but not the content of the letter. Encryption in this context is how to stop a third party from reading the message.

Encrypted traffic could be a one-off session using a unique key (e.g. TLS) or a more enduring connection such as IPSec between two or more geographic locations that have a high rate of traffic flow. This could be done through the browser using a form of digital certificate or symmetric encryption key (see Chapter 9), or it could be that IPSec is implemented from a network device such as a router without the user having any knowledge or needing any configuration of their device. In many cases secure communications originate from the user device, which is increasingly common with the use of laptops, tablets and apps on smartphones. Modern computing power, on-chip encryption capability and algorithms make this possible. The user has to identify and authenticate themselves in the usual manner. Once the identification and authentication is complete, the host and client agree on a secret key and the encryption process starts. From then on, the body of the data is encrypted and protected from eavesdroppers. The concept of the VPN can also be used to separate internal network traffic, as described in the previous section, to ensure it cannot be read by those without a need to know.

Control of third-party access

The concept of allowing a third party access to the organisational network is not a new one and it is no longer considered to be an unusual requirement. The original use cases included a dial-up connection from a supplier, used to remotely support hardware or software, or a 'road warrior' who travelled as part of their job, or a car manufacturer using a 'just-in-time' approach to manufacturing by placing electronic orders with suppliers for carefully timed deliveries of components and so on. Interconnection is the 'new normal' and for many younger people it has always been the case. In the internet age, the concept of electronic data interchange (EDI) is not the novelty or challenge it once was. What used to take networking experts days to make work reliably can now be done in seconds and the bandwidth is measured in gigabytes, not kilobytes. The world is ever-increasingly connected, and the advent of 5G mobile networks is gradually making the access to large quantities of data, such as streaming movies, an everyday, eventually an everywhere, occurrence.

It is also the case that the business model has changed dramatically, partly driven by the change in the way that network and communications technology has evolved. Economies of scale and efficiency mean that it's cheaper to host one's data and many service offerings in a cloud environment. It's also faster to implement because one doesn't have to get quotes, place orders, wait for hardware to arrive and a new data network connection to be installed (which could take months). An organisation can now 'spin up' a new cloud instance in a matter of minutes. All that is necessary is an internet connection and a payment card. In fact, a problem that is now faced by information security departments is the increasing use of 'shadow IT', where users are purchasing unofficial cloud services for use by their department without the knowledge of IT or security. This leads to unknown, and hence unmanaged, risks.

The use of progressively smart mobile devices also means that users are doing more on the move: interacting with data, live streaming material and communicating with each

other using social media. The term 'smartphone zombie' has come into use to describe the way people now behave; in Germany they coined the phrase *kopf unten*, which loosely translates as the 'head-down generation'. These mobile devices can be used to access either the internal network or organisational cloud services from anywhere in the world, and do it from an unmanaged device, which is why many security professionals came up with the unofficial translation of 'bring your own disaster' for the acronym BYOD.

It is becoming common practice that third parties are granted some form of privileged access to an organisation's IT over the internet. In this age of outsourcing and online support to drive down cost, the provision of services is more likely to be over a remote connection than somebody coming to the organisation's office. This is also driven by the increase in connected devices, such as the IoT where diagnostic data can be sent to a central monitoring facility. Rolls-Royce monitor many of the jet engines they have sold to airlines in real time. The engines send data while they are in flight to help their maintenance programme and improve fuel efficiency. The risk is that if there is a connection for legitimate use it might be compromised and used by an attacker to gain access for their own ends. This is an increasing problem, where the third party might be a weak link in an otherwise strong defence chain; in 2013 the USA retailer Target was successfully compromised for one of the largest data thefts in history through the connection from the company that supported their air-conditioning systems.

It is important to think about the risks from connections and to manage the associated risks. Ensure that the design is properly monitored and defended and segregate the network to limit access to only the subnets and assets with which third-party users need to interact, and protect the link with a VPN. It is also important to stipulate the level of effective information security that the third party must have in place before any link to them is created. This might be by requiring them to be accredited under an international standard such as ISO/IEC 27001 or it might be a specially written CoCo that defines certain requirements. However it is done, the contract should, where possible, specify the right to conduct no-notice compliance audits to make sure that they are doing what is required in the contract. This may not be possible for some smaller organisations with big service providers, such as the global cloud vendors. They offer a 'take it or leave it' service provision with a standard security model at a price so cheap that many organisations will accept the associated risks.

The partitioning of the network to limit the areas that the third party can access is another example of the need-to-know principle. They may be a business partner, but they probably do not need to know much about the organisation that isn't in the public domain. There may even be regulatory requirements governing this access (covered in a later section). The primary concern is to ensure that the access point can only be used by authorised persons or applications from within the third party. Identification and authentication are still required to stop attacks across the link by third-party staff or anyone who manages to find a way to connect into the link. The standard approach to protecting the link itself is through the use of cryptographic methods, such as VPNs, as described in the previous section, and for ID&A purposes. A good design will normally have the link to the third party located within a DMZ, protected by a firewall from the outside world and another one that only allows permitted and inspected traffic through into the organisation's inner network to access a specified server or asset and vice versa.

Network usage policy

The network usage policy exists to define the purposes for which the network may, and may not, be used. It will also define the individuals and roles who are allowed to use it and the official line on access control. This will include definitions of the user profile for each role – privileges, password lengths and strengths, renewal period and so on. This will be part of the ISMS for the organisation, but will also define the controls required to manage the risks of access by any third party.

Intrusion monitoring, detection and prevention

It has already been mentioned that networks are often attacked from the outside by unauthorised users or by authorised users within the organisation attempting to perform tasks for which they are not authorised. It is important that the network has some means of detecting, reporting and even blocking these attacks using automated tools. This is part of the role that is generally referred to as 'protective monitoring' or ProtMon. Modern ProtMon systems work either on a knowledge basis, using known signatures, or they look for usual behaviours, which can spot zero-day or highly advanced attacks. Both have their strengths and weaknesses.

The first task is to ensure that all relevant event log data are recorded securely in such a way that an attacker cannot change or delete the information in order to cover their tracks from investigators and auditors. This can provide evidence of what happened and be used to identify any damage done and how it was achieved. Many organisations will keep at least six months of log data so that they can go back to analyse it in case they discover a security breach. The data can be analysed for signs of when and how the compromise occurred and what the attacker might have done since then. Modern ProtMon technology can do more than just look at event logs. There are tools that can conduct Bayesian statistical analysis of behaviours of users and network traffic to identify unusual behaviour that might indicate signs of a compromise by an attacker or malware. These are powerful tools and can provide good insight into the state of the network or a user and detect activity that would be missed by an endpoint protection system on a system host.

Modern networks are so large and so powerful that it is impossible for even a team of people to inspect manually a fraction of the data, let alone correlate actions across devices and networks. That is why automated tools such as SIEM systems are used. They can collate and analyse data looking for unusual activity, often combining signatures and knowledge-based approaches to produce high quality alerts that require investigation by analysts. The primary functions of a SIEM system can be defined as:

- data collection and aggregation;
- correlation of data from different sources;
- reporting and alerting;

- data retention;
- analysis of compliance with policy and standards;
- tuning and development;
- continuously performing all of the above in a secure and reliable fashion.

As with any kind of product, none of them is infallible, but the good ones will detect and stop the vast majority of attacks. It should be said that this subject requires good knowledge and experience if it is to be performed well. There is no substitute for hours spent studying. Courses and external websites can be used to gain knowledge and keep current with new techniques. Know your enemy and their modus operandi.

Some organisations configure their protection monitoring (ProtMon) tools to do more than just detect and alert, because an attack can spread so quickly and cause huge disruption (e.g. ransomware). It is therefore possible to configure some to take automated action to respond to and to stop attacks quickly. This 'intrusion prevention' capability does run the risk of a false positive leading to an unnecessary response that disrupts legitimate operations, so their use must be carefully planned and subject to a lot of acceptance testing before they go live. Ongoing tuning will probably be required to improve their performance and reduce the risk of disruption.

Firewalls

A firewall is a network security device that is used to restrict access to assets such as data and systems in accordance with a defined set of rules. Modern firewalls can also apply contextual rules that depend on the direction of traffic, the user, the application and the port in use, to inspect or block traffic according to the policy and firewall rule set that has been created. It should be noted that firewalls can only restrict access by blocking communications, not enabling them.

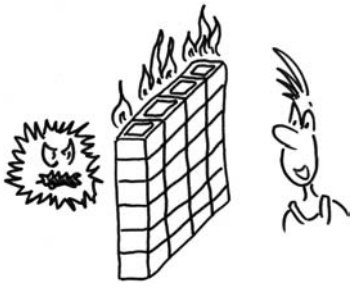

It is important to recognise that firewalls, although much more capable than they used to be, only offer limited protection and are just one part of a complete network security infrastructure. A point that is often forgotten is that firewalls cannot restrict network traffic that does not pass through them; they provide access control between network segments. The typical model is that the network on one side of a firewall is regarded as the inside and traffic from it is 'trusted'. Incoming traffic is 'untrusted' and all traffic must be inspected and validated before it is allowed access.

There are several different types of firewall, with Stateful Inspection and Next Generation being the two most effective at time of writing. Like any technology, they have strengths and weaknesses that must be taken into consideration when designing the network architecture.

Secure network management

The basic technical elements for network security have already been discussed, so it is time network management was considered from the perspective of the department manager and senior management. The task of managing a network securely is one of the most crucial aspects of IT service delivery. No network means no communications, no security means the organisation is open to loss of data, intellectual property, revenue and reputation. Any one of these can put an organisation out of business; an insecure network can easily cause several of these at once.

The huge increase in the use of networks, the connections outside the domain for ecommerce and the use of cloud services mean that organisations are more vulnerable than ever to attacks that occur on or across their networks. The speed of communications and an interconnected world means that a successful attack on a single system by whatever method can spread through the network like wildfire. The Slammer worm is believed to have spread around the world, compromising most of the 75,000 known infected systems in less than 10 minutes.

If the attack is ransomware, the impact can be huge. The shipping line Maersk was hit by the NotPetya attack in June 2017 and lost the use of their global IT systems, all of them. They had 50,000 infected systems in 130 countries. The systems took 10 days to recover to a basic level, but many months more to recover fully, and the loss of revenue is currently estimated at US$250–300 million.

The rate at which infections occur means that proactive patching, daily updates to malware signatures and preventative controls are the correct approach. There may not be time to patch or block a port if an alert is put out. It took just a few minutes for the NotPetya attack to reach over 50 countries, and reportedly less than 30 seconds to disable Maersk's worldwide Active Directory. The faster the processors work, the shorter this time will be!

The attackers are increasingly sophisticated at creating zero-day attacks and the release via WikiLeaks of a whole series of government agency attack tools has helped attackers around the world to raise their game. The formal assessment is that organised crime gangs are no more than four years behind the ability of nation states to conduct cyber-attacks, and that gap is closing.

The use of network monitoring tools is essential, and ones that have an ability to learn how the network behaves are vital, so that they can alert the organisation to unusual activity and take proactive action to contain infections such as ransomware. The segregation of networks into subnets and secure enclaves is also important to limit the spread of infections, the scale of any disruption and the amount of data that an attacker can access. The reality is that a network will be compromised at some point; not if, but when, and even then it might not be detected for some time. The 2017 Ponemon Institute report found that US companies took an average of 206 days to detect a data breach.[3] Data breaches are becoming an all-too-common event because personal and payment card data has significant financial value to criminals.

3 https://accenture.com/us-en/insight-cost-of-cybercrime-2017,

Some business sectors require minimum standards through legal and regulatory controls such as GDPR, HIPAA, Gramm-Leach-Bliley Act (GLBA) and so on. Others choose to implement them to comply with standards such as ISO/IEC 27001 and the EU's NIS directive. Network management can play a major role in managing risk and improving resilience for business continuity and compliance purposes.

In order to manage their business effectively, any organisation needs to have contextual information about their infrastructure, especially their:

- assets – physical and logical;
- architecture – systems integration and interconnectivity;
- risks – threats, impacts and vulnerabilities;
- countermeasures – logical and physical defences;
- dependencies on third-party service providers – support, cloud and so on.

In addition, a good management team will understand the:

- business processes that the IT systems support and service;
- organisational policies for IT, quality and conduct of operations;
- procedures or processes for all tasks;
- need for and value of effective communication routes within the IT department and with other departments;
- legal and regulatory requirements that apply to the organisation.

If this information doesn't exist, work will be required to create or develop them in agreement with all the business areas and then to implement and operate them. Ideally, this should be done to a recognised framework like the ISO/IEC 27000 series, the IT Infrastructure Library (ITIL), the SABSA matrix or the standard for the relevant industry sector.

It is highly advisable to follow the Plan–Do–Check–Act model (aka the Deming Cycle), as previously mentioned and as shown in Figure 6.1.

This provides a cycle of continuous monitoring and improvement in a demonstrable way that can be provided as evidence to external auditors. Don't forget that the management of anything requires metrics; you can't manage effectively what you can't measure – how do you know it is working, or how well? Networks are no different. Decide on how secure it needs to be and how you will know when you achieve it. What are your critical success factors and why – how do they support the operational objectives of the organisation? Make sure that you monitor and report on achievements against targets regularly to justify your budget and team.

Figure 6.1 The Plan–Do–Check–Act model

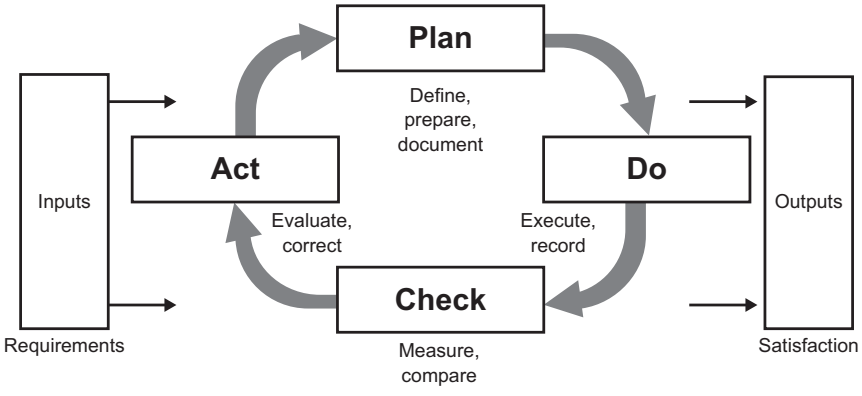

Considering all the data that GANT has to look after, it is clear that there are a number of occasions when encryption would be an appropriate part of the controls used to safeguard the information. The use of encryption would help to reduce the risk to confidentiality in particular, although it can also help with other aspects of security as well.

ACTIVITY 6.2

What advice would you give to the officers of GANT explaining why encryption should be considered? You should cover:

a. Which sets of information should be considered for encryption.
b. The benefits of using encryption.

OPERATIONAL TECHNOLOGY

LEARNING OUTCOMES

The intention of this section is to provide you with the basic knowledge to understand the security issues of operational technology, that is, technology that is increasingly being connected to networks and the internet despite not being designed with that purpose in mind.

There is a cumulative trend to connect ICSs, supervisory control and data acquisition (SCADA) systems and distributed control systems (DCSs) to networks and the outside world. Collectively these are referred to as operational technology (OT). They are the systems that control machinery, electrical power generation, transmission and distribution systems, oil and gas platforms, refineries and many other systems of that nature. They use very different protocols (e.g. Modbus, Open Platform Communications and Profibus) for control and to communicate and follow an architecture (often based on ISA/IEC 62443) that does not look like a standard enterprise network architecture. They have traditionally been the responsibility of the process control engineer to design and manage.

The constant management drive for 'faster, better and cheaper'[4] means that these systems are now connected either directly to the internet or their traffic passes through the enterprise domain to provide management/reporting data to ERP systems and global monitoring, using a follow-the-sun approach, and then on to the internet. The connection of OT systems introduces new challenges and vulnerabilities that are very different from those of enterprise IT. Information security practitioners who find themselves having to work with these systems need to consult with experts in order to manage the risks, which have different vulnerabilities and potentially safety-critical impacts to life and property if they go wrong. An example is the UK's Buncefield explosion of 2005, caused by multiple OT and process failures.

Vulnerability analysis and penetration testing

An even more demanding task is that of analysing systems for vulnerabilities and performing penetration tests (PenTests). PenTests are sometimes referred to as 'ethical hacking' because the testers will use many of the techniques that would be used by a hacker in order to identify any vulnerabilities in the network and applications. Only the most skilled and dependable of specialists should be allowed to conduct this kind of work as it is very easy to adversely affect the availability of systems and the data themselves if they don't have the right knowledge, experience or tools. There are also significant legal issues to consider before undertaking any form of penetration testing. It is advisable to use only professional, accredited PenTest organisations who have been independently assessed by an industry body such as CREST and have individuals with qualifications such as:

- Certified Ethical Hacker (CEH);
- Global Information Assurance Certification (GIAC) Penetration Tester (GPEN);
- Offensive Security Certified Professional (OSCP);
- Tiger Scheme certifications.

There are also vulnerability scanning tools such as Nessus and Nmap that can be used by penetration testers or administrators to look for unpatched parts of the operating system, open network ports and incorrect device configurations. As with Wireshark, training and practice is needed to use them well and they should not be used in some networks, especially those using operational technology/industrial control/SCADA

4 To which the normal response from IT is: 'We can give you any two of those. Which ones do you want most?'

hardware in order to avoid the risk of unintended disruption. They will also often only find the more common vulnerabilities, which will be useful, but not necessarily enough to provide full confidence. So, if in doubt, an expert should be involved in their use. This task is best done by someone who knows the network in conjunction with someone who understands security.

Vulnerability analysis is the process of examining the network for any vulnerabilities that could increase the frequency or impact of any threat. An example would be an unpatched network router or firewall for which there is a known exploit. Vulnerabilities are often not just weaknesses that allow access to data, but the ability to facilitate DoS attacks too.

Owing to the possible implications, there is a lot of paperwork to be completed before the work can start, including a detailed briefing document defining:

- the terms of engagement;
- what is in and out of scope for testing;
- acceptable levels of disruption (if any);
- level of social engineering allowed or expected;
- tools and techniques to be used;
- timing and format of reporting;
- secure deletion of data obtained during the test;
- actions upon finding a vulnerability – major and minor;
- use of a non-disclosure agreement.

Anyone who has not specified or managed a PenTest before is strongly advised to seek advice and guidance from someone who has.

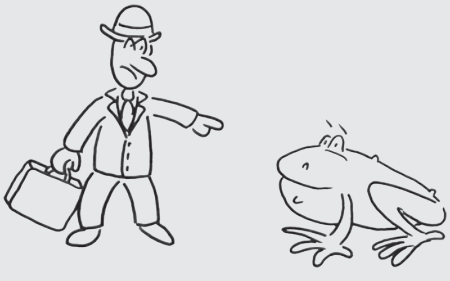

The success of GANT has led to the organisation growing in size and the recruitment of a team of wildlife surveyors to look for the toads across the country. These people are out in the field and need remote access to the IT systems for reference and reporting purposes.

In addition, there will be a national campaign to get members of the public to report sightings through a website into which they will enter data. Access to this must be secure enough to stop it acting as the start point for a remote attack, yet allow anyone to interact with it to input valid data.

This requires a new network structure and remote access capability – broadband and web-based methods will all be required.

ACTIVITY 6.3

One of the directors has been told about the ability to connect into the office from home by a friend in the pub, and wants to be able to do the same for GANT. How would you explain the security issues that surround the use of remote working to him?

ACTIVITY 6.4

There are concerns that the network is being accessed by people who do not have the necessary authorisation. How would you identify the right place to install an intrusion detection system and its sensors?

ACTIVITY 6.5

GANT has been approached by the directors of the Society for the Listing of Undiscovered Gastropods (SLUG) who are suggesting that their survey teams could work in conjunction with those of GANT to cover more ground. How would you design the security architecture for a data connection between the two organisations?

EXTERNAL SERVICES

LEARNING OUTCOMES

The intention of this section is to provide you with an understanding of the security issues surrounding services that use the network, often bought in from external suppliers.

Securing real-time services

The rapid rise in popularity of instant messaging (IM) services such as WhatsApp, Facetime, Skype and other forms of text, audio and video conferencing have added

another dimension to the challenges facing information security managers. There are already examples of IM being used:

- to extract data;
- to insert malware onto networks;
- as a channel for phishing attacks;
- for unauthorised purposes leading to legal action against the perpetrators.

Video conferencing isn't necessarily quite as vulnerable. Some organisations still use separate integrated services digital network (ISDN) or other data connections that are not linked to their data networks. The data can, however, still be the subject of eavesdropping, leading to a loss of confidentiality. Systems using webcams or sharing data connections have the same risks and threats as the data channel and can be used as an easy backdoor into the network if not properly segregated and protected.

Other real-time services, such as ordinary telephony, voice over internet protocol (VOIP) and CCTV feeds, are also possible avenues of attack. VOIP is especially vulnerable if it is integrated into a single messaging system. Those with data connections can be used as a route into the organisation's data networks. Ordinary private automatic branch exchange (PABX) systems can be the subject of various technical attacks (some of which are known as phreaking and dial-through fraud), leading to losses in the millions if they are not configured, protected and monitored effectively. Just because it isn't like other data formats, in documents for example, does not mean it won't be attacked. The enterprising attacker has known for a long time that anything related to telephony is vulnerable to attack. All an attacker has to do is find the right number, dial it and they have a connection.

A similar tactic is also used to identify unencrypted Wi-Fi access points, in which attackers search for available Wi-Fi networks, and can easily see from their smartphone, tablet or computer whether or not the access point carries any encryption. If it does, they can tell if this is the older wired equivalent privacy (WEP) protocol, which is easy to overcome, or the stronger Wi-Fi protected access (WPA) protocol, which has a number of variants, and is much more difficult to subvert.

Quite often, attackers will use a tactic called war dialling, which is to ring every number the company has and see which ones have a modem attached or which will allow access to the main telephone exchange control system. War dialling can also be a useful tactic for the security manager; security auditors have quite often used this technique and found unauthorised modems connected by users that the IT department knew nothing about. However, dial-in modems and ISDN connections are much less common since the introduction of broadband internet connectivity.

Since many of these services are quite new, the technology available to protect them is also new and may not be as mature as products that protect against other threats. That means they may still have vulnerabilities that can be exploited. Attackers could well target these as being the weakest spot in the defences.

Securing data exchange

The exchange of data over the network needs to be protected against threats to confidentiality, integrity and availability. Data must arrive without being altered, copied or subjected to eavesdropping. The ability to send data whenever required must also be maintained. It doesn't matter what forms this data takes the same principles apply. It is merely the countermeasures used to protect the data that will vary. Cryptography and security protocols can be used to perform this function for data in transit. The key issue is to ensure that all parties protect the data to the same standard. If one does not, then they risk being identified as the easy target and the additional protection at the other locations will count for nothing.

The last point to note is that, once data arrives, they must be checked for any signs of malware or compromise before being allowed access or given any credence as legitimate traffic. This should be conducted in the DMZ, described previously, before passing through into the inner network.

The protection of web services and ecommerce

In business-to-business relationships there is normally a lower degree of risk when EDI occurs. A level of trust is often established by some means before EDI begins. Security architects must remember that the users of web services and ecommerce are often members of the public and so organisations have no control over the configuration and integrity of the PC being used to access the service being provided. It is therefore important to consider the possibility of malware such as infectious Trojans or key loggers being installed on the user's PC and to design security to protect the servers providing the functionality.

There is also the obvious issue that websites are normally public facing and therefore open to attack by anyone with an internet connection. It is estimated that as many as one in three of all websites have been compromised with malware at some time. Protection must be present to stop attackers from extracting data, entering false data and adding their own code to the site, either for propaganda purposes or to add malware that is downloaded by any visitors.

The most obvious form of cryptography that most people see and use is TLS (as previously discussed), or SSL – when a user connects to the website, their browser and the website set up a SSL channel to protect the data from being read by a third party as they travel across the internet, providing the encryption for access to websites, especially ecommerce, to protect financial data such as credit card numbers. Although the name and abbreviation are still in widespread use, SSL was updated some years ago to become TLS, which offers a more robust level of security for data being exchanged during transactions. Additionally, the secure hypertext transfer protocol (https) is increasingly used where websites are secured by an SSL or TLS certificate.

In business, the increase in mobile working has caused a steady rise in the need for VPNs, also described earlier in this chapter. This is another way of encrypting (protecting) traffic that travels over a public connection, which could be the internet, fixed or wireless

broadband connection. As mentioned earlier, the risk is an untrusted system over which the user's data must pass. It is possible for a third party to compromise the channel and eavesdrop on traffic in transit.

The system uses VPN client software on the remote system to contact the host server over a public channel. The user has to identify and authenticate themselves in the usual manner. This is all done in plaintext but once the ID&A is complete, the host and client agree on a secret key and the encryption process starts. From then on, the body of the data is encrypted and protected from eavesdroppers. The concept of the VPN can also be used to separate internal network traffic, as described in the previous section, to ensure it cannot be read by those without a need to know. VPNs are becoming used more often on private computers as an additional means of ensuring privacy of connections over the wider internet.

Protection of mobile and telecommuting services

More and more people are spending time out of the office travelling or working from home. This increase has been facilitated by new technology that allows improved remote access, with broadband at home and in hotels, and wireless networking. While still supporting the older services, such as Global System for Mobile Communications (GSM; 2G), General Packet Radio Service (GPRS), Universal Mobile Telecommunications Service (UMTS; 3G), High-Speed Downlink Packet Access (HSDPA) and Enhanced Data Rates for GSM Evolution (EDGE), the mobile phone companies have rolled out 4G and long term evolution (LTE) services, which deliver considerably greater bandwidth. The next (fifth) generation, or 5G service, is now being deployed, and, while it promises to deliver much greater capacity and functionality, it will also potentially increase dramatically the security issues, not least with the 5G technology itself, that has been the subject of intense debate as to whether some of the manufacturers of the 5G core and radio networks are able to intercept voice and data transfers and pass the traffic to foreign powers.

Securing the systems in the office that receive this kind of traffic has already been discussed, so in this section emphasis will be on the elements that are 'out on the road'.

The three main problems facing assurance practitioners here are:

- The connection uses network infrastructure that does not belong to the company, so traffic can be more easily viewed, altered or deleted by an attacker.
- The users take their IT and communications equipment away from company premises, making it more vulnerable to theft, loss or compromise.
- Ensuring that connections are only used by authorised employees.

The first problem can be defended against with encryption. Creating a VPN tunnel from the user device back to the office can defeat all but the most determined attacker if it is implemented properly (as described in previous sections).

The second challenge can be partly safeguarded with encryption to protect data held on devices carried off site. This can either be at file level or, where possible (a much better solution), the whole of the device, usually by encrypting the entire hard disk drive.

If hardware is stolen, the attacker cannot login to the device and read the data, so all they have done is stolen a device to reformat and sell, not access valuable company data. The other part of the equation is to make sure that the users have received appropriate security awareness training about mobile working and are issued with good physical locks to secure their equipment. Part of the awareness training should be about working in unsecured environments: who can see the tablet or laptop screen and paperwork or overhear sensitive conversations?

The last part is to make sure that any communications ID&A process includes a PIN or token code, and that devices capable of remote communications can have their service disabled quickly. This stops the attacker from being able to access an organisation's network and from running up big bills with the service provider.

The ISO/IEC 27000 series of standards has been enhanced to include ISO/IEC 27010 – Information security management for inter-sector and inter-organisational communications.

Secure information exchange with other organisations

The process of securing a connection to a third-party organisation has already been covered, but there are more than just the technical issues to consider. It was briefly mentioned that there may be regulatory or legal requirements governing data interchange, and now is the time to go into more detail. The main legislation to consider in the UK are:

- DPA;
- GDPR;
- Human Rights Act (HRA);
- Financial Services Act (FSA);
- Official Secrets Act (OSA) for government and defence projects;
- Markets in Financial Instruments Directive (MiFID);
- Freedom of Information Act (FoIA);
- Computer Misuse Act;
- Communications Act.

Without doubt, the most important of these are the DPA and the GDPR. These define very clearly how personal data are to be protected and used, taking into account the rights of the data subject as defined in the DPA. Other legislation will relate only to the financial industry (e.g. FSA and MiFID), but is equally important to them.

When two or more organisations plan to work together, the important start point is for those organisations to agree and sign a protocol that specifies all of these matters as part of a legally binding contract where all parties agree to common standards for the processing and protection of data each provides to the other. Each party is then bound under law to a duty of care. All parties are then said to have shown due diligence and have defence in law (and usually the right of redress) against wrongdoings by the other.

Service management considerations

Many organisations are now moving away from in-house IT departments, and outsourcing their requirements to third-party organisations for the design, deployment and support of all the services they require. This can also include the provision of cloud-based services, described in more detail in the next section.

One of the most important considerations for an organisation in this situation is that of the service management contract, which must include the overall security requirements to be incorporated into the project. Occasionally, these contracts are developed by a financial team, who may not regard security as an integral part of the contract, and the organisation's security team should always have an input to this document in order to save potential problems further down the line.

The directors of GANT have decided to open up an ecommerce site to sell toad-related merchandise and host a forum dedicated to amphibians in general. This will be in partnership with several other wildlife groups working with other amphibians native to the UK. The economies of scale have been recognised and welcomed by all parties.

In order to monitor stock levels and pass orders back to the right group for dispatch, there need to be secure links and data sharing agreements created.

ACTIVITY 6.6

The directors want to know how to protect GANT against malware contained in messages posted to the proposed forum. What would you advise them to do?

ACTIVITY 6.7

As their advisor on assurance, you need to make sure that GANT don't fall foul of the Data Protection Act or the General Data Protection Regulation when exchanging information with their new partners. What do you suggest to them?

ACTIVITY 6.8

Thanks to an unexpected grant, GANT has acquired a video-conferencing system and you have been asked to link it into the network so that anyone can watch the participants of a meeting from their desk. What threats do you think you should protect against?

CLOUD COMPUTING

LEARNING OUTCOMES

The intention of this section is to provide you with the basic knowledge needed to understand the information security issues faced when utilising cloud computing facilities. Once completed, you should be aware of the issues and be able to identify approaches to reduce risk.

Introduction

Cloud computing is a generic term used to describe on-demand, off-site and location-independent computing services. There are a variety of ways that cloud computing can be delivered, and they generally fall into the categories of providing software services, platforms or infrastructure. They are accessed either via the internet or, in the case of larger organisations, through direct connections into a cloud provider's network.

Most people are already using cloud-based services in their personal lives, such as hosted email, photo sharing and social media; however, cloud computing is taking an increasingly prominent role within the workplace too. Organisations are eagerly taking advantage of cloud environments enabling them to rapidly implement technical solutions to meet business needs. For smaller organisations, cloud solutions can provide access to powerful computing tools that would previously have been out of their financial reach. These include such offerings as Microsoft's™ Office 365®, in which, rather than buying multiple copies of the MS Office suite of programs, individual users, small businesses and large organisations can subscribe to the online service, which includes an element of cloud storage and information sharing between teams.

In cloud computing there are a number of common terms, such as software as a service (SaaS), platform as a service (PaaS) and infrastructure as a service (IaaS), which are used to describe the types of service offered by the cloud provider. The terms public and private clouds are also used. In simplistic terms, public clouds are shared environments where the service provider makes resources such as applications and storage available to the general public over the internet. Private clouds describe environments where computing resources are used by only one organisation or, more commonly, where the organisation's information is completely isolated from other clients' – the term 'private cloud' is considered by some as a misnomer because of this. The term 'hybrid cloud' is sometimes used to describe where an organisation has some elements of their computing services within a private cloud, from which they can then access other resources held in public clouds.

Typically, a cloud supplier provides a service that is based on the public cloud model and utilises an infrastructure shared by many organisations and individuals, harnessing economies of scale to keep unit costs down and to enable higher levels of availability. To achieve this, information may be located in various facilities across a number of legal jurisdictions and be handled by a number of third-party service suppliers. Depending on the cloud environment being used, an organisation may not know precisely where

its information is kept or have few rights or control over what safeguards are in place to protect it.

Some cloud services cannot be customised and have to be taken as they come, while others can be tailored to meet organisational requirements and constraints. This is generally more applicable to the PaaS and IaaS models. The variety of cloud services available is extremely diverse, as are the cloud suppliers. These can range from multinational corporations to small start-up software companies. Therefore, the levels of control are variable, and it is vital to understand what is being offered and how it is being delivered.

Legal implication for cloud computing

It can be relatively easy for a business or end user to enter into a cloud services contract. For example, an end user can purchase or take up an application over the internet. By pressing the 'accept' button, they will be bound by the suppliers' terms and conditions (whether or not they have been read), so services can be obtained without the security implications being fully assessed. Essentially, when a business or end user signs up to a cloud service, the organisation has agreed to the terms and conditions and entered into a formal contract, which may limit the organisation's legal rights. This can have important implications later on.

Even in more formal contractual arrangements, it is essential that an organisation understands the cloud services they are using and the agreed contractual arrangements in place to control them. If not, the organisation may be in danger of breaching legislation, exposing confidential information and putting their intellectual property at risk.

For instance, information may be held by the cloud supplier in jurisdictions that are either undesirable or not legally permitted as specified by the legislation local to the organisation (for example in the case of data protection). Most countries have legislation controlling the location for the storage of PII and it is essential that information held in the cloud meets those legal and regulatory requirements.

The contract may give the cloud supplier important rights over the information held, including the right to use it commercially or to be able to disclose it to third parties. This could impact on the ownership and the value of the organisation's intellectual property or result in disclosure of personal information to unauthorised parties. Both scenarios could have far-reaching legal implications.

The contract may legally allow the cloud supplier to subcontract the delivery of part or all of the service onto other third-party organisations. Again, the controls and handling of the information by these additional third parties may not meet the organisation's requirements and put their information at risk.

There may be little or even no rights to audit the service being provided or to regain control of the information should the supplier go out of business or if the contract is breached or terminated. Some providers maintain the right to change their terms and conditions without prior consent, which may degrade an organisation's rights and control over the information being held.

Selecting a cloud supplier

A cloud service provider is a third-party supplier and good third-party security practices must be applied when engaging with them. Whether purchasing a PaaS, SaaS or IaaS, a risk assessment should be carried out to understand the potential implications to the organisation. The impacts associated with loss of confidentiality, integrity and availability equally apply to the cloud environment.

A security breach on a cloud-based service could result in commercial or reputational damage. The organisation must understand the financial and operational impacts if the cloud service is suddenly withdrawn or becomes unavailable, its information becomes compromised or is disclosed, and what protection measures are provided by the supplier. There may be compelling arguments against using a cloud-based solution for these reasons.

Some cloud providers do not provide adequate security arrangements and a common pitfall is that incorrect assumptions are made about the service and the assurance levels that will be provided. When choosing a cloud supplier, the organisation must ensure that the supplier can meet the organisation's security requirements and that they fully understand how the service will be delivered to them.

The organisation must consider all stages of the information life cycle and gain explicit assurances that key security issues are being addressed to an adequate level. For example:

- What safeguards are in place to prevent commercially sensitive information being disclosed to a competitor sharing the same platform?
- Will the organisation's information be used for any other purpose or disclosed to other organisations?
- Are committed service levels in place?
- How frequently is the data backed-up and to where?
- How is data transferred, stored and deleted?
- What levels of support will be provided?
- Where will the data be stored?
- What handling arrangements are in place?
- What are their infrastructure standards?
- What is the procedure for ending this contract and perhaps moving to a new supplier?

The levels of control should be proportionate to the risk to the organisation and the value of the information being held and must obviously meet legal and regulatory requirements.

Suppliers may restrict or not allow customers to audit or monitor the services being provided, making it difficult to implement effective governance processes to gauge how well the information is being protected. The supplier may claim rights over the

information held. These conditions might be unacceptable to an organisation, so it is important to understand these constraints before entering into a contract.

The usual business processes must be followed, including due diligence checks and references being taken up, to determine whether the supplier is reputable and stable. Finally, the service must be covered by a contract reviewed by a legal specialist. Clauses will vary from contract to contract, but the following areas should be covered as a minimum:

- the levels of privacy and confidentiality that will be applied to the data;
- any restrictions on the legal jurisdictions where the data can be held;
- any restrictions controlling the subcontracting to third parties for all or part of the service ('flow through clauses');
- agreed service levels to be provided and the penalties if they are not;
- the rights to review, audit or monitor the service;
- the process for dealing with any security breaches;
- how changes to the service will be controlled, including notifications and modifications;
- options to decline changes or modifications and the ability to terminate the service if they are unacceptable;
- service termination arrangements, which must include return of the organisation's information together with any indexing metadata and its subsequent destruction from the supplier's systems;
- the arrangements for the termination of the contract;
- supplier indemnity and liability levels and arrangements.

Comparing the risks of conventional and cloud-based solutions

With 'classical' organisations and architectures, the ownership and management stay firmly within the organisation. The organisation retains total 'end-to-end' control of the service, but it also retains the overheads and costs of selecting, implementing, maintaining, securing and upgrading the various components. Facilities and equipment will be purchased or leased; internal staff, and/or third parties need to be directly engaged and managed to deliver this model. There may be compelling business, industry sector or regulatory reasons that would advise against the organisation putting any information in a cloud environment – particularly business-critical and sensitive information. This may be the correct and sensible operating model from a cost and risk perspective for the organisation.

However, for other organisations, using cloud services for some activities may make economic and business sense and enable the organisation to achieve strategic goals. Two of the main drivers for implementing cloud solutions are the potential cost savings and speed of implementation. Cloud suppliers can respond quickly as requirements change, such as providing additional functionality or capacity. Economies of scale can

be harnessed and, depending on the services being purchased, the management of environments, hardware and software can all be handled by the cloud supplier. These benefits may be extremely attractive to some organisations, especially smaller ones where it fits their risk appetite.

Distinguishing between supplier commercial risk and purchaser risk

The key risks to the provider of cloud services will largely be commercial. If they fail to deliver the contractually agreed service, their commercial model will fail and serious commercial issues ensue. They can mitigate some of these by ensuring there is an appropriate degree of resilience in the systems providing the service and that they take all necessary and appropriate precautions to guard against failures. However, if their systems are breached in some way, allowing unauthorised access to, for example, personal information, while they could be embarrassed, they are unlikely to suffer the same consequences as the owners of the data. The purchasers of the service will have to deal with the customers whose data they have failed to look after and it is not likely that those customers will be satisfied by statements along the lines of: 'It wasn't us who failed – it was our supplier.'

As discussed in Chapter 2 on information risk management, this is, in some ways, an example of risk sharing – the organisation may choose to outsource its data to a cloud supplier, but it must always retain responsibility for ensuring that the data are properly protected.

The major systems failures of recent years in many countries have mainly left the owners of the data with a more serious reputational problem than the suppliers whose systems failed. The risk assessment and the business impact analysis for the purchasing of such services must be undertaken with as much rigour (if not more) as if the service was being provided in-house.

The risks to a purchaser of such cloud computing environments are generally associated with lack of control over information and incorrect assumptions about the service and the safeguards being provided by the cloud supplier. When entering into a relationship with a cloud supplier it is essential that there is a clear understanding as to what will be provided and how it will be delivered, as outlined previously.

An organisation can be particularly put at risk by the unauthorised purchase of cloud services. As it is relatively easy to buy and implement certain cloud services, end users may begin to access services without management approval and without appropriate security controls being considered. A clear and well-communicated policy on the use of cloud services should be put in place and procurement policies and processes should ensure that information security management approval is obtained prior to purchase. This should be supported by monitoring facilities within the organisation to identify any services that have been purchased without authorisation.

Whether using a cloud or a classical architecture, overall it is critical to understand that ownership of risk to the data still remains with the organisation. Even in a 'classical' structure there is still likely to be reliance on third-party services, where the corresponding risks need to be understood and managed to an acceptable level. If using

a cloud service, the nature and levels of both commercial and operational risk will vary depending on the type of services taken up. PaaS and IaaS offerings can provide more control, with the private cloud providing the greatest. However, as services become more customised and dedicated, the costs will generally rise, and this also needs to be taken into account.

Advice and guidance on managing risk within the cloud is becoming more available as cloud computing matures. There are several organisations that provide standards and guidelines to help improve risk management of cloud services. These include NIST; the Cloud Security Alliance (CSA); the ENISA; and the ISF.

GANT's Ms Jackson has been approached by an organisation that can provide a web-based software package that would help to manage membership administration and their payments and also provide a platform for members to share information. It appears to be extremely cost-effective and will reduce internal administration, management and maintenance overheads. Apparently, it is a cloud-based service and would require minimal tailoring to meet GANT's needs. She has asked you to join her in discussion with the vendor.

ACTIVITY 6.9

What key information assurance issues would you highlight to Ms Jackson that would need to be considered before meeting with the vendor?

ACTIVITY 6.10

What information assurance issues would you raise during the meeting with the vendor?

IT INFRASTRUCTURE

LEARNING OUTCOMES

The intention of this section is to provide you with an understanding of the security issues surrounding security of IT Infrastructure and the content of associated documentation.

Separation of systems to reduce risk

A simple, yet very effective, way to manage risk and provide assurance is to keep systems separate. Although there are advantages to joined-up systems that share data, it is not always necessary. In some cases, it may be decided that the risks outweigh the advantages and that it should not be done. An alternative is to allow very limited functionality to pass between systems through an inter-domain connector (IDC) or to allow data to pass only one way, through some form of data diode or specially configured router.

Another advantage is that separate systems are less complex to manage and easier to assess for risk because of the reduced complexity. Increased complexity usually means more cost to implement and support the IT infrastructure. If the functionality cannot be shown to provide a positive business benefit, why do it?

Conformance with security policy, standards and guidelines

There is no point in having standards for the design, implementation and operation of the systems if they are not followed. Having said that, if the procedures are not aligned with the processes and requirements of the business, the staff will not follow them. The same is true of the security policy, standards and guidelines. They have to be aligned to the operational needs of the business – for day-to-day operations and for effective business continuity and disaster recovery. This is a complex subject and one that may need expert advice to get it right.

Accreditation to ISO/IEC 27001 will require that all the relevant controls have been identified, documented, implemented and then followed. Regular internal and external audits will be needed to confirm this. The hierarchy of these documents is as follows:

- The policy defines the overall information assurance goals of the organisation and must be supported by the board and chief executive to provide authority.
- The standards define the minimum acceptable criteria for achieving that policy in the key areas (e.g. the control groupings in ISO/IEC 27002).
- The guidelines advise how to design and implement workable procedures and countermeasures to meet the standards and enable the business to manage risk.

Access control lists and roles, and control of privileged access

The concept of ACLs and roles has been mentioned earlier in this book. The learning point here is that there is no point having such controls if access to the ability to update or change those controls is not also protected. An attacker who finds that they cannot access certain material may well turn their attention to finding a way to subvert those controls. The obvious place to start is by seeing if they can grant themselves the necessary privileges, perhaps by creating a new account for themselves about which the system administrator may know nothing. Many organisations use extra safeguards for the accounts that can grant these sorts of privilege. There will probably only be

one or two accounts with these rights and they often have longer passwords (e.g. 12 characters long instead of 9) to make them even harder for an attacker to try and break.

Principles and requirements for correctness of input and accuracy of stored data

There is no point in having the most secure system in the world if the data it contains is inaccurate and of no use. While attackers and malware can subvert stored data, the most common cause of incorrect data is either incorrect user input or errors in software design or coding. One of the main controls in the UK's DPA, which is mirrored throughout the European Union, is a requirement for PII data to be accurate. This means it is not just good business practice, it is also the law.

There are several ways to promote data accuracy and they all need to be used in conjunction with each other.

- Make sure the design of the software and database is correct, so that values reflect the right information, and relationships have the right meanings.
- Use proven code review techniques when developing and testing the application before it goes live.
- Use defensive coding, which checks for values within acceptable ranges and looks for correct relationships with other fields before accepting an update command to change the database.
- Train the users in how to use the application properly. Make sure they understand the meanings of the fields and their relationships.
- Audit the system regularly to look for anomalies. Automated tools can help this process.
- Have a means whereby it is easy for those who have data in the system to report any errors and have them resolved as soon as possible. Try to identify how the errors occurred and then how to stop them happening again.

Principles of recovery capability, including backup and audit trails

Having produced a means of creating and holding data securely, it is vital to be able to recover it should anything go wrong. Problems can range from the theft of a laptop or the failure of a hard disk through to the entire building being lost to fire. It has been shown many times that any organisation that loses access to its data for more than 10 days is very likely to go out of business. In some cases, more than 48 hours is enough to signal the end, or at least invoke severe financial penalties and major loss of goodwill.

It is absolutely essential that backups exist for all data, and not just a current backup. Use of an approach such as the GFS approach (maintaining at least three generations of the backed-up data) to allow recovery of data back to a previous point in time is highly desirable. There have been occasions when organisations have discovered a piece of malware that has been present for months, quietly changing data values at random. The only way to resolve the issue is to roll back the system to a point in time before the

malware was present and rebuild the data from paper records. Without a GFS backup approach, this is not possible.

It is also important to consider how the integrity of the restored data can be checked. It is all very well to bring back a database of information, but if there has been some corruption of that data, from whatever cause, deliberate or accidental, there is a lot of work to do to clean the data up before operational use is possible again.

The cause of the risk may be outside the organisation's control, and they may not even be able to use their own premises (think of 9/11 in the USA in 2001 or Buncefield in the UK in 2005), but they still have to be able to recover data and operations. It might be useful to consider a disaster recovery contract or having the ability to relocate the data centre to another company site in times of emergency. Whatever is done, a copy of the backup must be kept in a secure location off site. More detail on this aspect of information assurance is provided in Chapter 8.

An audit trail has four main uses:

- To understand the current status – what is complete and what transactions need to be rolled back or re-entered?
- To identify what happened and who did it.
- For compliance with standards and legislation and demonstration of due diligence.
- As a deterrent against internal attack.

Collecting and keeping transaction and event logs is often referred to as protective monitoring, because it is a means of doing all of the above tasks. Treat the logs in the same way as all data backups. With the right tools and training, the audit data can provide powerful insights into what is going on.

Principles of intrusion monitoring and detection methods

Intrusion detection systems and intrusion prevention systems (IDSs and IPSs) use automated tools to analyse log data, system activity and network traffic in an attempt to identify and, in IPS, to block unauthorised users or malware from causing a security breach. There is so much log data and system activity, especially in large systems, that it is impossible for any one person to monitor it all in real time, and uneconomical for any organisation to pay sufficient numbers of people with the right skills to do the work. The only practical solution is automation.

These systems can capture data from network traffic and devices such as routers and firewalls, network intrusion detection systems (NIDSs) and from system hosts, host intrusion detection systems (HIDSs). The analysis is done either by an application or a hardware device, many of them using statistical techniques and a tool such as Snort to analyse the data, looking for changes in system configuration or operation, or for known types of behaviour often referred to as a signature.

The problem with these systems is the number of false positives that they often return, especially when first installed. It takes a skilled user to configure them correctly and to

educate the system to understand what is, and is not, normal activity. The IPS solutions cause the most problems, because they tend to stop authorised users from working when they block a false positive.

Installation of baseline controls to secure systems and applications, and the dangers of default settings

Baseline controls are standards used to define how systems should be configured and managed. The intention is that any new system in any location should be built using the settings and guidelines contained in this document. In this case, the concern is about configurations for information security.

The contents will include details on:

- which versions of operating systems to use;
- which parts of the operating system to install;
- the patches required;
- additional applications such as anti-virus software, intrusion detection agents and so on;
- settings for password length, ACLs and so on;
- network configuration.

Baselines are a good start to implementing security, but there is no 'one size fits all' answer. The configuration for an email server is different from that for a web server, which in turn differs from those for file and print servers. A further danger is that people assume that a machine built to the baseline is secure. The problem with this is that new vulnerabilities and exploits are discovered all the time and new patches are issued for the operating system and applications. The assurance provided by the baseline does not last beyond the first new patch and so the baselines must be continually reviewed and updated.

Part of the initial default installation process for many software applications and hardware devices is a default password. These are configured on the basis that a password of some sort is better than no password at all. Unfortunately, there are many sites on the internet that list literally hundreds of default passwords. Once an attacker identifies the infrastructure in use, they can try the default passwords, which will often give them administrative privileges and provide an excellent basis for an attack. It is very important that all default passwords are changed as soon as the installation is complete. Since they are for administrative use and provide significant administrative rights to the user, that password needs to be longer and stronger than ordinary user passwords, making it much harder to break.

Worse still are hard-coded passwords, which some hardware manufacturers and some software developers insert, thinking that they will only ever be used *in extremis* and not understanding that, once their presence becomes known, they represent a major threat to any organisation that has deployed that software or hardware.

Configuration management and operational change control

The topic of configuration management follows on logically from that of baseline controls. It is the process of monitoring and controlling the configuration of devices and documentation within the infrastructure. The configuration documentation should describe the baseline that is in place and it can then be used to identify any changes made.

Change control management requires the effective process of configuration management as an essential element. The documentation can be used to help assess the requirements for changes and the impacts these changes may have before granting approval for the change. It is important that the documentation is kept up to date to reflect any changes made. The documentation can also be used as part of the auditing process, for quality, assurance and operational purposes.

The protection and promotion of security documentation

If an organisation has any links to third parties or external suppliers, such as managed service providers or outsourced operations, it is very important that they are required to work to the same information assurance standards and adopt the same working practices, or at least to those that are clearly compatible. If they do not, they may become the weakest link in the chain and can invalidate much of the good work done in-house. The use of working protocol documents and contractual clauses can require them to do so, and should allow auditing to ensure compliance. It is becoming more common to see third parties required to have an accreditation such as ISO/IEC 27001 before they can work for an organisation. This provides a degree of confidence in their assurance, including the quality and content of their documentation.

Having produced a set of security documents, it is most important that they are protected against unauthorised access and loss. They may be physical, electronic or both, and all must be safeguarded. The contents of these documents describe how the countermeasures and procedures in place work to protect the assets of the organisation. Knowledge of the content would make life much easier for an attacker to find a vulnerability in the infrastructure and gain access. Access to the documents, physical or logical, must be very strictly controlled and monitored to prevent abuse. It is often worth considering the introduction of a protective marking system to allow such documents to receive extra protection and safe handling.

GANT continues to grow and now has more IT infrastructure than can reasonably be supported in-house. The economics do not justify employing the necessary specialists to manage this, yet the skills are required to be available when necessary. The time has come to issue an invitation to tender (ITT) to third-party suppliers of IT support and other services to provide managed services and IT support to the organisation.

ACTIVITY 6.11

You have been tasked with ensuring that the ITT documentation contains the necessary statement of requirements for information assurance and professional standards of work. What would you include?

ACTIVITY 6.12

The members of the board are aware that they ought to have a formally documented information assurance policy and supporting documentation, but they are not clear on the structure that it should take. How would you explain to them the purpose of each kind of document and their hierarchy?

ACTIVITY 6.13

Another requirement to be included in the third-party ITT is for the baseline builds for the systems to be implemented and supported as part of the contract. What requirements would you include for the builds, documentation and change control?

SAMPLE QUESTIONS

1. **Encryption is used for several reasons. Which of the following is *NOT* a reason for using encryption?**
   a. Confidentiality.
   b. Availability.
   c. Non-repudiation.
   d. Integrity.

2. **A firewall is designed to do what?**
   a. Stop unauthorised traffic into an organisation's networks.
   b. Stop unauthorised traffic out of an organisation's networks.
   c. Stop all traffic in and out of the organisation's networks.
   d. Stop unauthorised traffic into and out of the organisation's networks.

3. **Which of the following is one of the main uses of an audit trail when reviewing a backup procedure?**
   a. To help to identify what happened and who did it.
   b. To measure how long a system has been operating before an incident.
   c. To help to determine if authorised software is in use.
   d. To determine if there are any unauthorised internet connections.

4. **If an organisation decides to store all its personally identifiable information (PII) in a cloud storage facility, under GDPR who has the overall legal responsibility for its security?**
   a. The organisation whose information it is.
   b. The organisation providing the cloud storage service.
   c. Both the owners of the information and the cloud storage provider.
   d. The provider of the software used to store the information in the cloud.

5. **Which of the following is *NOT* a type of malware?**
   a. Trojan.
   b. Worm.
   c. Bug.
   d. Rootkit.

7 PHYSICAL AND ENVIRONMENTAL SECURITY

Information security managers need to have a good appreciation of associated physical security issues and the controls that they might use to make sure there is a seamless information security management system across the whole organisation.

As mentioned in previous chapters, there are three main types of operational control:

- Procedural – for example checking references for job applicants.
- Product/technical – for example passwords or encryption.
- Physical – for example locks on doors and secure cabinets.

Physical and environmental controls are often the ones most overlooked and yet can be the most cost-effective: just physically stopping people from getting into a room can remove, or seriously reduce, the need for additional technical controls to restrict access to sensitive systems or information.

LEARNING OUTCOMES

Following study in this area, you should have a sound understanding of the environmental risks to information in terms of the need, for example, for appropriate power supplies, protection from natural risks (fire, flood, etc.) and in the everyday operations of an organisation. You should also understand the operational controls that physical and environmental security can offer to enhance overall information security and assurance.

PHYSICAL SECURITY

As discussed in Chapter 5, physical security relies on the presence or otherwise of physical limitations to the activities that a criminal or other unauthorised person might wish to carry out.

It is usual for physical security to be the first line of defence in many organisations. Stopping people from entering the building, or at least having a reception area where they are 'detained' until suitable authority or other arrangements are made, is the usual first point of security. However, these are only useful and successful if the reception is properly manned and it acts as a 'guard' for the only entrance into the building. Fire

escapes, back doors left open for those escaping the office for a breath of fresh air and other similar entrances must also be suitably protected. For the more adventurous, climbing walls and fences to gain access to roof lights, upper-floor windows or stairs in a shared building is also a source of security risk that must be assessed and managed.

On a smaller scale, once inside a building the use of locks for offices, server rooms and other sensitive areas is again commonplace. As are locks on filing cabinets, desks and other document storage facilities, and they are entirely suitable as one layer of security whilst accepting that they will not deter the most determined criminal for long. It should go without saying that the keys to such locks also need to be protected and not left in a glass on the desk!

Equipment needs to be appropriately protected and often this leads to reinforcements of the usual building materials: impenetrable layers within walls and cages around equipment, including above ceilings and below floors. In high risk and high security situations, this may even include crash-proof barriers outside the server locations to prevent vehicles being driven into the building and destroying the equipment.

Electronic door locks, swipe card readers and the like can be included as technical controls to protect and enhance physical security. These technical devices are more secure than their basic mechanical counterparts, but may still suffer from similar problems. Using a fire extinguisher to keep a door open has been seen many time over the years and tailgating – the art of following a legitimate entrant into a building very closely before the door has had time to close again – is a well-known issue. The provision of individual 'pods' with automated doors, allowing only one person in at a time, can counteract this problem. They are obviously a more expensive option, but may be a necessary approach.

Once inside, if a visitor is left to wander the building unaccompanied, this then presents further security issues. Monitoring and detection tools such as motion detectors, CCTV cameras and intruder alarms can be installed to alert and record that the physical security has been compromised.

Protection of equipment

To take the idea of different types of control a little further, we might consider the threat of equipment being stolen or taken off site by unauthorised staff. Another possible scenario is the loss of use of some critical equipment, such as filtration or air-conditioning units becoming ineffective through fault, power failure or other cause.

If it is clear from the business impact analysis that the consequences of these problems would be significant to the organisation in some way, perhaps loss of revenue, loss of credibility, reputational losses or equivalent, then it may be decided that there should be some measures taken to reduce the impact or perhaps even remove the cause of such a problem. This might mean marking all equipment in some way to lessen its attractiveness to potential theft from staff or outsiders. It might be deemed appropriate to have a stand-by

power generator connected through a UPS so that if the mains power fails then the generator automatically cuts in to maintain the supply; this is frequently used for air-traffic-control systems, hospital life-support systems and critical financial and manufacturing systems. Simply having appropriate maintenance contracts and service level agreements, with realistic timings for the provision of an engineer, are further ways to enhance the security of the assets with which there may be concern.

Provision of alternative power or data supplies is an approach that can achieve the necessary level of service provision, but it will come as no surprise that these come at a price. Sometimes that price will be too great, and the organisation may have to accept a lower level of support or a longer period of downtime before a remedy is in place.

Realistic assessment of the risk and its impact on the organisation is the key. As an example, when one of the authors was undertaking a full risk assessment for a public body, each section was asked to define how long they could continue to operate effectively if all the computers were unavailable. Some sections mentioned minutes, up to a couple of hours, while others could survive perhaps a day. In the personnel department, though, they were quite happy to work on for up to a month because much of what they were doing at that time was paper based. It wasn't until the staff needed paying that the computer systems were critical and so this meant the acceptable downtime could be as long as several weeks, but might only be a couple of days if it was coming up to pay day.

How critical elements of any system can be protected is very much determined by physical and technical options. The main power supply to a building might be the only one. If an over-enthusiastic road worker digs up the tarmac outside and breaks the cable with his excavator when putting in a new drain, it might take some days for that to be repaired. Would it be appropriate for the organisation housed in the building to go to the expense of putting in a secondary backup supply, either in the form of a second cable or a stand-by generator? Probably not in the case of GANT, but it certainly would be if the organisation was a high-street bank with millions of transactions to be dealt with each day, even if this provision was only for the duration of the roadworks outside.

Inside a building, the protection of cables and similar equipment can be equally important. It is not unknown for an ordinary mains electricity socket in the passageways of buildings to be used for network equipment. If the cleaner decides to use that socket to plug in a vacuum cleaner, significant problems could arise. Some companies wire their buildings with two or more separate circuits: one for critical electrical equipment like computers; one for other electrical equipment that can be unavailable during an electrical supply outage without serious impact on the business; one for normal lighting circuit, desk-lamps and the like; one for emergency lighting in case of fire, power outage and the like. To do this is not cheap, but will provide a high level of availability if the circuits are properly used. Often, different coloured or different types of sockets are used to prevent the wrong things being plugged into each circuit.

The security of network cabling is another area of concern. In high security environments it is often a requirement that all such cabling (including fibre connectivity) be housed in transparent ducts to ensure that no additional extraneous connections are put in by some unscrupulous person. With fibre this is much more difficult to achieve than with the more traditional copper, but it is not impossible. Various means of checking the data

being transferred along a cable can be used in cases where there are no suitable ducts available, but these can also be expensive and difficult to manage.

ACTIVITY 7.1

Looking at the records and information held by GANT, suggest how long these sections of the group might be able to continue operating effectively if they lost their main computer system, which contains all their details.

1. membership secretary;
2. general enquiries on the Natterjack toad breeding ground details;
3. forthcoming planning application where there was interest in the toads;
4. financial information.

Suggest factors that might affect the timescale.

Clear screen and desk policy

With so much data processing now being completed on computer systems, it is inevitable that there will be times when computers are left unattended with sensitive information displayed on the screen. It may be for a short period while the staff member collects a cup of coffee, but should they then get called away or the fire alarm goes off, the information might be displayed for some time. With the trend for open-plan offices, where the screens of fellow workers are easily visible to anyone wandering by, it becomes all the more important to protect that information. It may be reasonable to assume that all those walking the floors of such offices are legitimate members of staff with not only a right to be there but suitable authority to see such sensitive information. In other circumstances, though, it may be less certain who those wandering by might be.

The use of third-party companies to carry out routine tasks, often some of the more menial, including cleaning, presents further issues to the security manager. The supplying company may use staff who have little security clearance and yet they may be free to enter all offices or working areas more or less unsupervised and, while there, can gain access to the unlocked filing cabinet, logged-on computers and papers left lying around or in rubbish bins.

In one UK government department, due to the sensitivity of most of the information they process, the decision was made to prevent anyone who was not suitably authorised to even enter the building. Instead, they established a visitors' centre where guests could be hosted and meetings take place outside the confines of the main building. However, even in this climate of protection for their information, a clear screen policy is enforced to avoid those without a need to know getting to see information they shouldn't.

There has to be an acceptable compromise here. If the time default of the clear screen is set too short it becomes a real annoyance to workers who continually have to enter

their password to re-access their computers, while if it is set too long, too much damage could be done by inappropriate people seeing sensitive information. Nevertheless, a suitable time for the screen to be cleared by a standard screen saver system should be determined and, on almost all occasions, employed on every computer without allowing the users the option of turning it off. However, it may simply be better to ensure all users lock their keyboards and invoke the screen saver whenever they leave their desk, relying on the automated system as a fallback system only.

'If a cluttered desk is the sign of a cluttered mind, what is the significance of a clean desk?' goes the quotation from Laurence J. Peter, a US educator and writer who died in 1988. A tidy or clean desk might be more pleasing to the eye, easier to work at and have all sorts of other benefits, but, most importantly, it allows those items that are sensitive or valuable in some way to be properly looked after. If a desk is cluttered, sensitive documents may well become covered by other papers. Then, when the owner of the desk takes a quick look before leaving the desk for a while, they don't see the potential security breach hidden on the desk.

This is all the more important at the end of a day's work. It is then that a clear desk is much easier to check for important documents that should be locked away. This does include documents that might not be regarded as sensitive in many circumstances; perhaps a directory of all the staff is left out as it is considered not to be very sensitive. However, if the policy of the organisation is to use something akin to an individual's name as their logon name for the computer systems, as can be very common, this directory becomes very helpful to a potential hacker. This could be the contract cleaner who has time to try their hand at getting into the system while ostensibly cleaning up each night. All they then have to do is watch the unwary user typing in their password when they are working late one night, and the potential hacker has all they need to hack into the system and do whatever takes their fancy!

ACTIVITY 7.2

The study in which the GANT secretary works is as cluttered as anyone has seen. There are piles of paper and books everywhere with filing cabinets left open and windows unlocked. She argues that no one would be interested in her study and, anyway, if she can't find anything how would anyone else?

Suggest three reasons, with some justification, as to why she should consider the implications of a clear desk policy.

Suggest why other security measures might be appropriate and how she might achieve a secure working environment.

Moving property on and off site

The control of an organisation's property both on and off site is another area of critical interest to those concerned with assurance. Apocryphal stories abound, along with the

genuine ones, of: laptops and phones left here and there; phone camera pictures of new models of car being released to the press before the official launch; CDs being lost in the post; memory sticks left on the train; classified rubbish being found at the roadside; a salesman being sacked and taking the company's database of clients with him; and so on. In the past, when the equipment itself had more value, there were stories about thieves backing lorries into offices and removing all the computers, or all of the memory chips being stolen out of office computer systems by cleaners. What is considered valuable changes over time, but valuable stock items, whether they be jewels or the latest smartphone, have been the target of crime almost since the beginning of the consumer society, and criminals will adapt their methods as the world changes.

The way in which property is securely moved around is very dependent on the nature of the property and its value, both intrinsically and to business operations. It might be a very cheap item, but if it is lost or stolen and it takes a while to find a replacement and have it delivered from the other side of the world, its loss could be critical to the continued business operations of the organisation.

There are some fairly obvious ways of reducing such risks. First, a good start is marking all assets with an indelible mark that uniquely identifies it. This then allows a full inventory of all assets to be taken and maintained. As equipment is exchanged for newer items, then the register must be maintained. Only allowing those with the appropriate authority and skills to move equipment around is another useful control. This is especially pertinent to technical equipment such as computers or sophisticated printer-copiers. These often need specialist skills or resources in order to set them to work effectively after a move and anyway they may require additional work such as extra power points or data connections to make the move possible. When such an item is moved, it should be automatic that the asset register is updated with the new location so that the next time the service engineer arrives, he can be directed immediately to the right office.

Taking equipment off site should also be controlled. The idea of a staff member being able to remove from an organisation's site a laptop containing all the details of the customers of that organisation without any control, suggests that the organisation is not too concerned about their customers and that it may be out of business before too long. Clearly, laptops, tablets, smartphones and the like are essential tools for the mobile staff member, but the control of the information they contain and to which they might give access is also very important. Limiting the facility to take copies of company databases, controlling access to company intranets when accessing the network from off site and other similar controls can be very useful.

Laying down clear procedures, to which staff members must sign up, concerning how they use equipment, where it must be stored when not in use and so on is another area where security and good practice coincide. Outside the scope of this book are the health and safety aspects of the use of company or personal equipment at home and elsewhere too, and to wrap all these up in one well-drafted policy for the use of such equipment can be very effective.

This needs to extend to all equipment and access to any asset, information or otherwise of the organisation. Rules governing the use of mobile phones are now fairly widespread within many organisations, both private and public sector, controlling their use inside

buildings. Some organisations reserve the right for all mobile phones to be surrendered before entering buildings, especially where commercial or very sensitive information might be visible to the casual visitor. For instance, it was the lack of such a control that allowed the pictures of a new car to find their way onto the internet long before the official launch of the vehicle.

There is an ever-increasing need for organisations to be aware of the desire and often the strong business need for individuals to use their own devices for work purposes, what has become known as 'bring your own device' or BYOD. Notwithstanding the legal issues that this immediately entails, the practical issues, such as the security of the devices, the prevention of spreading malware from these devices to the corporate environment and vice versa, must be considered and managed. There is a clear business benefit of reducing the organisation's overheads by not having to buy all the separate devices, but the other side of the cost/benefit analysis has to consider the extra security measures that will be required to be implemented and maintained.

Overall, the world of information is now a very mobile one, where people expect to be able to gain effective access to almost any information relevant to their work or daily lives at any time, anywhere and on a variety of media types. While technology has rapidly allowed this to happen, it is less clear how well the assurance world is keeping pace. While significant advances in security have been evident in recent years, there is still a long way to go. Dealing with the compromise that is inevitable between full and effective assurance and the availability of information when and where it might be needed is still a difficult judgement call.

Procedures for secure disposal

The disposal of equipment or other information assets that are no longer required has often been another source of good stories: the confidential files being found in second-hand filing cabinets; the valuable or sensitive files found on hard disks in computers sent for disposal that found their way into the second-hand market; the classified waste bags found in open rubbish tips, and so on. Having procedures that ensure that any filing cabinet, desk drawer or other container is properly checked by two competent people before it is allowed off site is a good start, but there are more issues to consider.

Policies and procedures for the secure disposal of any piece of equipment or other asset, including wastepaper, are crucial. Simply writing a policy is not really enough. A study by the Blancco Technology Group in 2019[1] found files containing PII on 15 per cent of the hard drives purchased through the online retailer eBay. Of the 159 hard drives purchased in the USA, UK, Germany and Finland, more than 40 per cent contained sensitive data. It is critical, therefore, that those disposing of equipment that might contain sensitive information must be careful to ensure that they or their contractors clean the devices effectively.

Electronic media are a particular problem. It is often assumed that simply pressing the delete key on a computer will remove the information completely from the system, but that is unfortunately very far from the truth. The way data are stored and used on a

1 https://www.blancco.com/resources/rs-privacy-for-sale-data-security-risks-in-the-second-hand-it-asset-marketplace/

computer means that destroying them is not like taking a paper file out of the drawer and shredding its contents in a decent cross-cut shredder. That will make reconstitution of the information on paper virtually impossible, but with a computer the data are retained long after the delete key is pressed. There are a number of ways to completely delete the information, including writing random data to the same data store a number of times or physically destroying the media itself by cutting-up or shredding hard drives, for example, but most need some special technical equipment or knowledge. Thus, it is again important that the security professional is ready and able to call in specialist advice and guidance when necessary to ensure appropriate measures are taken.

Consider the contractor engaged to remove classified waste from an office who brought a very smart, well-secured vehicle to collect the waste documents but, instead of taking them for secure destruction as expected, the driver simply dumped the sacks of waste in open rubbish skips where anyone could open them and take the contents. It is the responsibility of the organisation whose information assets have been sent for disposal to follow a contract through to ensure all is well and that the contractor is actually doing as expected.

Reputable organisations that specialise in the destruction of obsolete paper, media and computing equipment will provide certificates of destruction. They will be able to provide clear explanations about the various levels of crunching they deploy to destroy the information held. Some of these facilities can be mobile and brought to your location so that the destruction can be witnessed for yourself.

Security requirements in delivery and loading areas

Delivery and loading areas are often remote from the main buildings of organisations. This brings with it additional security issues that must be addressed. Many of the biggest raids in the UK have been at least initiated by attacking the staff in the receipt or dispatch area of a warehouse or factory. It is common for those working in such areas to have lower security clearances than other members of staff, reflecting their lack of access to critical business information, but this may be short-sighted. If those responsible for the receipt and dispatch of goods are regarded as the last check before sending items for disposal, then they need to be cleared appropriately, since they will potentially have access to a great deal of business-critical information. If the inbound goods dealt with are assets of the company, new computers for example, then they need as much protection in the receipt and dispatch areas as they receive in all other parts of the organisation. Simply the sight of heaps of polystyrene, plastic and cardboard can indicate the arrival of a new batch of desktops, which could encourage thieves to take a closer look.

Indeed, it might be suggested that they need even more protection since it would be easy for an unscrupulous worker receiving such items to misappropriate them, declaring they had not been received, were damaged on receipt or use some other explanation for their absence. In a similar way, if old items of IT equipment are to be dispatched for secure disposal via a contractor, it would not be difficult for a devious worker to create a means by which some of those goods never reach the appropriate destination but are sent off for sale on the internet or local car boot sale. It is generally believed that this is a fairly common example of company information being released inappropriately.

DIFFERENT USES OF CONTROLS

For each of the three types of operational control mentioned at the beginning of this chapter, there may be one of three uses: preventative action, detective action or reactive action. These three uses, referred to as tactical controls, can apply equally to each of the three types in so far as it is possible to use the different controls in different ways.

Physical security is mainly intended for use as a preventative control – to stop unauthorised people getting into a building for example. However, it is also possible to use such controls as detective controls – intruder alarms for example – and, although perhaps not so desirable, to use them as reactive controls, such as electrified fences, so that a potential intruder is 'rewarded' for their trouble by being detected, arrested and therefore prevented perhaps from trying it again.

Similarly, technical controls can be used in three ways, and anti-virus software is an example where in fact the same system can be used in all three. The software tries to prevent any malware being loaded on a computer system (preventative), will routinely run checks to ensure there is none installed (detective) and then provides a system for virus removal should something get through (reactive).

An example of procedural controls might be a non-disclosure agreement to protect the intellectual property rights of an organisation. This is designed to prevent unauthorised disclosure by warning of the consequences if such an event takes place and advising staff that they should not release such information without authority (preventative). It might include details of the measures taken to prevent such disclosure, including numbering of copies, limited and controlled distribution on signature, disabling the copying of large database records of clients and so on (both preventative and detective). It then might include details of the possible consequences that can be used in the event of a breach – dismissal, legal action or similar (reactive).

Clearly, there needs to be a combination of controls to provide an effective coverage. The tools used to prevent incidents can and do fail and so a 'backup' system to detect such events and the potential consequences of such a breach must also be in place. It is no use, for example, having an intruder detection system installed if the system is not monitored routinely and effectively at all times. The five minutes when the security guard pops out to make a cup of tea is inevitably going to be the moment when an intruder makes their entrance.

There is also a link between the differing types of controls and their uses. The procedures may define the technological protection measures that are used behind some physical barrier. All these not only need to be present but also to be consistent and appropriate. It is pointless protecting information assets with very expensive disaster recovery plans, including hot backup sites for example, if the information is not critical to the organisation and work can continue more or less as normal for a few days without great inconvenience. Once again, this brings us back to the business impact analysis, risk assessment and cost–benefit analysis to determine what is appropriate and necessary in any given situation.

When staff leave the employment of an organisation, how they are looked after during the interval between being told they are no longer required and actually leaving the

buildings can also be a tricky period. If the separation is not voluntary, in other words the staff member has been told to leave, there is every likelihood that they might try to leave with either some valuable asset of the company or after doing something to some asset of the company.

Once again there are stories of disgruntled employees putting malware onto computer systems that runs after the employee has departed, of deleting or copying important files in the interval between being sacked and actually leaving, and so on. It is now common practice, and often laid down in the personnel department's procedures, that such soon to be ex-employees should be escorted straight off the premises and be told there that their personal belongings will be forwarded to them shortly. The removal of their access to all systems, including doors and computers, the changing of passwords or codes to which they have had access, all has to happen in a similar time frame if the security of the organisation is not to be compromised. It is always worth considering how you would do this if the person to be sacked is your senior, or the only system administrator.

SAMPLE QUESTIONS

1. **If security guards are used to patrol the perimeter of an establishment, what sort of control could they be considered to be?**
    a. Technical control.
    b. Preventative control.
    c. Procedural control.
    d. Adaptive control.

2. **If there is a remote loading and unloading area for goods supplied to or by an organisation, from a security perspective, why might it be bad practice to have a single person working there?**
    a. They could be lonely.
    b. They would not be able to take a break.
    c. They could be tempted to steal goods.
    d. They would be a single point of failure.

3. **When disposing of unrequired filing cabinets who is the best person to check the cabinets before final disposal?**
    a. The last user of the cabinet.
    b. Anyone.
    c. It is not necessary to check.
    d. An independent person.

4. **A clear desk policy would be regarded as which type of security measure?**
   a. Procedural.
   b. Technical.
   c. Physical.
   d. None of these three.

5. **Which of the following describes an appropriate manner for the effective deletion of information from a computer system?**
   a. Pressing the delete key several times.
   b. Checking to ensure the directory no longer contains an entry for the files.
   c. Writing random data to the same data file for at least seven cycles.
   d. There is no effective way of ensuring effective data deletion from computer systems.

8 DISASTER RECOVERY AND BUSINESS CONTINUITY MANAGEMENT

Even in the best-prepared organisations, problems will arise. Hopefully, these problems will have been anticipated in some way and preparations will have been made to deal with them. This chapter looks at what the security manager needs to understand in order to deal effectively with the inevitable problems that can, and most likely will, arise.

LEARNING OUTCOMES

Following study in this area, you should be able to define and explain the terms disaster recovery and business continuity management and be able to describe their appropriate use.

RELATIONSHIP BETWEEN DR/BCP, RISK ASSESSMENT AND IMPACT ANALYSIS

A business continuity plan (BCP) is, as the name suggests, all about maintaining the continuity of business operations. Problems will always occur in any organisation, no matter how well run it might be, and these problems will adversely affect the operational capability in some way. It might be something seemingly very simple, such as the main printer running out of toner, through to something much more serious such as a power outage. In either case the situation will be made much more acceptable to both users and senior managers if someone has thought through what might happen and put in place some tried and tested plans to deal with it. Maintaining normal operations as effectively as possible while resolving the issue is the approach that must be taken.

There is a second issue. In an ideal world, if concern over the security of information is very high, the way to achieve the 'ultimate secure environment' is to lock all information in a large safe and lock the door of the room holding it. This would have the effect of ensuring that no one would be able to access the information inappropriately, but would also have the seemingly disastrous result of making the information virtually impossible to use. The solution is therefore some form of compromise, which, as mentioned earlier, is what information assurance is all about. The availability of information is

critical, and business continuity planning is part of the mechanism to allow operations to continue, come what may.

Undoubtedly, there may be times when the problem that has arisen is so major or significant that normal operations are damaged or disrupted beyond reasonable or rapid repair. This is when disaster recovery (DR) takes over, and plans for dealing with the most major of issues is a key area of information assurance. This entails having to do things in a significantly different manner as a result of some very major problem. It might be short term or long term before 'normal operations' are restored. Indeed, it may be that normal operations are never restored, in which case the DR plan may become the new 'normal operations'.

Naturally, as has been discussed earlier in this book, the keys to planning of this sort are the risk assessment and BIA. It is possible to consider major DR plans involving backup sites, significant investment and major planning and testing, but if the anticipated problem is only likely to happen once in 100 years, and only then as a result of some event so significant that half the Western world has been thrown into uproar and disarray, then it is probably not sensible to expend significant amounts on a DR plan – unless you are the organisation that is required to deal with that situation of course!

The degree to which a DR plan is developed has to depend entirely on the BIA. If the impact of any specific or general event is so severe as to increase significantly the likelihood of the organisation not being able to operate effectively ever again, then this must be anticipated and appropriate plans developed. On the other hand, if the event is likely to cause minor disruption and can be dealt with in the first few hours of the situation to get back to something close to normal, significant investment is possibly not required. There is, though, another factor to consider. If this seemingly minor event occurs each week, then the cumulative effect of the event may raise its importance and impact and, therefore, the likely acceptable expense on its BCP.

A key difference between BCP and DR is the scale of the plans invoked in any specific situation. If the plan calls for minor adjustments to normal working practices or, at the most, a comparatively small change in normal operations, then this should form part of the BCP. DR, on the other hand, is generally focused on contingency planning for IT systems, and may be part of a larger organisation-wide BCP that includes Human Resources, Sales, Research and Development, Finance, Production, Warehousing, Distribution and other key departments within the organisation. There is also a clear link to financial consequences as well, both in terms of the disruption to operations and the cost of the implementation of the plans.

The distinction between DR and BCP in each operational environment is fundamental to the way in which they are treated. It is critical that each of these aspects is considered and processes, safeguards, provisions and other activities are based firmly on the risk assessment and the consequential BIA. Inevitably, this will mean that some organisations place great emphasis on the DR side of their planning while others might choose to effectively ignore that and deal only with the BCP side. While at first glance this may seem foolhardy, it may actually be the most sensible way of dealing with the issue in the most cost-effective manner.

For example, a business may conclude that the only real risk to its continued operation in the medium term – the area covered by DR for example – might be that there is a major disruption of all services in the office block in which they operate. However, if those services are commonplace, water, telephones, electricity, computers and so on, and there are many similar establishments within an appropriate distance that could provide these facilities, then it could be entirely acceptable simply to ensure that all records are duplicated and maintained in secure storage somewhere else without the need for a DR plan that is any more complicated.

On the other hand, in a business driven by the rapid turnover of cash, as a result, say, of retail esales, a short-term loss of its website could mean a major downturn in productivity and hence finance. In this case, the emphasis might need to be on very significant BCP and DR planning if they are to be able to cope with any eventuality. The operational environment in which the company works must be one of the most important factors to consider when looking at the possible outlay on DR and BCP.

RESILIENCE AND REDUNDANCY

As has been seen in Chapter 2 on risk management, one method of treating risks is risk avoidance or risk termination. In the context of business continuity and DR, this can be achieved in one of two ways: resilience or redundancy.

Resilience involves ensuring that there are no so-called 'single points of failure'. This means that there are no systems or services within the organisation's infrastructure that can bring all or part of the overall operation to a standstill, or degrade it so severely that it cannot continue to provide the level of service expected by the organisation's customers and stakeholders. In practical terms, this normally involves the deployment of an additional number of facilities; for example, multiple web servers with load balancing across them so that if any node fails the remainder will take up the load and continue to provide service. Naturally, resilience has a financial cost, which must be balanced against the potential financial losses if the service were to fail.

Redundancy, on the other hand, has a slightly different approach, in that there is always a standby system or network connection that can take over if the active system or network connection fails. This normally involves a duplicate of the active system or network connection, often located in a different location, with the added benefit of providing additional resilience. A standby system may be termed as 'cold', in which case it must be configured from scratch and it may take some time to bring this into service. A 'warm' standby system may be partly configured and may even have data loaded – usually to a known backup point – this will still require some work to bring it into service and fully up to date. Both cold and warm standby systems are frequently constructed from test systems, so that they have a dual use. It is also common (where practicable) to use one cold or warm standby system to act as a redundant system for more than one active system.

So-called 'hot' standby systems are normally fully configured, will contain up-to-date data and can be brought into service very quickly. Beyond this, there will be 'high availability' systems, where switchover from the failed system to the standby is instantaneous, and there is no threat of any data loss. This type of system is the most

costly and can work in one of two ways: asynchronous replication involves transmitting data from the active to the standby system without waiting for any acknowledgement that they have been received; synchronous replication involves transmitting data from the active to the standby system, but waiting for acknowledgement that they have been successfully received and written to the system's storage.

Asynchronous replication is much faster since there is no waiting time, but there exists the possibility of some data loss; while synchronous replication is somewhat slower, but avoids the possibility of data loss. The choice is one of balancing response time against the chance that one or more transactions may be lost and will depend on the nature of the service being provided. Those that are simply providing information may be better using asynchronous replication, while those that provide financial services or order placement will benefit from synchronous replication.

It goes without saying, of course, that resilience and redundancy must apply in the context of the entire IT infrastructure, including buildings, power supplies and environmental systems.

When looking at the need for resilience, organisations must consider the potential impact of cyber-related incidents. They must take into account that system failures may not be the only cause of service delivery problems, but that such things as viruses and denial of service attacks could also impede their ability to deliver service. In that case, business continuity and DR plans and preparations must include the ability to manage cyber-attacks – often in real time. An excellent example of this is that of the Wannacry ransomware virus, which in May 2017 affected more than 20,000 computers around the world that were running the Microsoft Windows® operating system. The virus encrypted all data on the computers' hard drives, and, although the attack was stopped within a few days, the impact was very significant on many organisations, including hospitals.

It is useful to note that for any 'hot' standby system, there might have been the additional problem that the virus infection also encrypted the backup servers, rendering them useless. An off-line backup is therefore important, and systems should be developed utilising the concept of taking a 'snapshot' of the database at frequent and regular intervals, but the system should only be connected to the main database for as short a time as possible.

APPROACHES TO WRITING PLANS AND IMPLEMENTING PLANS

There are a variety of ways in which BCPs and DR plans can be developed, some more effective than others, but the most effective is very often simply the one that works for the organisation itself. The first step will be to ensure that the risk assessment has been completed effectively, considering all those 'unlikely' events. It is not necessary to consider all possible events, but more the consequences of the event. For example, there are a number of potential events that could make a building unavailable: bomb, flood, fire, aircraft crash, building decay, animal infestation, power outages disrupting supplies for lifts, doors, lighting and so on. Some of these may not affect your building but occur in a neighbouring one, with the same consequences. All these will result in similar outcomes with only slightly different impacts, one of which might be the duration of the disruption.

Nevertheless, plans can be developed to deal with all these eventualities by addressing what to do if the building is unavailable for a significant period. This is where the procedures for implementation become significant. As a result of the specific event occurring, those in charge would need to decide which aspects, if any, of the DR plan should be implemented (what is known as the invocation decision) and to what degree. This will be determined by the magnitude of the problem, the anticipated duration of the problem and the impact on normal business operations.

It is often good practice to involve a number of key staff members in a workshop to determine what would really have a major effect on their work. It is sometimes overlooked that the supply of a seemingly minor part of the business process could have a major impact by disrupting several other aspects of the operations. This may have been done at the original risk assessment, but a slightly different approach may be required in order to consider these more unlikely and perhaps fanciful events.

It is often questioned how far this fanciful suggestion of events should be taken. There is no rule, but what is clear is that in the recent past some organisations have been caught out by not going far enough. When the World Trade Center in New York was first attacked in February 1993, with bombs placed in the underground car park, many of the companies working in the building at the time never traded effectively again. This was the result of the building being declared unsafe and hence closed to all access until a full building survey had been completed – an activity that took several months. The main problem for the companies involved was the loss of access to vital records, notably those pertaining to cash flow.

The lack of access is a major issue regularly overlooked. It is often the case that this is caused by a problem elsewhere – in the building next door or perhaps some distance away. The need for security cordons of varying sizes, up to some kilometres in certain circumstances, means that the drafting of BCP plans must consider who else is in the neighbourhood and the possible consequences of their presence. Those working near the Buncefield oil depot in the UK should at least have considered the possibility of an incident affecting them, even if they hadn't expected the major events that actually transpired in December 2005.

What is important, though, is to keep a sense of perspective. A very experienced member from the Institute of Advanced Motorists (now IAM Roadsmart) used to say to potential members that one should never be surprised by anything that happens on the road – expect the unexpected. This was until he talked to a fellow driver who had experienced a Boeing 747 jumbo jet trying to land in front of him on the A4 road out of London instead of the parallel main runway at Heathrow Airport. Perhaps that is beyond reasonable expectation. That an aeroplane could 'fall out of the sky' onto any specific location is possible, but there are clearly places where this is far more likely – under the main flight paths into and out of major airports for example. So, should this threat be ignored elsewhere? Since it is likely that the business impact will be high, then it is worthy of further consideration, but then consideration of the likelihood should come up with a reasonable assessment of the risk and suggest appropriate countermeasures.

Implementing these plans (making the plans available for use as and when required) also requires some considerable and detailed planning. Simply issuing a document to all staff and assuming that is good enough would be naive. The implementation needs

to be accompanied by a significant awareness and education programme to ensure that all staff end up being fully aware of the plans, as they affect them, and what actions they need to take in the event of an issue arising. It is almost inevitable that staff will say they will check on the details 'tomorrow', which, almost as inevitably, will be too late.

THE NEED FOR DOCUMENTATION, MAINTENANCE AND TESTING

Documentation is vital and can be the difference between a successful conclusion to an event or a disaster. It is vital because everyone involved with and affected by the event must have the same understanding and expectation of the response. If anyone 'does their own thing', it is likely to cause even more problems and possibly counteract the good work being done by others. Advice can be sought from the professionals in this area, who will help to draft and test documentation and plans for organisations, particularly where the issues are complex or likely to prove expensive. Even for the less grand requirements, professional advice at the start can help to ensure the development is on a firm footing.

Just documenting the expected actions and procedures, though, is not enough. If the documents are not available to those who need them at the time they need them, in a convenient location, they are almost worthless. It is clearly not sensible, for example, to have all the documents stored in a nice secure place in the office if one of the potential events is a lack of access to the office. Giving them to any of those staff members to look after at home is a better idea unless that staff member is critical to the emergency actions and they happen to be on leave when the event happens. There is also the concern about the overall security of the plans that may well be commercially sensitive; once again there needs to be the compromise between confidentiality and availability.

Maintenance of the plans is another area that can cause problems. Organisations have in the past spent considerable amounts of time and effort getting the plans written and checked in preparation for some significant potential problem, such as the millennium bug. They then leave the plans on the shelf for the next few years without further attention. The plans are not used in anger and so they become invalid through lack of attention and maintenance. Then, when they are required again, the plans are out of date, incomplete, don't work properly and likely to cause more problems than they solve.

Contracts for the provision of a DR facility are now commonplace, with companies offering standardised facilities of desks, chairs, networked computers and printers, telephones and so on that can be made available at very short notice (usually within a matter of hours) and configured quite quickly (usually within about 24 hours) to emulate the normal working environment in which a company is used to operating. The advent of cloud technology (covered elsewhere in this book) has made this all the more accessible and a suitable option for companies. The investment in such a provision must be made after completing and accepting the findings of the BIA in order to justify the expenditure.

Routines for testing and checking the details of the plans must be comprehensive, but, again, there is a need for a reality check. To close a factory, even for a day, could have very significant implications for the company concerned and, while it might provide an

excellent test of the BCP, may actually do more harm than good to the profitability of the company.

Testing can be carried out in a variety of ways. The usual first step is to do a desk check of the plans. This involves the key people sitting round a desk pretending to do the activities required of them in the plans. This will often sort out the major issues with plans and will enable updating and checking to be completed. As a result of this and the appropriate updates, the next stage might be to do a limited walk-through of all the parts of the plans, but in manageable chunks. Many people are familiar with the fire drills required for organisations large and small. This type of practice for similar incidents – chemical spills, power outages, bombs and so on – can provide a good reality check. Is it really possible to evacuate a 15-storey building with no power for the lifts or main lights within the stipulated time in the event of a fire? What happens if one of the staircases is not available? This also will raise further issues as well as acting as a reminder for staff of the existence of the emergency plans.

The next step may be a full-scale enactment of the plans. This can be critical if the testing so far has only been carried out in parts. When checking the plans for one organisation based in a three-storey building it became clear that each floor was expected to evacuate to a specific building in the neighbourhood in the event of a major incident. This was fine until it was recognised that each of the three floors was expected to go to the same building, but that building was not large enough to accommodate all the occupants of all three floors simultaneously. One floor at a time was fine, but not all three.

If the full enactment is to be effective it needs significant planning and must be co-ordinated by a central control organisation, often called the incident management team, who in reality would also control the implementation of the BCP or DR plan. Often, such facilities are set up and maintained in case of an emergency, but clearly they need to be somewhere that is likely to be available in the event of the most serious incident. Housing them in the basement of the main building is fine provided that the unavailability of the building is not one of the major incidents being considered.

The reality may well be that this central control facility will need to be set up wherever is available in the event of an incident. Therefore, it may be more practical to consider temporary facilities and to define the essential requirements in terms of communications, office space and the like. Then, having all the necessary information readily available in a portable format should provide the necessary resilience. Identifying a number of potential facilities in other sites belonging to the organisation, sister organisations or publicly available facilities, such as community halls, sports centres or the like, might be considered prudent. Clearly, *assuming* such facilities would be made available is not sensible. A short discussion about the facilities with their management team will ensure they would be compliant with a request should the appropriate situation ever arise.

It is often considered appropriate to simulate possible emergencies in some way. One way to do this is to use the 'brown envelope' technique. This entails setting up a scenario of a major incident and then drafting a number of instructions or information sheets given or sent out to relevant staff in 'brown envelopes'. The relevant staff members are instructed to open the envelope at the appropriate time and to take the necessary

actions in accordance with the information supplied. This could be to make a telephone call, to invoke a particular element of a plan or some other action, including reacting in some specific manner. The control of this exercise must be outside and independent from the 'normal' control of the incident in order to ensure it is as realistic as possible and engenders the correct responses from the incident management team.

This technique can be used to test the plans for one specific location or area of an organisation, or indeed for the whole organisation if it is deemed appropriate and necessary. It may be that the co-ordination of actions in a significant number of locations (shops or branches located all over the country or similar) may require the whole organisation, or at least a major part of it, to participate in this level of testing. This may need to be done over a weekend, for example, when the normal work is unlikely to be severely affected. To try and run such a major exercise during normal operations may be deemed too expensive or difficult to do without seriously affecting normal operations. It is also important to pick a suitable time from the business perspective. Choosing the busy season or at the financial year-end is unlikely to result in wholehearted buy-in to the tests and so will reduce their effectiveness. Nevertheless, whether it is a small-scale trial of the plans in one location or a full-scale whole-organisational simulation, it is critical that the exercise properly tests the appropriateness, effectiveness and comprehensiveness of the overall planning. Guarding against the ignoring of issues that are outside the control of the incident management team is also vital. The weather, rush-hour travel, pandemics and so on should all be considered and the potential impacts measured appropriately.

NEED FOR LINKS TO MANAGED SERVICE PROVISION AND OUTSOURCING

Any plans for dealing with emergencies must naturally cover those services supplied by third parties as well as those the organisation itself carries out. This must start with the contract, which has to include some mention of the expected level of service provision in the event of an emergency. It would not be helpful, for example, if the telephone contractor said that it would take three days to enable the telephones in the incident room in the event of an emergency.

All contractors must then be closely involved with the development and testing of any set of BCP or DR plans. The services they provide may be critical to the overall success of the DR and may be essential in dealing with minor incidents as part of the BCP. Indeed, it may be that, within the contract, the responsibility for the management and resolution of minor incidents as part of the BCP could be passed over to the supplier. It would naturally be prudent to check that their plans work and are consistent with those of the client organisation.

Where managed services are under consideration, the contracts must provide the facility to ensure that these services will continue to be provided under the changed circumstances of the DR or BCP. While it may not be feasible to envisage all circumstances that might arise, and hence ensure every angle is covered, it should be possible to ensure that changes to the contract don't require lengthy procurement processes that will take longer to sort out than the original problem.

NEED FOR SECURE OFF-SITE STORAGE OF VITAL MATERIAL

As already mentioned, the access to BCP and DR plans is critical. It is useless to spend a lot of time and money on developing and testing the plans if they are not available at the crucial time. It is vital that the plans are available to those who need them whenever and wherever they are. They must always be consistent – everybody must have the same and latest version – and they must be in a usable format. It is inevitable that there will be a significant quantity of personal data in the plans – contact details for all the key players and indeed possibly for all staff. The requirements of the DPA and the GDPR must be acknowledged.

In the past there have been 'war chests' containing all the plans, contact details and the like for the management of a possible situation. These had to be stored somewhere and it would be the responsibility of one or more people to take the chest to an appropriate location when required. It has been known for key individuals to have to store these chests at home as being the place least likely to suffer a major incident.

With the more modern technology now available, a better solution might be to provide them to key people on an encrypted memory stick in a format that requires no major application to read them, such as Adobe Acrobat Reader™ – the small application could be on the stick too. This could be used anywhere there is a computer with an internet connection (now an acceptably common requirement) and, when accessed, with a simple piece of software it could be made to check a central, perhaps cloud-based, repository for the latest version. This meets many of the issues raised above, including the issue of security, particularly if the stick is securely locked in some way in case it should fall into the wrong hands.

The alternative is to find somewhere secure to store all the plans that is always going to be accessible in the event of a problem at the main building. This may prove problematic. A sister organisation or perhaps another branch of the same company might be an appropriate place, but it is more likely to be geography dependent than anything else. There are accepted distances for the cordons that might affect the accessibility to a building. It is important to take these sorts of distances into account when deciding where the most appropriate storage location might be. Naturally, among the possible events leading to the use of the BCP or DR plans, a terrorist incident or fire will be one of the major factors to consider. The Buncefield and World Trade Center incidents may make people think again about the impact of other businesses in the locality, and consequently where an appropriate place might be to store the emergency plans safely.

It may be necessary to consider how critical information or other assets might need to be stored elsewhere to enable 'normal business' to be maintained. This could mean taking backup tapes from computer systems off site each night to ensure their availability or could entail having a secure store of critical supplies (maybe drugs for a hospital for example) somewhere appropriate. Yet again, it is important to have reality checks on these types of stores. An organisation could end up with a complete ready-use store elsewhere that in turn means further BCP issues, more expense and ultimately distracting the organisation from their main business activities.

The widespread availability of cloud-based services has had a major impact on BCP and DR planning. In many ways it makes it much easier. Assuming the cloud-based

facility itself is reasonably resilient, and the technology is designed specifically to try and achieve that, then the latest versions of documents could just be stored in the cloud. The provision of suitable alternative office space is then simplified – essentially anywhere that has an internet connection will do, with appropriate capacity being the only major concern. Indeed, maybe there is no need to consider alternative office space at all; simply telling staff to work from home on the cloud-based systems used every day might be sufficient.

In other ways, cloud-based facilities have increased the task of the BCP or DR manager. What happens if the cloud-based facility is no longer available for whatever reason now becomes a much more serious concern. For instance, banks that have had major problems after a seemingly routine software update have suffered the wrath of users when they were unable to undertake day-to-day transactions. Consideration of the wider supply chain in all its formats is now a critical part of DR and BCP planning.

NEED TO INVOLVE PERSONNEL, SUPPLIERS AND IT SYSTEMS PROVIDERS

It is vital, naturally, that all staff, full- and part-time, temporary and permanent, must be fully aware of the workings of the BCP and DR plans, as all are affected by them. This may also affect visitors to the site, but that may be best dealt with by the host staff member being responsible for the security, safety and well-being of their guest. The requirement for an ongoing education and training programme is clear, but again this must be tempered against the risk and compared with the outcome of the BIA. It is good practice to ensure there is an induction course covering all the basic requirements and perhaps highlighting the individual's responsibility in the event of an incident. This could then be supported by a series of 'exercises', reminder sessions, updated leaflet or email distributions and the like.

The extent of such a programme of education and training must always consider the other key players, including suppliers and outsource companies. Staff who routinely work within the establishment of the client, manning an IT helpdesk or other such facility for example, must be involved and actively engaged in a manner that doesn't affect their productivity but helps to ensure they do not become a liability. In the event of a fire, the first question asked by the arriving fire-fighters will be: 'Is anyone inside?' Unless an unequivocal 'No' is given in response, their first efforts will be targeted at finding and evacuating people. If only part of the organisation has any method of recording those who are working in a building, it will be very difficult to ensure that a full picture of the evacuation is available. This in turn could not only compromise the safety of the fire-fighters unnecessarily, but could also affect the fire-fighting efforts to reduce the damage caused.

When it comes to DR, the need for suppliers of crucial services to be involved is also paramount. Their responsibility should be towards their client, but it may take second place if the operational effectiveness of the company itself is in danger. The joint understanding of client and supplier is then all the more critical. How this is dealt with in a contractual manner will be determined by issues and requirements outside the scope of this book. Nonetheless, it is critical that those who have to manage and work with the contracts on a day-to-day basis must understand the detail of how any service level agreement will work in the event of a major problem.

On these occasions, IT companies are always the first ones that come to mind, but it must be remembered that all the other suppliers, ranging from telecommunications and mail through food and stationery supply to cleaning and waste disposal services, must be considered as well. While not all of these may need to be dealt with within the first 24 hours of an incident, it is highly likely that they will have to be dealt with within a fairly short period of time in the event of an ongoing problem. Workers in a new building used as a DR site will still need to be fed and watered, supplied with paper and have the wastepaper cleared.

The supply chain is a particular issue, and there are a number of options that could be considered. These might include: holding stocks of critical elements in an alternative location; ensuring that all the suppliers themselves have up-to-date and workable BCP and DR plans; and ensuring the contracts cover the more likely eventualities.

RELATIONSHIP WITH SECURITY INCIDENT MANAGEMENT

Incident management is the term used to describe the work done to deal with the incident itself. It is usual to have a team specifically trained and ready for this work since it is often quite technical. There are often legal aspects to consider too – if the building is the site of a crime, then the police will close it to all and may impound anything contained in the vicinity. Forensic readiness is the work of ensuring that, when an incident takes place, crucial evidence such as temporary files on a computer, fingerprints (electronic or human), audit logs and a wealth of other materials are not destroyed, whether intentionally (by the perpetrator) or accidentally (by turning off a PC or server or walking through a crime scene in dirty boots).

Once again, planning is the key and it is vital that due consideration is given to all the materials (information, tapes, etc.) that will be required to maintain business as usual and/or activate a DR site. Where these materials are stored has been mentioned earlier, but it is worth considering the possible eventuality that there is no access at all under any circumstances to 'normal business information'.

The relationship between the teams responsible for dealing with the incident, the BCP and the DR plan must be a very close one. There is a wide range of overlapping areas and as such there must be no chance of anything being missed. It is more likely that doing it twice is a better option than running the risk of it not being done at all. Clearly, there are issues in many areas here and, once again, it is at the planning stage that these must be talked through and worked out in such a way as to be applicable in any circumstances. Table-top walk-throughs often identify problems in this area when one team can say '... and then I would do this' and another team can say '... but we would have done that already'; or, 'You can't do that because we will have already done this ...'.

One of the most difficult areas to work through is that of prioritisation. It might be very important for the business to continue in certain areas, but, if the incident management team wants to seal off a specific location, it may not be possible to get the necessary information or equipment to continue. Someone (frequently a board director or a member of the senior management team) must make the necessary and, most importantly, timely decisions.

COMPLIANCE WITH STANDARDS

When the syllabus for the BCS Certificate in Information Security Management Principles was devised, and this book was envisaged, a specific decision was taken not to follow or use any one standard. The reason for this was to ensure that the resulting qualification was as generic and applicable around the world as possible. It is for that reason that little mention has been made of specific standards, although there are many that could have been covered.

There are a number of standards in the UK and elsewhere in the world that cover some or most aspects of the management of BCP and DR. The international standard ISO/IEC 22301:2019 is the main standard for business continuity requirements, while the related ISO/IEC 22313:2014 contains additional guidance and ISO/IEC 27031:2011 provides guidance for IT readiness for business continuity. It is worth noting that many of the standards mentioned here are, at the time of writing, undergoing a refresh and new versions are expected in due course. The main standard for DR services is ISO/IEC 24762:2008.

There are a number of other documents that provide useful information: Publicly Available Specification 77 (PAS 77) – IT Service Continuity Management; the ITIL guide on service management best practice now produced by AXELOS; and the Business Continuity Institute produces its own good practice guidelines that are an excellent source of information – the current version at the time of revising this book is from 2018. You would be well advised to consult the available appropriate documentation to ensure that what you implement in this area is based on good practice and therefore likely to be successful. The real problem with this area of information assurance is that, like a teabag, one is never going to know how good it is until it is put into hot water!

ACTIVITY 8.1

Ms Jackson has realised GANT have no real backup of the members' database. Indeed, they are not really doing any backup at all. Advise her on the minimum requirements of a backup system and the frequency of performing this necessary task.

SAMPLE QUESTIONS

1. **Which of the following should inform the decision to invoke a business continuity plan?**
   a. Risk assessment.
   b. Security policy.
   c. Business impact assessment.
   d. All of the options above.

2. **If the solution to dealing with an issue that has arisen is to move to alternative office accommodation, what would the plan used likely to be called?**

   a. Disaster recovery plan.
   b. Business continuity plan.
   c. Alternative accommodation plan.
   d. Business disaster plan.

3. **An organisation decides to have two separate connections to the internet through different providers and physical connections. The work is shared between the two connections. What would this arrangement be called?**

   a. Redundancy.
   b. Reduction.
   c. Resistance.
   d. Resilience.

9 OTHER TECHNICAL ASPECTS

In this chapter you will gain an understanding of the important aspects of incident investigation and how the forensic evidence may be preserved. You will learn about the basic concepts of and uses for cryptography, and threat intelligence and its role in a modern organisation's defences.

INVESTIGATIONS AND FORENSICS

It has already been mentioned that, even in organisations with very effective governance, there will be occasions on which it is necessary to investigate activity and use forensic techniques to discover and preserve evidence for later use. Part of incident management is about the ability to identify answers to the questions: who, why, what, when, where and how? Some of this has been described previously and you will be referred back to that material where appropriate. You are advised to read the section on 'Security incident management' in Chapter 3 if you have not already done so. Much of the knowledge for this section is described in that material.

LEARNING OUTCOMES

Following study in this area, you should be able to define and explain each of the following terms and processes and be able to describe their appropriate use as applicable.

Common processes, tools and techniques for conducting investigations

If it does become necessary to work with an external law enforcement organisation following a security incident, they are going to want to collect evidence for use in further investigations and possible prosecutions. The UK's Police and Criminal Evidence Act 1984 (PACE) defines very strict standards of conduct in order to allow the police to demonstrate that the evidence is valid and admissible in court. The required course of action where IT assets are concerned can be very complex and it is very easy to render evidence inadmissible. It is strongly recommended that the incident response team is properly trained in how to deal with these requirements and work with law enforcement representatives to achieve the desired outcomes. A good source of advice on how to

conduct this activity can be found in the latest version of the NPCC (previously ACPO) guidelines for computer-based evidence.

Whether or not there is the possibility of involving a law enforcement organisation, it is good practice to observe the same high standards of rigour when investigating an incident. The findings could be used for a prosecution or an internal disciplinary hearing. A hearing that leads to the dismissal of an employee could then go to an employment tribunal and having evidence that is legally admissible, through following good practice, will be a great help in winning the case.

It is advisable to have someone designated as the evidence custody officer. This person is responsible for collecting and securely storing evidence while maintaining a good documentary record to preserve what is often referred to as the 'chain of evidence'. This is a clear, unambiguous and indisputable record of everything that has happened to an asset (regardless of its nature) from the time it was located to the time it is presented in court or elsewhere as evidence in a criminal case.

There are forensic tools available to collect and examine evidence from IT systems. They should only be used by skilled and properly trained investigators because of the ease with which evidence becomes contaminated and inadmissible. Many organisations will not have this kind of resource in-house, but if a list of forensic specialists has been prepared in advance, they can then be located quickly and easily and contracted-in when required. Some organisations may consider it appropriate to set up a 'retainer' or framework agreement in advance to ensure a rapid response when it is required.

Legal and regulatory guidelines

Chapter 3 described the greatly increased legal and regulatory requirements that have been introduced in many jurisdictions worldwide for corporate governance and accountability. Part of this requires the ability to investigate incidents and attribute responsibility to one or more individuals. More importantly, the investigation must be conducted in a manner that preserves the evidence in a form that is compliant with legal procedures. This is also outlined in Chapter 3 and describes the challenges in collecting and preserving the evidence. Part of the section on 'Relationship with security incident management' in Chapter 8 explains the need to practise incident response and investigations to help identify any flaws in the plans. Time spent planning, training and practising is never wasted.

One final tip: whatever is done, make sure the appropriate rules of evidence for the relevant jurisdiction are well understood, in particular regarding the chain of custody for the evidence. Sometimes legislation such as GDPR may require the consent of individuals for the collection and storage of material that may contain their personal information and data about the behaviour of users. It is important to ensure that appropriate consent has been recorded in advance and ideally this should be specified in a staff member's contract of employment or the terms and conditions of use of a service.

Need for relations with law enforcement

Chapter 1 has a section entitled 'The role of information security in countering hi-tech and other crime', which explains that there are times when it may be necessary to

engage with specialist law enforcement organisations who work in the computer crime area or those working to protect what is often generically called the critical national infrastructure of a country. There may be agencies, such as the NCSC, CPNI and NCA in the UK and the Department for Homeland Security (DHS) in the USA, who wish to work with an organisation to help improve the level of security if it is considered important to the 'national interest'. This includes utility, energy and communications companies, among others.

There are also various emergency response teams, such as CERT in the USA, GOVCERT in the UK and in other countries, and the Forum for Incident Response and Security Teams (FIRST). Find their websites and read their guidance documents for much valuable information. The section 'Processes for involving law enforcement' in Chapter 3 describes the mandatory requirement in the UK to report certain crimes or activities to law enforcement agencies. It is, for instance, mandatory in the UK to inform the police if there is a suspicion of terrorist activity or that child pornography has been viewed or processed through the IT systems of an organisation. UK legislation also requires the reporting of suspicious financial activities. The law enforcement agencies often have specialist staff who can offer advice and guidance to any organisation that feels at risk from logical or physical attacks. It is always worth contacting them for any material they can provide. In the UK this may be provided through the NCSC. There are times when it will be necessary to involve law enforcement or other similar organisations in the response to an incident. If there is any likelihood of criminal activity or other deliberate action, the appropriate authorities should be notified at an early stage. It is important that senior management has a good understanding, in advance, of the legal requirements for reporting certain events.

Another area of crime where law enforcement agencies may be involved is that of attempted extortion and blackmail by use of techniques such as denial of service attacks or ransomware. In this case, the NCA is the appropriate body in the UK to contact. Activity of this sort, or malware (malicious software) discovered in UK government departments, should lead to a report being passed to the NCSC. In other countries there will be similar organisations and law enforcement agencies who will be the key point of contact for such incidents.

One last possibility is that an organisation may be visited by law enforcement officers conducting an enquiry into activities of which management has no knowledge, as discussed in detail in Chapter 3, 'Processes for involving law enforcement'.

Security issues when procuring forensic services and support from third parties

Chapter 3 described the reasons why organisations must develop policy, processes and procedures, developing standards, guidelines, operating procedures, and so on, internally and with third parties. These reasons include the need to be ready to investigate information security incidents and possible criminal offences. Although the events that initiate this kind of response may vary greatly, the response itself will

often be very similar in nature. It is important that the process has been developed and checked with a specialist in this area before it is used. The response needs to be prompt, appropriate and valid. Any mistake can render the findings inadmissible in a court or employment tribunal. It is no use bringing in law enforcement agencies if the evidence has already been contaminated and cannot be used. It is also vital that, when an investigation starts, the incident team consults senior management to decide whether the organisation intends to prosecute (e.g. if the organisation has been hacked or suffered a loss) as this determines whether or not to involve criminal justice organisations and the level of effort that will be required to gather the evidence, for example getting external specialists on site to make images of hard drives and the contents of memory for evidence.

Most organisations choose to outsource forensics and investigations to a third party. This is because they are too small to have their own skills in-house, because they cannot justify the cost of employing and maintaining the skills themselves, because it is not considered to be a major risk or for any number of other reasons. Organisations that do decide to have their own in-house resources must make sure that the products and skills they acquire are sufficient for the task and will provide legally admissible evidence. There are well-known products with very good reputations and there are others that make bold claims but don't meet the standard. It is important to do some homework before buying goods or services. Once they have been acquired, it is essential to make sure that assets and skills are tested and updated regularly to ensure they are ready when actually needed. One simple mistake can invalidate everything else done correctly.

If the organisation needs to go outside its own resources in order to complete the investigation or response, there are some important considerations to take into account. One of the main ones is that of NDAs. Earlier mention was made of a framework agreement. This sets out in advance all the contractual issues. Such an agreement can take time to negotiate and agree, so it is preferable to put it in place before an incident occurs. In addition to the normal contractual terms about payment, provision of service levels, dispute resolution and so on there also needs to be a section that puts in place an NDA between the signatories. This document provides legally binding confidentiality so that the third party is under obligation to provide the same level of confidentiality to the information seen or facts discovered as a permanent employee of the company. It should be noted that these agreements are still subject to the overriding requirements of the law. The third party is legally required to notify law enforcement, for example, of any suspected child pornography, danger to life or terrorist activity, even if they have signed an NDA.

The document should also define the requirements for the following:

- standards required in preserving evidence and accompanying documentation to those defined in PACE/NPCC guidelines or national equivalent for legal admissibility;
- the handover, assured destruction or secure erasure of all materials obtained by the third party at the end of the investigation;
- participation in any review at the end of the incident to help improve the response process.

One area that must be addressed is that of timeliness. In order to preserve evidence, devices must not be used, even if it is a critical server or business asset, until forensically investigated. That means that any evidence must be collected as soon as possible to allow a prompt return to normal operations. It is most important to gain agreement and support from senior management in advance as to the method by which third-party support is procured, how any such work is to be done and the means by which their help should be invoked. There isn't time to identify possible suppliers, go out to tender, negotiate terms and schedule the work. Incident response needs to be very prompt. A framework agreement is usually the best answer.

ROLE OF CRYPTOGRAPHY

Information security managers need a sound appreciation of the role of cryptography in protecting systems, services and assets, including awareness of the relevant standards and practices.

LEARNING OUTCOMES

Following study in this area, you should be able to explain and justify each of the following concepts:

- Basic principles of cryptographic theory, techniques and algorithm types, their use in confidentiality and integrity mechanisms and common cryptographic standards and protocols.
- General policies for cryptographic use, common key management approaches and requirements for cryptographic controls.
- Principles of link, file, end-to-end and other common encryption models and common public key infrastructures and trust models.
- Common practical applications of cryptography – for example, for digital signatures, authentication and confidentiality.

Cryptography is a very wide-ranging and potentially detailed area of information security. No attempt has been made within this chapter to cover the whole topic in depth, as this is better researched from standards and dedicated works on the subject, and a detailed study of cryptography is also part of many InfoSec MSc programmes.

Basic cryptographic theory, techniques and algorithm types

Much of what takes place in the context of exchanging information is based on establishing trust and a secure communications channel between two or more parties. If the parties are meeting face to face and have already established some form of trust between them, then subsequent transactions should hold no major risks. However, when some distance separates the parties or they have never met, such as in ecommerce, they must satisfy themselves that the trust relationship can be established and maintained for as long as necessary. Specifically, each party must usually ensure that, with regard to any information that passes between them:

- it is kept secret, assuming that this is a key requirement of the relationship (confidentiality);
- it is not changed by third parties while in transit (integrity);
- the origin of the information (person or system) is assured (authentication);
- the originator cannot deny having sent the information (non-repudiation);
- it can, if required, protect the identity of the user (anonymity).

There are two similar, but separate, needs to provide confidentiality. The first is to secure information stored in a system against unauthorised access – a process frequently achieved by use of password protection, but occasionally by some form of encryption of the information itself; file and whole disk encryption are examples of this. The second, dealt with here, is to secure information while in transit between the sender and recipient so that unauthorised parties are unable to understand the information even if they are able to intercept it.

In order to provide confidentiality, information or 'plaintext' may be encrypted – changed into 'ciphertext' so that the original plaintext cannot be read or inferred – and then sent to the recipient, who reverses the process by decrypting the message to recover the original plaintext.

Encryption may be used in several ways. First, for example, it may be used to encrypt information during transfer from one computer to another. Second, it may be used to encrypt a number of files on computer media. Third, it may be used to encrypt an entire hard disk drive including the operating system, applications and configuration information as well as the data.

Secret (or symmetric) key cryptography

There are two methods of encrypting information. In the first, information is encrypted effectively one bit (as in binary digit) at a time and each encrypted bit is transmitted to the recipient, who decrypts it in real time. An example of this is that of mobile phones in which the speech, text or application data are encrypted in the handset and transmitted to the point in the mobile network where decryption takes place. The speech or text is then delivered to the recipient as plain speech, text or data. The same key is used to encrypt and decrypt the data, so it is 'symmetric' as it is a single key. This is illustrated in Figure 9.1.

Encryption that works in this way is referred to as a stream cipher, an example of which is A5, used for GSM (2G) mobile encryption.[1]

In the second method, information is encrypted in one or more blocks (normally 64 bits) of data and the entire message is sent as these blocks to the recipient. Decryption need not be in real time, but may take place sometime later, and is carried out on the blocks of encrypted data. This method of encryption is referred to as a block cipher. Examples of block ciphers are Triple-Data Encryption Standard (Triple-DES), Blowfish and Advanced

1 GSM encryption can also be thought of as 'link' encryption, as only the connection between the handset and the base station is encrypted. Secure telephones, which also use stream ciphers, encrypt the entire connection from handset to handset and are therefore said to perform 'end-to-end' encryption.

Figure 9.1 Symmetric key encryption

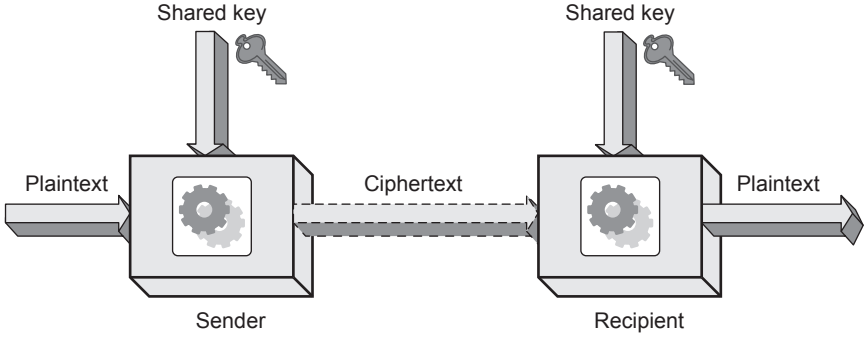

Encryption Standard (AES), which is used, among many other purposes, for IPSec tunnelling. There is also TLS, which is implemented to secure internet ecommerce transactions and the communications for most of the smart energy metering networks.

Apart from the information to be encrypted, the processes of encryption or decryption require two things: a computational method known as an algorithm and a key. Algorithms tend to be few in number as only those that deliver strong encryption can be used with any degree of certainty. The weaker algorithms are weeded out by a process of cryptanalysis in which different types of attack are made against the algorithm in order to try to recover the plaintext. Over time, older algorithms may be retired as increasing computing power and advances in cryptanalysis techniques render them 'weak' and no longer sufficiently secure.

Both sender and recipient keep the symmetric encryption key as a shared secret. Keys are simply a string of bits and, in general, the greater the key length, the stronger the key. Cryptanalysis attacks on keys generally involve an exhaustive key search or 'brute force' attack in which each possible combination of bits that make up the key are tried in turn. Eventually one will work, and by making the key greater in length, the number of possible permutations will be increased, resulting in an attacker taking longer to identify the valid key and decrypt the message, often extending the process well beyond 'practicable' timescales. The weakest part of any cryptographic protocol is normally key management, and an attacker will seek to steal a copy of a key before they resort to trying to break it by brute force.

Another consideration is the 'cover time'. This is the minimum time for which the information must remain secret. It follows therefore that if an attacker can recover the key by brute force in less than the cover time, a stronger key is needed. That said, most key lengths in commercial use today are sufficiently strong to withstand all but the most determined attack by a very powerful computer or a large number of smaller syndicated computers, as happened with the DES some years ago. DES used a 56-bit key, which has 2^{56} possible combinations. This was eventually broken by an exhaustive key search using a purpose-built computer in less than one week. DES security was later improved by tripling the key length to become 'Triple-DES' using a 168-bit key (2^{168} combinations). This has now been widely superseded by the AES algorithm, which is even stronger and also a block cipher. The usual key AES lengths are 128, 192 and 256.

Once the recipient has received the encrypted message, it can be decrypted using the same secret key and the same algorithm as that with which it was originally encrypted, resulting in an exact replica of the original message. However, as has been seen, the cover time for the information may be critical to the two (or more) parties, and if successive messages are to resist unauthorised interception and decryption by an exhaustive key search, it is vital that the encryption key is changed at intervals – perhaps daily or even for each successive message. An attack of this sort was successfully conducted on TLS some years ago.

The problem now is how to pass on the new key to the recipient. Consider for a moment that a man-in-the-middle attack[2] is taking place and that an attacker has managed to recover the message key by exhaustive key search. If the sender included the new key with a message encrypted with the old key, the new key would already be compromised, so another method must be found that will permit the new key to be sent securely. There is a protocol called Internet Key Exchange (IKE), which is widely used to negotiate session keys for IPSec sessions, that allows the agreement of a secret key without it being revealed to a third party. This approach is used to ensure that every new session has a fresh key that has no relation to previous keys. In TLS v1.3 this is accomplished via a much-shortened (compared to previous versions of TLS) TLS handshake protocol that resists rollback and key compromise attacks.

Public key (or asymmetric) cryptography

Long before IKE was developed, the 'key exchange problem' had been a cause of concern. For many years the most secure keys had to be distributed by physical means, which is slow and expensive. This problem was solved initially in secret at the UK's GCHQ, but only after the formation of a commercial company – who also worked out the mathematics that make PKI possible (RSA Security) – was it possible to exploit the technology on a commercial basis.

The end result of solving the key exchange problem is that it is possible to exchange secret keys between sender and receiver without them being compromised and at the same time allow people who have never met or communicated before to do so securely from the very first message. The recipient can be assured that the new key has originated from a trusted source and not from a man-in-the-middle attacker.

In the asymmetric model, there are two entirely different keys, known as the public key and the private key, both of which relate to an individual or entity. As the names suggest, the public key is intended to be used by anybody – it is not secret and is shared with anyone who needs to use it. The private key, on the other hand, is intended to be kept secret by its owner. These two keys are produced in the same operation and are mathematically linked, but in such a way that it is virtually impossible[3] to deduce the private key from the public key, although this may change with the advent of true quantum computing. Anyone can encrypt data to send to a recipient using that person's public key. In normal usage the sender encrypts a message using the public key of the recipient, which can then only be decrypted using the private key, unlike a symmetric key algorithm. This provides confidentiality of

2 In which an attacker has succeeded in inserting themself between the true sender and recipient, and looks to each exactly like the other.

3 Virtually impossible in this context means virtually impossible to most of us. It should be assumed that the security agencies of major governments will have the method, means and motivation to recover private keys.

Figure 9.2 Asymmetric key encryption

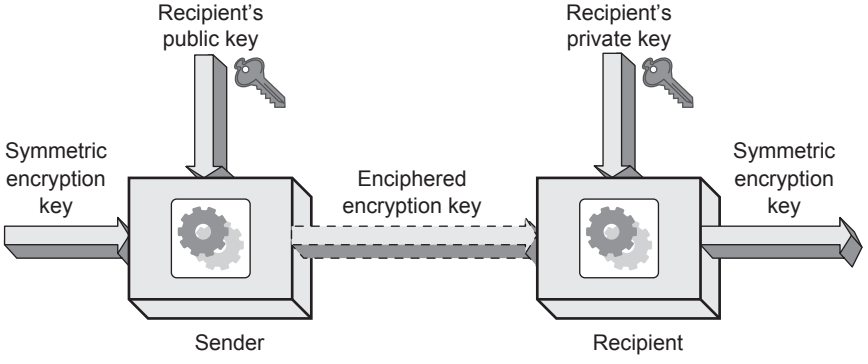

the message. It is also possible to send a message encrypted with one's own private key, which can then be decrypted using the senders public key. This provides authentication of the sender, as only they could have encrypted the message. It does not provide any form of confidentiality as anyone with access to the senders public key can decrypt the message. What normally happens is that, in the digital signature, a person encrypts the entire message with the public key of the recipient and 'signs' the hash of the message data with his private key so that the sender can prove their identity while maintaining the confidentiality of the message contents. This also provides non-repudiation, as the sender cannot deny having sent the message at a later date. Only they could have signed the message with their private key. Figure 9.2 illustrates asymmetric key encryption.

A further problem that still remains is that the both the sender and the recipient need to authenticate the owner of the public key they are about to use to make sure it's not a 'masquerade attack' where the actual recipient pretends to be somebody else. The recipient also needs to be able to check that any digital signature of the message, signed using the sender's private key, really does belong to the claimed identity of the sender. This defends against both impersonation and a man-in-the-middle attack.

Public key cryptography allows the use of 'digital signatures', as mentioned above. These can provide proof of message integrity, authentication of identity and non-repudiation of the sending of the message. This mechanism is strong enough to ensure that a digitally signed message is accepted as a legally binding document in most countries.

Verification of the integrity of a message is achieved by a process known as hashing, described in more detail later in this section. A hash used for PKI is known as a message digest. It is produced when a message is passed through an algorithm that will always carry out the same action on a message to create a numerical value that is totally dependent on the message contents to derive that value.

If a message digest is produced from the original message, encrypted with the sender's private key and sent to the recipient, they will be able to decrypt this encrypted or 'signed' message digest using the sender's public key. They can then produce their own message digest and compare that value with the received message digest. If the two are identical, the integrity of the message has been proven as being unchanged since it was sent.

Examples of message digest algorithms are SHA-256 and RSA, two of the most common hashing algorithms currently used for digitally signing documents and messages. This is illustrated in Figure 9.3.

Figure 9.3 Producing a signed message digest

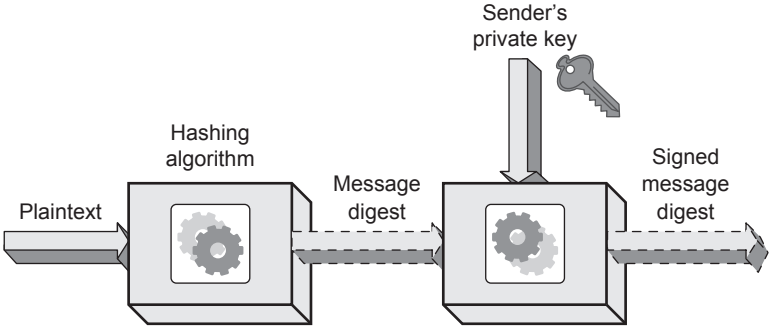

There remains one more link in the chain to make PKI work, that of authenticating the identity of the owner of a key pair, which is done using a digital certificate. The participants have to obtain a digital certificate from a CA, of which there a number of commercial organisations. They have to provide the CA with proof of identity, similar to that required to obtain a passport, in order to obtain a digital certificate. In return the CA certifies that the digital certificate signed with the applicant's private key authenticates the entity's public key by signing it with a key of their own that can be checked as being genuine. It is important to note that anyone can apply for a digital certificate – governments, commercial organisations or private individuals. They simply have to provide the appropriate proof of identity. It allows the user's certificate to be verified as being issued by a genuine CA and confirms that they are who they claim to be. This equates more or less to the situation in which a notary witnesses the written signature of a person on a legal document and by doing so certifies that the person has proved his identity to the notary by using a passport or equivalent identity document.

The certificate chaining process establishes trust in the identity of the holder of the certificate. Most or all of the above components comprise what is generally referred to as the PKI, and can be accommodated in a single transaction that might contain:

- the message, encrypted with a symmetric key, providing confidentiality of the information to be transmitted;
- the symmetric key itself, encrypted with the recipient's public key, providing confidentiality of the message key;
- the message digest, encrypted with the sender's private key, providing an integrity check on the encrypted message, an authentication check on the sender and non-repudiation of the information transmitted;
- optionally, a digital certificate, providing authentication of the sender and non-repudiation for the message sent.

Figure 9.4 illustrates how message integrity is verified.

Figure 9.4 Verifying message integrity

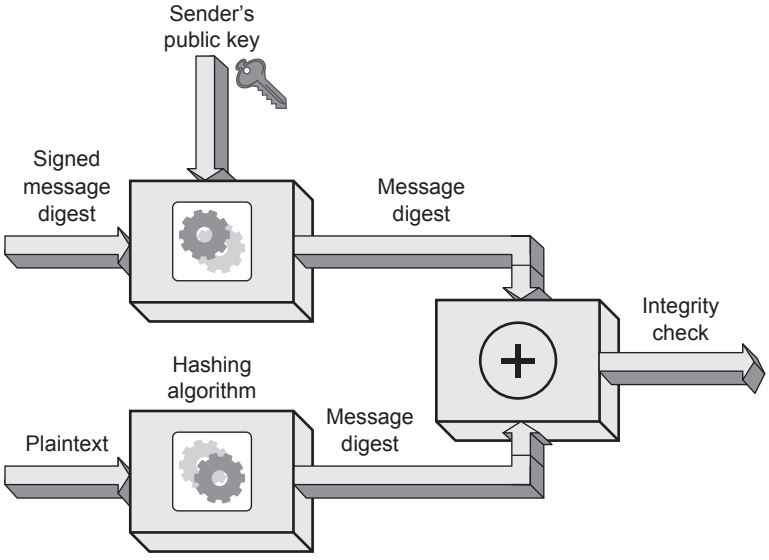

One question that is frequently asked is: 'If asymmetric key cryptography does not suffer from key distribution problems, why not use that all the time?' The answer is quite straightforward. Symmetric key cryptography can be processed very quickly, especially if implemented in dedicated hardware. Asymmetric key cryptography, on the other hand, uses complex and repetitive mathematics to encrypt and decrypt information and this takes longer to carry out, especially on computer systems with relatively slow processing capability. Asymmetric key cryptography is therefore used for encrypting and decrypting shorter items of information such as keys and message digests, as previously discussed, leaving symmetric key cryptography to encrypt and decrypt larger volumes of information. PKI can be used to distribute securely the symmetric keys, solving the key distribution problem.

Pretty Good Privacy

The need for reliable, available encryption prompted an American, Phil Zimmermann, to develop Pretty Good Privacy (PGP) in the late 1980s. It contained all the elements described above: symmetric keys to encrypt information and asymmetric keys to secure the symmetric keys and to sign message digests. It also included the ability to encrypt files on disk and to exchange secure email messages.

PGP is available commercially in a number of forms and it provides for digital certificates. It is also available as OpenPGP and Gnu PrivacyGuard. Instead of relying on CAs to verify identity, the PGP model operates a so-called 'web of trust', which is ideally suited to small networks of users. In this model, each user acts as a CA for other users whom they trust. By extension therefore if someone trusts this user, they will also inherit their trust of the other users. This approach is used by IdenTrust, which is a group of

organisations who have developed their own secure communications domain for doing business.

Hash functions

A cryptographic function, which has grown in importance with the advent of digital cryptocurrencies and other 'blockchain' services, is the hash function. This was mentioned in the earlier section on public keys, but it is also used to create the encrypted versions of passwords stored in computers for the authentication of users when they log on and in single sign-on systems such as Kerberos.

A cryptographic hash function is used to create a fixed-size numerical value from any piece of data that can be of any size. The numerical value returned is known as the 'hash', or 'digest' as mentioned earlier. The function is designed only to work one way, so the original text cannot be derived from the hash value. It is referred to as a 'one-way trapdoor' function because there is no way to recover the original text from a strong message digest – imagine trying to reassemble an egg once it has been scrambled. The output value of the hash function also varies considerably for changes as small as a single bit of data being flipped from a 1 to a 0. This make it almost impossible to change a message so that it still has the same value for the message digest if the message is changed in transit.

This property is also used to sign the transactions in the log of a blockchain, and the contents of the previous transactions are added to the data to be hashed as verification of the transaction. The idea is that an attacker would have to try and modify all or most of the transactions in the blockchain in order to modify a single transaction. That is computationally very difficult, especially if there are multiple copies of the transaction ledger.

Policies for cryptographic use

Policies for the use of cryptography will be largely based on the results of risk assessments, to a certain extent based on legal requirements such as GDPR, and on regulatory requirements that may be sector-specific, for example in financial services.

Policies must take into consideration:

- the method of storage of information;
- the method of transmission of information;
- the required cover time for the information;
- the relative strength of the encryption algorithms and key sizes when offset against the cost and time of processing;
- the relative risks presented by the loss of confidentiality or integrity of information;
- the laws of the countries in which it will be used.

It follows, therefore, that the policies will also relate closely to an information classification/protective marking scheme and the more sensitive the data, the greater the strength of the encryption that must be used to protect it in transit and at rest. Policies

will also determine such things as the frequency at which keys must be changed and how (and where) keys are stored and managed, which in turn will involve processes and procedures, roles and responsibilities and segregation of duties. Suggestions for further reading or study on the subject of cryptography are given at the end of the chapter.

Cyber threat intelligence and vulnerability data

There is an old phrase 'forewarned is forearmed', which is very valid for information security practitioners. The more one knows about the most likely sources of threat and their techniques, the better prepared one can be to deter, defend, detect, respond and recover from cyber incidents.

In an age where attackers are becoming increasingly sophisticated and capable, the time taken for exploits to appear after a vulnerability is discovered or a patch is released is now measured in hours and days, not weeks. It is important for every organisation to understand the threats they face and to monitor for any vulnerabilities in the operating systems, software and hardware assets that they use in order to operate.

THREAT INTELLIGENCE

The primary purpose of threat intelligence is to keep an organisation informed of the risks that they face and how best to protect against them. There are specialist organisations that are able to provide threat intelligence as a service.

Threat intelligence assesses and validates information from a range of sources on current and potential threats, analysing trends and highlighting security issues that are particularly relevant to the organisation. Good threat intelligence should provide context to enable an organisation to make informed decisions on how to protect itself by understanding who is potentially attacking them (the threat actors), the motivations and capabilities of the threat actors, and how their systems could be compromised. Properly applied, it can help an organisation to stay up to date with the often overwhelming volume of threat information, which can then provide greater insight into the threats, enabling a more targeted response with better deployment of resource.

Threat intelligence solutions generally use machine learning and security analytics for Big Data to automate data collection and processing, taking unstructured raw data from multiple and disparate sources about emerging or existing threat actors and threats. The vast amount of information collected means it can be difficult to see the wood for the trees. Therefore, the information needs to be interpreted by experienced and trained individuals to ensure biases, deceptions and uncertainties are identified and managed, as it relies on deploying a rigorous way of thinking and structured analytical techniques. The analysis has to look at the threat actors, their intent and their capabilities, their tactics, techniques and procedures (TTPs), motivations, and access to the intended targets. The analysed data then need to be filtered to produce threat intelligence feeds and management reports on predictions, priority of threats and methods of attack to enable the organisation to make informed, strategic, operational and tactical decisions to protect itself.

Threat intelligence is not a one-off activity. It needs to be ongoing and cyclical to redefine requirements as situations change. It can inform threat modelling activities by identifying new vulnerabilities and threat agents and form part of the overall risk management process.

Cyber threat intelligence

The term cyber threat intelligence (CTI) is really a subsection of the threat intelligence that relates specifically to the internet. The term is used to describe the identification, collection, collation and analysis of information gained from third parties such as:

- nation state intelligence and security agencies;
- commercial threat intelligence services;
- Open Source Intelligence (OSI) and other sources, including social media;
- the internet;
- the dark web in particular.

This information can be used to help identify and quantify the internet-based threats to the organisation from nation states, organised crime, hacktivists, journalists and hackers in general.

The information gained can be very useful in helping to justify recommendations for security policy, process and controls together with the budget required to implement and operate them. It can also provide input into an ongoing cyber security awareness programme, highlighting any new threats of which staff need to be vigilant. The data may also be shared with organisations and third-party service providers with whom the organisation works to ensure they are warned and alert for any activity that might affect them or that may be used as a conduit to attack the organisation's own systems.

There are open source tools and techniques that can be used to acquire data, but these do require training and experience to be effective. For higher-threat environments, specialist third-party CTI organisations (e.g. Digital Shadows and Recorded Future) can be contracted to provide information where appropriate.

Threat intelligence tools can work with SIEM tools to provide alerts that indicate unauthorised activity happening within the enterprise. These alerts can also come from other threat intelligence sources. Whatever the source, they are classed as indicators of compromise (IoCs) that require further action. This starts with analysis to verify that it is an actual incident and then to work on the determination of what has happened.

These IoCs take various formats that are not standard. This means that effective threat intelligence requires organisations to understand their own threat environment and have a human element to process and assess the IoCs as part of the process to confirm or deny their relevance.

It can be challenging only to collect data that are relevant, to avoid overloading the tools and to allow effective analysis for threat identification and response. The outputs can be used by automated toolsets and analysts to search for evidence of an attack

or successful intrusion. There are standards, such as OpenIOC, which can be used to produce a standard format of CTI that can be shared internally and with third parties.

Threat intelligence also has a vital role to play during incident response. The function can:

- Act as the interface with external specialist law enforcement and intelligence agencies, allowing the incident response team to concentrate on dealing with the incident. An agreed communication plan will define what information about an incident can be shared with authorised bodies during and after the event.
- Conduct rapid research into any successful attack form to identify sources, impacts and ways of mediating, controlling and stopping the attack to help the incident response teams.
- Perform specialist analysis that requires human interpretation and conjecture. Automation, analytics and various tools can drastically increase the effectiveness of analysts, but there must always be analysts driving and controlling that process.
- Liaise with other organisations and third-party service providers with whom the organisation works to ensure they are warned and alert for any activity that might affect them as a result of the incident.

As well as sharing CTI, it is important to share and be aware of any vulnerabilities that may be found in operating systems, firmware, networks, software and hardware assets that are in use to conduct business operations and deliver goods and services. New vulnerabilities are being found every day and, in most cases, patches are then issued. The standard source of information for these is the Common Vulnerabilities and Exposures (CVE) database, hosted by the USA's NIST. This database can be searched for vulnerabilities in specific products, so it is possible to build up a list for an organisation's own assets to allow easy discovery of any new vulnerabilities and patches that need to be considered for deployment to manage the risk of exploitation.

Other sources of information may well include product vendors and manufacturers, especially if a support agreement is in place. Academic journals and conference papers (e.g. Black Hat and RSA) might also reveal new vulnerabilities, and there are numerous blogs and newsfeeds that contain useful vulnerability information. It is important to remember that the bad guys also monitor the various sources, looking for new vulnerabilities that they can exploit, so checks should be made frequently and prompt action taken when considered necessary. Having an effective patching/update process is essential.

Co-operative threat intelligence

Most organisations, large or small, now have partners of one type or another ranging from the basic outsourced suppliers through to strategic partnerships to run specific operations or handle major pieces of work. These partnerships should be seen as an excellent way to increase the sharing of threat intelligence such that all parties can reduce their overall risk of a successful attack being perpetrated against them. Threat intelligence sharing can take many forms, from the simple passing of information that might be helpful through to the more formal Cyber Security Information Sharing

Partnership (CiSP) organisations set up in the UK and elsewhere to enable organisations to share relevant cyber security information within their own particular industry sector.

An organisation needs to set up a two-part threat intelligence facility. The first part will be to gather and collate relevant information from their own internal and external sources and related areas, including receiving information from others. This information must then be analysed and assessed to ensure that it is correct (factually), that it is current, that it has credibility and relevance, and that the supplier of the information (where relevant) is trustworthy. There then needs to be a facility to share it with their partners, either those with whom they have a contractual agreement or through the CiSP type mechanism. When done effectively, this can reduce the time it takes an organisation to deal with a threat, real or potential, very significantly. They can begin to address proactively the potential threats and so reduce the chance of real attacks being successfully completed.

In the UK there are now a number of CiSP organisations established and working very effectively in close partnership with the NCSC and other government bodies, and more are being set up as their value becomes more evident. They provide the opportunity of getting a little ahead of the attackers while also learning from the mistakes of their colleagues in the same business sector. The system has a strong ethical boundary such that the originator of any information that is shared is never identified, so preventing any form of 'blame game' being generated. CiSP is sometimes referred to as part of the UK government's Active Cyber Defence strategy, and evidence shows it can be a very effective way of reducing successful attacks on organisations large and small.

If we consider GANT's situation, there is a great deal of information contained in the computer system that could cause problems if it were to fall into the wrong hands. As we have seen in Chapter 1, there are two areas of particular interest:

1. Information currently kept openly on the GANT website identifies the locations where the Natterjack toad breeds and lives, which could be of interest to property developers and illegal collectors.
2. Financial information about the members' fees and grants provided is also kept on GANT's computer and, while disclosure of this information itself might not be a major issue, it could render members and sponsors vulnerable to intimidation, as some of their personal details (aside from those that they wish to be made public) might become openly available.

At the same time, GANT does need to publicise its work through its website, and clearly there must be some form of segregation of information that is available through this – keeping the public information openly available, while securing the confidential information and allowing bona fide members to access it through secure means.

It is also necessary from time to time for members (especially committee members) to access and exchange this information electronically using remote access over the internet.

ACTIVITY 9.1

Describe the steps GANT might take in order to protect its databases from unauthorised access and alteration beyond those afforded by simple password protection of the user area on their computer system.

Suggest how a commercial product such as PGP, TLS or IPSec might be used in order to secure sensitive information passing between GANT committee members.

CONCLUSION

Throughout this book every effort has been made to give you enough information to help you to reach a level of understanding and knowledge that will prove to be an excellent basis for taking the BCS's Information Security Management Principles foundation examination. The authors have also tried to provide a general grounding for you in the fundamentals of the subject. It is not intended that this should be the only book you study if you are to take up a role in information assurance – far from it – but it is hoped that it whets your appetite for the immensely complex, interesting and rewarding field of information management and its assurance.

If you choose to go on to more demanding and detailed areas of study or employment, you will find a vast array of texts available that cover all of the areas in this book in much more detail. If your interest is less detailed, then it is hoped that this book has provided enough information to allow you to make appropriate decisions on your activities, further study, preparations and possible career in information assurance. If you have responsibility for all or some of the aspects covered here, you may well find it useful as an introduction to help to sell the concepts to senior managers.

In all your work, the authors wish you well and hope that all your preparations prevent the worst from happening – that is the best they can offer you.

SAMPLE QUESTIONS

1. **Symmetric key encryption systems are those in which:**
   a. Sender and recipient have completely unrelated encryption and decryption keys.
   b. Sender and recipient both share the same encryption and decryption keys.
   c. Sender and recipient have different but mathematically related encryption and decryption keys.
2. **Production of a message digest enables the recipient to:**
   a. Verify the integrity of the message content and authenticate the sender.
   b. Verify the integrity of the message content only.
   c. Authenticate the sender only.

3. **Asymmetric key encryption is not generally used for encrypting large messages because:**
   a. It only works on very short message lengths.
   b. It is less secure than symmetric key encryption.
   c. It takes much longer to carry out the encryption and decryption processes.

REFERENCES AND FURTHER READING

There are numerous books and papers on cryptography, and it is impossible to list them all here. Instead, the authors have identified some of the key publications and websites dealing with the topic.

One of the most illuminating introductions to cryptography is *Cryptography: A Very Short Introduction* by Professors Fred Piper and Sean Murphy of Royal Holloway, University of London. A rather heavier tome is *The Codebreakers* by David Kahn, which was originally published in 1997, so does not contain information on more recent developments in cryptography, but remains an excellent reference work covering the history of cryptography and that of cryptanalysis, especially during the First and Second World Wars. Another highly readable book, *The Code Book* by Simon Singh, was serialised on television and contains a ten-stage cryptography challenge. Finally, for those who wish to explore present-day cryptography in greater detail, there is *RSA Security's Official Guide to Cryptography*.

Publications

Anderson, R. (2001) *Security Engineering: A Guide to Building Dependable Distributed Systems*. John Wiley & Sons.

Beker, H. and Piper, F. (1982) *Cipher Systems*. Van Nostrand.

Bernstein, P.L. (1998) *Against the Gods: The Remarkable Story of Risk*. John Wiley & Sons.

Burnett, S. and Paine, S. (2001) *RSA Security's Official Guide to Cryptography*. McGraw-Hill.

Diffie, W. and Hellman, M.E. (1976) New directions in cryptography. *IEEE Trans. Inform. Theory*, 22 (6). 644–654. https://cs.jhu.edu/~rubin/courses/sp03/papers/diffie.hellman.pdf

Diffie, W. and Landau, S. (1998) *Privacy on the Line: The Politics of Wiretapping and Encryption*. MIT Press.

Ford, W. and Baum, M.S. (1997) *Secure Electronic Commerce: Building the Infrastructure for Digital Signatures and Encryption*. Prentice Hall.

Kahn, D. (1997) *The Codebreakers: The Comprehensive History of Secret Communication from Ancient Times to the Internet*. Scribner.

Martin, Keith (2012) *Everyday Cryptography: Fundamental Principles and Applications*. Oxford University Press.

Menzies, A.J., van Oorschot, P.C. and Vanstone, S.A. (1996) *Handbook of Applied Cryptography*. CRC Press. (see also under Websites).

Piper, F and Murphy, S. (2002) *Cryptography: A Very Short Introduction*. Oxford University Press.

Rivest, R.L., Shamir, A. and Adleman, L. (1978) A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21. 120–126. https://people.csail.mit.edu/rivest/pubs/RSA83a.pdf

Salkind, N.J. (2004) *Statistics for People Who (Think They) Hate Statistics*. SAGE Publications.

Schneier, B. (2015) *Applied Cryptography: Protocols, Algorithms and Source Code in C*. John Wiley & Sons.

Singh, S. (1999) *The Code Book: The Science of Secrecy from Ancient Egypt to Quantum Cryptography*. Fourth Estate.

Smith, R.E. (1997) *Internet Cryptography: Evaluating Security Techniques*. Addison Wesley.

Websites

The Handbook of Applied Cryptography website (from which it is possible to download the book in pdf format): https://cacr.math.uwaterloo.ca/hac/

RSA Security Labs website with, among other things, a cryptography FAQ in the 'Historical' tab: https://rsa.com/rsalabs

The US National Institute of Standards and Technology (NIST) cryptography website: https://csrc.nist.gov/

APPENDIX A:
INFORMATION SECURITY STANDARDS RELEVANT TO CISMP, PCIRM AND PCBCM EXAMINATIONS

BUSINESS CONTINUITY STANDARDS

ISO/IEC 27031:2011 – Information technology – Security techniques – Guidelines for information and communication technology readiness for business continuity

PD 25111:2010 – Business continuity management – Guidance on human aspects of business continuity

PD 25222:2011 – Business continuity management – Guidance on supply chain continuity

PD 25666:2010 – Business continuity management – Guidance on exercising and testing for continuity and contingency programmes

ISO 22301:2014 – Societal security – Business continuity management systems – Requirements

ISO 22313:2014 – Societal security – Business continuity management systems – Guidance

ISO 22318:2015 – Societal security – Business continuity management systems – Guidelines for supply chain continuity

ISO 22322:2015 – Societal security – Emergency management – Guidelines for public warning

BS ISO 22324:2015 – Societal security – Emergency management – Guidelines for colour-coded alerts

BS 11200:2014 – Crisis management – Guidance and good practice

The Business Continuity Institute Good Practice Guidelines 2018 – The global guide to good practice in business continuity: https://thebci.org

DATA PROTECTION STANDARDS

BS 10012:2017 – Data protection – Specification for a personal information management system

UK Data Protection Act 1998: https://opsi.gov.uk/acts/acts1998/ukpga19980029en1

General Data Protection Regulations (GDPR): https://eugdpr.org

RISK MANAGEMENT STANDARDS

Institute of Risk Management's 'A Risk Management Standard': https://www.theirm.org/media/4709/arms_2002_irm.pdf

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

ISO/IEC 27001:2017 – ISMS – Information technology – Security techniques – Specification for an Information Security Management System

ISO/IEC 27002:2017 – ISMS – Information technology – Security techniques – Code of practice for information security controls

ISO/IEC 27005:2018 – Information technology – Security techniques – Information security risk management

ISO Guide 73:2009 – Risk management – Vocabulary – Guidelines for use in standards

ISO 31000:2018 – Risk management – Principles and guidelines

ISO/IEC 31010:2010 – Risk management – Risk assessment techniques

UK PRIMARY LEGISLATION

Official Secrets Act 1989

Computer Misuse Act 1990

Freedom of Information Act 2000

Regulation of Investigatory Powers Act (RIPA) 2000

Communications Act 2003

The Police and Criminal Evidence Act 1984 (Codes of Practice) Order 2008

Data Retention and Investigatory Powers Act 2014

Data Protection Act 2018

General Data Protection Regulation (GDPR) 2018

INFORMATION SECURITY STANDARDS

ISO/IEC 13335-1:2004 – Information technology – Security techniques – Management of information and communications technology security – Part 1: Concepts and models for information and communications technology security management

ISO/IEC 15408-1:2011 – Information technology – Security techniques – Evaluation criteria for IT security – Part 1: Introduction and general model

ISO/IEC 15408-2:2008 – Information technology – Security techniques – Evaluation criteria for IT security – Part 2: Security functional components

ISO/IEC 15408-3:2008 – Information technology – Security techniques – Evaluation criteria for IT security – Part 3: Security assurance components

ISO 15489-1:2016 – Information and documentation – Records management – Part 1: General

ISO/IEC 27000:2018 – Information technology – Security techniques – Information security management systems – Fundamentals and vocabulary

ISO/IEC 27001:2017 – ISMS – Information technology – Security techniques – Specification for an Information Security Management System (this replaces BS 7799 Part 2)

ISO/IEC 27002:2017 – Information technology – Security techniques – Code of Practice for Information Security Management (this replaces BS 17799)

ISO/IEC 27003:2017 – Information technology – Security techniques – Information security management system implementation guidance

ISO/IEC 27004:2016 – Information technology – Security techniques – Information security management – Measurement

ISO/IEC 27005:2018 – Information technology – Security techniques – Information security risk management

ISO/IEC 27006:2015 – Requirements for bodies providing audit and certification of information security management systems

ISO/IEC 27007:2017 – Information technology – Security techniques – Guidelines for information security management systems auditing

ISO/IEC 27008:2019 – Information technology – Security techniques – Guidelines for auditors on information security controls

ISO/IEC 27010:2015 – Information technology – Security techniques – Information security management for inter-sector and inter-organisational communications

ISO/IEC 27011:2016 – Information technology – Security techniques – Information security management guidelines for telecommunications organizations based on ISO/IEC 27002

ISO/IEC 27013:2015 – Information technology – Security techniques – Guidance on the implementation of ISO/IEC 27001 and ISO/IEC 20000–1

ISO/IEC 27014:2013 – Information technology – Security techniques – Governance of information security

ISO/IEC 27015:2012 – Information security management systems – Information security management guidelines for financial services

ISO/IEC 27016:2014 – Information technology – Security techniques – Information security management – Organisational economics

ISO/IEC 27017:2015 – Information technology – Security techniques – Code of practice for information security controls based on ISO/IEC 27002 for cloud services

ISO/IEC 27018:2014 – Information technology – Security techniques – Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors

ISO/IEC 27019:2013 – Information technology – Security techniques – Information security management guidelines based on ISO/IEC 27002 for process control systems specific to the energy utility industry

ISO/IEC 27023:2015 – Information technology – Security techniques – Mapping the revised editions of ISO/IEC 27001 and ISO/IEC 27002

ISO/IEC 27031:2011 – Information technology – Security techniques – Guidelines for information and communication technology readiness for business continuity

ISO/IEC 27032:2012 – Information technology – Security techniques – Guidelines for cybersecurity

ISO/IEC 27033-1:2015 – Information technology – Security techniques – Network security – Part 1: Overview and concepts

ISO/IEC 27033-2:2012 – Information technology – Security techniques – Network security – Part 2: Guidelines for the design and implementation of network security

ISO/IEC 27033-3:2010 – Information technology – Security techniques – Network security – Part 3: Reference networking scenarios – Threats, design techniques and control issues

ISO/IEC 27033-4:2014 – Security techniques – Network security – Part 4: Securing communications between networks using security gateways

ISO/IEC 27033-5:2013 – Security techniques – Network security – Part 5: Securing communications across networks using Virtual Private Networks (VPNs)

ISO/IEC 27033-6:2016 – Security techniques – Network security – Part 6: Securing wireless IP network access

ISO/IEC 27034-1:2011 – Information technology – Security techniques – Application security – Part 1: Overview and concepts

ISO/IEC 27034-2:2015 – Information technology – Security techniques – Organisation normative framework

ISO/IEC 27034-3:2018 – Information technology – Security techniques – Application security management process

ISO/IEC 27034-5:2017 – Information technology – Security techniques – Protocols and application security controls data structure

ISO/IEC 27034-6:2016 – Information technology – Security techniques – Case studies

ISO/IEC 27034-7:2018 – Information technology – Security techniques – Assurance prediction framework

ISO/IEC 27035-1:2016 – Information technology – Security techniques – Information security incident management – Principles of incident management

ISO/IEC 27035-2:2016 – Information technology – Security techniques – Information security incident management – Guidelines to plan and prepare for incident response

ISO/IEC 27036–1:2014 – Information technology – Security techniques – Information security for supplier relationships – Overview and concepts

ISO/IEC 27036–2:2014 – Information technology – Security techniques – Information security for supplier relationships – Requirements

ISO/IEC 27036–3:2013 – Information technology – Security techniques – Information security for supplier relationships – Guidelines for information and communication technology supply chain security

ISO/IEC 27037:2012 – Information technology – Security techniques – Guidelines for identification, collection, acquisition and preservation of digital evidence

ISO/IEC 27038:2014 – Information technology – Security techniques – Specification for digital redaction

ISO/IEC 27039:2015 – Information technology – Security techniques – Selection, deployment and operations of intrusion detection systems (IDPS)

ISO/IEC 27040:2015 – Information technology – Security techniques – Storage security

ISO/IEC 27041:2015 – Information technology – Security techniques – Guidance on assuring suitability and adequacy of incident investigative method

ISO/IEC 27042:2015 – Information technology – Security techniques – Guidelines for the analysis and interpretation of digital evidence

ISO/IEC 27043:2015 – Information technology. Security techniques. Incident investigation principles and processes

ISO/IEC 24762:2008 – Information technology – Security techniques – Guidelines for information and communications technology disaster recovery services

ISO 38500:2015 – Corporate governance of information technology

ISO/IEC 18028-1:2006 – Information technology – Security techniques – IT network security – Part 1: Network security management

ISO/IEC 18028-2:2006 – Information technology – Security techniques – IT network security – Part 2: Network security architecture

ISO/IEC 18028-3:2005 – Information technology – Security techniques – IT network security – Part 3: Securing communications between networks using security gateways

ISO/IEC 18028-4:2005 – Information technology – Security techniques – IT network security – Part 4: Securing remote access

ISO/IEC 18028-5:2006 – Information technology – Security techniques – IT network security – Part 5: Securing communications across networks using virtual private networks

ISO/IEC 17788:2014 – Information technology – Cloud computing – Overview and vocabulary

ISO/IEC 17789:2014 – Information technology – Cloud computing – Reference architecture

ISO/IEC 29100:2011 – Information technology – Security techniques – Privacy framework

ISO/IEC 29101:2013 – Information technology – Security techniques – Privacy architecture framework

ISO/IEC 29147:2014 – Information technology – Security techniques – Vulnerability disclosure

ISO/IEC 29190:2015 – Information technology – Security techniques – Privacy capability assessment model

ISO/IEC 30111:2013 – Information technology – Security techniques – Vulnerability handling processes

Good Practice Guide for Digital Evidence (V5) (NPCC formerly ACPO): https://www.app.college.police.uk/app-content/investigations/forensics/

The Information Security Forum Standard of Good Practice: https://www.securityforum.org/?page=downloadsogp

British standards may be obtained in PDF or hard copy formats from the BSI online shop: www.bsigroup.com/Shop

ISO standards may also be obtained through the BSI, or directly from the ISO online shop: www.iso.org/iso/store.htm

ACTIVITY SOLUTION POINTERS

CHAPTER 1

Activity 1.1

There are a significant number of threats, vulnerabilities and risks to this organisation. You may have come up with others, but here are three of the most serious ones. It is most important that you fully appreciate the differences between the three categories as well as being able to make some specific suggestions.

Three threats
These are areas where there is a potential for some adverse consequences if this threat should arise. In this scenario three threats might be as follows.

1. Information about members might be accessed by unauthorised people.
2. Information about the habitats of the Natterjack toad might be used by those who are not inclined to support its ongoing existence.
3. The website might be compromised with unofficial messages added to it.

Three vulnerabilities
These are weaknesses in the system that might allow a threat to materialise. In this scenario and building on the threats given above, the vulnerabilities might be as follows.

1. The records of the members are maintained in a variety of ways including paper and unreliable computer systems.
2. The information about the toads' habitats is maintained on an old internet-based server with very limited assurance in place.
3. There is no firewall between the website server and the internet.

Three risks
There is a large number of risks resulting from the threats and vulnerabilities listed above. Three of them might be as follows.

1. There is a risk that unscrupulous property developers might gain access to the personal details of members of GANT and take action against them or their property.
2. There is a risk that a habitat of the Natterjack toad might be destroyed by someone who is not interested in the existence of the animal.

3. There is a risk that someone might gain access to the code of the GANT website and change the messages to information that is offensive to those interested in nature conservancy.

Activity 1.2

The cost-effectiveness or cost–benefit analysis for such an implementation would include many areas. Three of the most significant, following on from the suggestions given above for Activity 1.1, might be the following.

1. Members of GANT could be injured, or their families and property adversely affected in some way. The cost of protecting the members and their families would be excessive and could not be found through the membership of GANT alone.
2. The cost of reintroducing Natterjack toads into the wild after their habitat has been destroyed would be very considerable. This could be the consequence (impact) of allowing unauthorised access to the details of the toads' habitats.
3. GANT relies very heavily on the goodwill of other nature conservancy groups and donations from interested commercial companies. If they were embarrassed by the content of the website, they might reduce or withdraw their support for an organisation they saw as unprofessional and poorly organised. This could be devastating for the existence of GANT.

CHAPTER 2

Activity 2.1

Given several types of threat such as loss of membership records, details of the Natterjack toad's breeding grounds or the financial records of GANT, an impact analysis should look at the short-term, medium-term and long-term consequences.

You should also try and look at various types of impact – including financial loss, GANT's reputation and the possible effects on the toads themselves. You may find that the impacts may change over time – improving or worsening according to the type of impact, and this may (later) affect your recommendations as to how the risks might be mitigated.

Activity 2.2

Whereas impact analyses can be relatively straightforward to conduct, the likelihood assessments can be more complex as they can be very subjective if using a qualitative approach. To gain a less-subjective assessment, you should carry out the risk assessment using a quantitative approach, which may involve gathering statistical information such as frequency of previous events, for example, which can sometimes be a time-consuming process.

For the purposes of this activity it is suggested that you take a qualitative approach and use your best judgement to come up with a low/medium/high likelihood rating.

You can number each threat and mark the numbers on the completed risk matrix diagram. You then need to decide what action(s) to take for each threat, and it might be useful to make a note of why you consider this to be the optimum approach – in the real world you would have to justify your recommendations, possibly with a cost-based business case, so it is worth getting into the habit early on.

Try to allow logic to influence your recommendations rather than emotions. You may feel strongly about a particular course of action, but others may hold a different view, and you might have to put up a convincing business-based argument to bring them round to your point of view.

CHAPTER 3

Activity 3.1

The plan should include the following components.

- A senior member of staff, for example Ms Jackson, should be given board member responsibility for information assurance and provide high-level sponsorship for assurance.
- Someone should be given the role of day-to-day co-ordination of information assurance across GANT – full-time responsibility. Suggest yourself.
- Detail what you will be able to deliver within the assurance function, for example:
  - co-ordinating assurance activities across GANT;
  - providing advice and guidance on assurance;
  - producing a security policy;
  - monitoring the effectiveness of assurance controls;
  - reporting on the effectiveness of controls;
  - raising awareness about assurance within GANT and ensuring that people understand their responsibilities.
- Identify people who can help to support you within GANT, decide what they need to do and request that assurance responsibilities are built into the roles.

Activity 3.2

An end-user code of practice for GANT could include statements on:

- access to systems;
- protection of passwords;
- leaving information or systems holding information unattended;
- measures required to protect information about GANT members;
- protection of information and equipment if taken out of the office;

- acceptable behaviour when using GANT systems;
- use of the internet;
- any particular requirements for the protection of PII;
- use of GANT systems for personal use.

Activity 3.3

The following types of activities will help Ms Jackson to demonstrate to auditors and regulators that assurance is being managed effectively:

- establishing a governance process with regular governance reviews that can be chaired by Ms Jackson and supported by a security forum;
- setting up a risk register to record GANT's information risks and documenting how the risks are being treated;
- defining a schedule to show that assurance is regularly reviewed, for example compliance and policy reviews;
- setting up a documentation library to demonstrate that assurance review work and planning has been completed and formally recorded. The library could contain current security polices, outputs from risk assessment, results of audits and compliance reviews, minutes from governance reviews, GANT's risk register, risk assessments, incident reports, dispensations to policy and so on.

Activity 3.4

A sensible approach would be:

1. To carry out a high-level risk assessment to identify where the main risks are and which areas cause the most concern and why. This should involve major stakeholders such as Dr Peabody and Ms Jackson.
2. To develop a high-level plan, including timescales and effort, to address the most pressing issues quickly with tactical improvements and longer-term strategy to improve control overall. This should be supported by a business case that communicates the benefits, including outline costs.
3. To produce a simple presentation in business terms to communicate this information to Ms Jackson so that she can decide on what work should take place.
4. Once agreed, a more detailed plan can be put in place to carry out the work and be regularly reviewed to check progress.

Activity 3.5

As GANT is a small organisation, a simple process for managing assurance breaches will be adequate. The main elements should include:

- staff (with deputies) should be nominated to manage assurance breaches;

- a procedure for reporting, recording and managing incidents should be developed so that they can be dealt with in a consistent way;
- all staff and members should be told who they should contact and what they should do if they suspect a breach;
- personnel responsible for managing incidents should be trained to understand how they should deal with potential incidents and when they should engage with either specialist third parties or law enforcement agencies.

Activity 3.6

To ensure that GANT complies with appropriate personal information legislation, the following types of activities should be carried out:

- review local legislation and how that applies to GANT;
- review whether personal information is passed to any other legal jurisdiction and if so what their requirements are;
- understand if local laws require that the holding of personal information is registered and if so is the registration up to date and accurate;
- identify if there a policy in place to specify how personal information should be held and check if it is up to date;
- identify what personal information is collected by GANT and the use made of it;
- identify which systems are holding personal information;
- identify what controls are in place to protect personal information being held;
- identify who has access to the information and check they understand their responsibilities;
- identify if any information is shared with third parties and if so what controls are in place;
- understand what monitoring takes place and does this comply with local legislation;
- identify whether the enterprise has communicated to individuals what monitoring takes place;
- document finding where there are compliance issues and make recommendations to resolve;
- discuss findings and recommendations with a legal expert.

CHAPTER 4

Activity 4.1

The report's specification must include:

- A requirement for data that provides evidence of performance against SLAs.
- Timescales for reporting and regular deadlines to ensure timeliness of reports.

- Standards for reporting incidents.
- Compliance with necessary privacy legislation:
  - escalation procedures for both parties;
  - any level of protective marking for such a potentially sensitive document.

Activity 4.2

The internal audit process will require the checking of the use of the controls and policies against actual practice. A good tip is to pick a process and follow it through from start to finish, talking to people who actually use it regularly, and also looking at audit and event data to make sure that all the elements agree.

Activity 4.3

The most important point is to be able to follow the guidelines for the collection of evidence in such a way that it is admissible in a court or tribunal. These are contained in guidance issued by NPCC in the UK. The process often requires some specialist hardware and software.

For a small organisation like GANT, probably the most important part of the plan is to state that you would bring in an external consultant with specialist knowledge in how to do this, or have a contract in place with an organisation that can do it for you.

Another important aspect is to make sure you know the activities about which you are legally required to notify the police:

- suspected paedophile images or activity;
- terrorism;
- danger to life of an individual.

Activity 4.4

There are a series of mandatory and desirable criteria, which should be listed in the invitation to tender (ITT) document. Tenderers should be required to answer all these questions as part of their initial submissions:

Mandatory

- Supplier to agree to no-notice quality and security audits by GANT.
- All source code and development material to be placed in escrow in case of business failure by the supplier.
- All staff who will work on the project to have been vetted for honesty and reliability.
- Supplier to agree to sign a legally binding data sharing protocol as part of the contract.
- Supplier to have suitable indemnity insurance.

- Supplier to provide three referees for previous work of this type.
- Deliverable product and all outputs capable of being inspected by standard malware scanners.

Desirable

- Supplier has achieved level 5 on the CMM scheme.
- Supplier has achieved a quality accreditation such as ISO 9001.
- Supplier is accredited to ISO/IEC 27000 series for information security.

Activity 4.5

Before any kind of change can be made, the suggested modification must be formally reviewed and approved by a change management team. This is in order to protect the existing information assets and business processes against any adverse impact to service delivery.

If the request is approved, the users will be asked to help design the change. The work will be done on the development system by the developers and then the users will be asked to test the new and existing functionality to make sure there are no unexpected results.

Having said all this, in a properly configured information security architecture, the users would not have the necessary privileges or access to the development tools and source code to make changes. They would have to go to the development and system administration teams.

They should check that the change will not be carried out during a peak processing period (such as year-end) or critical commercial window (such as Christmas trading for a retailer).

Activity 4.6

There are several rating and assessment schemes for the security of IT products. The level of product accreditation required is proportional to the impact on GANT and the toads themselves if there were an issue with the information assets. It is not anticipated that a high level of accreditation will be required in this instance.

The second point to note is that, while the level of assurance goes up from using such products, the range of choice goes down and the price often goes up. This is because not all manufacturers go through the assessment process, and those that do need to recoup the costs involved. It may be that an unacceptable proportion of the grant would be consumed by the assured products and not the application, but only a quick review of the requirements and available products would confirm this.

The final point to make is that one should never rely on the claims of the supplier to have met a particular standard. This should always be verified with the organisation that performed the assessment and the website for that particular standard to make sure that they are listed as having passed.

CHAPTER 5

Activity 5.1

The need for information assurance is not intuitive. People need to be made aware of the risks and threats, together with some basic guidance on what to do. Make sure that the following are in place:

- security policy document, defining what to do;
- relevant terms and conditions in contracts of employment;
- a clear and easy-to-follow acceptable use policy;
- material for assurance awareness training;
- code of conduct for staff and volunteers;
- responsibilities for protecting personal information.

The organisation is starting to grow and now is the time to define the information assurance culture that will be of great benefit to GANT in the years to come.

Activity 5.2

Remember that transferring the work to a third party does not absolve GANT from responsibility for the assurance of their data. GANT is as culpable as the third party if a breach occurs and GANT can be shown not to have exercised their duty of care, often known as due diligence, in ensuring that the third party:

- was aware of its obligations;
- implemented and maintained appropriate countermeasures;
- submitted to periodic compliance audits;
- ensured that their suppliers are also adhering to these requirements.

Activity 5.3

The policy documents need to include a specification for access control and minimum requirements for identification and authentication:

- for users in the office;
- for any remote access users.

This should take into account any data from risk assessments performed for GANT.

The documents should also define requirements for:

- separation of roles and responsibilities – for example, users and administrators;
- implementation of user groups based on roles, often grouped by job function;
- enrolment of new users and deletion of access rights of those who leave.

Activity 5.4

The first task is to identify all repositories of data and work outwards to identify the means by which they are accessed internally and externally (if any). This should include any cloud-based platforms being used.

The next task is to decide how to manage the risks each of these means of access brings to GANT.

Activity 5.5

1. The initial training campaign should focus on confidentiality of information generally. All staff should understand the importance of protecting members' information and the possible threats to the organisation in terms of sabotage, theft and so on. Specific training should be given to any staff handling personal records so that they can understand their specific responsibilities in protecting this type of information. The types of messages that could be included are:
   - responsibilities when handling information about people;
   - enterprise's requirements for handling information, especially about people;
   - the importance of password protection and maintaining a clear desk policy;
   - protection for information that has to be taken out of the office.
2. Face-to-face training would be the preferred option as the enterprise is small, probably with one-to-one training within the UK with Dr Peabody and any other staff handling membership or other confidential information. Producing an induction pack could ensure that any new personnel are made aware of their assurance responsibilities. Producing some posters to be put on the wall would also be helpful to remind people to be careful about protecting GANT information.

CHAPTER 6

Activity 6.1

The important thing to remember is the need to balance the risk of malware against the costs of purchasing and implementing countermeasures. A good base set of recommendations would be as follows:

- Have combination anti-virus and personal firewall programs on each PC that have a connection to the internet or are networked to another computer that does. Choose a product that will also restrict the executables that will run on that computer to a known list of authorised products.
- Change all the default settings in operating systems, applications and browsers, for example passwords, configurations, open ports and so on, to make it harder for malware to compromise the computer.
- Apply patches to the operating system and applications promptly.

- Educate users about the threat and some of the tactics used.
- Have a backup policy, make backups regularly and then test them regularly.
- Consider a firewall with intrusion detection or intrusion prevention capability for a network connected to the internet.

Activity 6.2

All information that GANT holds should be considered for encryption, but in particular the membership details should be encrypted since that is classed as PII. Other sets of information to be considered would include the details of the toad habitats, all financial information, particularly when it includes details such as bank account numbers, and any sensitive information regarding studies, potential developments or other aspects of looking after the toads.

The benefits of encryption should be shown under the four main areas of:

- protection of confidentiality or secrecy;
- data integrity – preventing unauthorised changes to information;
- user verification – ensuring only those with the appropriate need have access to the information;
- proving the identity of someone who has made changes to a set of information.

Activity 6.3

Start by explaining that it is like putting a door with a lock in what is currently a solid wall round the systems of GANT. It allows access to people who have the right key, but that lock can also be picked by a skilful attacker.

It is also worth explaining that it is possible for sensitive data to end up being saved on the home PC, which is not as well protected as the office systems and therefore is a weak spot vulnerable to attack. A network usage policy or technology (or both) needs putting in place to ensure that this does not happen.

Another important point is to explain that the connection can be eavesdropped, just like a telephone conversation, and the data passing backwards and forwards can be copied. That might include usernames and passwords. A means of protecting the traffic and login data must be put in place, preferably involving some kind of one-time password system such as a token.

Activity 6.4

The first thing to recognise is that this is usually a job for an expert, so consider asking for some outside help or training. Like a firewall, these are tricky systems to configure properly and are of little value unless working effectively. Once you have the data, you then need to be able to understand what they are telling you.

The most important thing is to identify what is and is not allowed to happen on the network and define these as rules in the IDS. The second thing is to identify the points in the network where monitoring is best used. Obvious locations are any connection to external networks and the internet, and any point used to separate networks where an attacker might be trying to gain access to a partitioned area containing sensitive data.

Activity 6.5

Actions include the following:

- Identify the most cost-effective means of linking the two organisations, depending on the number of users, volume of data and the frequency with which the connection will be used. The options include asymmetric digital subscriber line (ADSL) (broadband), VPN and leased line.
- Decide how to secure the connection:
  - How should each user identify and authenticate themselves to the other organisation?
  - What kind of protection do the data travelling across the connection need?
  - What kind of data sharing protocol is needed for legal purposes (e.g. Data Protection Act or General Data Protection Regulation)?

Activity 6.6

The forum will have the same kind of threats as email and websites, in that it is possible to hide malware in the messages and it is also possible for an attacker to try and take control of the website itself, to hide malware in the code for the pages. You recommend that the forum should not be made available without appropriate countermeasures being in place to screen any postings to it for malware. You recommend that specialist advice be sought from an expert in this field.

Activity 6.7

The DPA and GDPR can be complex to understand. It is important to read through all the material, work out which parts apply to GANT and its partners and then write a protocol to govern the data sharing with the third parties. Get the document checked by a lawyer who has sound knowledge and understanding of the DPA and GDPR before using it.

Activity 6.8

The video-conferencing system will provide an easy access route into the GANT data network unless careful thought is given to the protection of the link. The protocols allowed through from the system must be tightly controlled and the link may need monitoring with an intrusion detection or prevention system of some sort to defeat attempts to hack into the network.

Activity 6.9

Answers should include:

- The nature of cloud environments vary, and this may mean that inadequate security controls are in place or data may be held in undesirable jurisdictions.
- The importance of understanding the contractual conditions as they may result in a loss of control over GANT's information, its use and its disclosure.
- The potential reputational and legal risks to GANT if a security breach occurred.

Activity 6.10

Some key issues that should be covered include:

- What controls do they have in place to protect your information?
- Where will information be stored and in which countries and do you have any control over jurisdictions where information is stored?
- Will the cloud provider have any ownership or disclosure rights over your information?
- Do they involve third parties and what controls do they have in place to ensure that information remains fully protected?
- Will you have the right to audit their services?
- Are you able to opt out of any changes to service?
- What termination arrangements are in place?
- Can they provide you with their standard contract conditions?
- Can they provide reference organisations that you can contact directly?

Activity 6.11

The ITT documentation should make sure that the invitees are aware of:

- The levels of confidentiality, integrity and availability required for GANT assets.
- Any legal requirements, such as the Data Protection Act and General Data Protection Regulation, that will apply.
- Any in-house standards for information security, such as ISO/IEC 27001, to which the supplier must adhere.
- The requirement for the successful bidder to sign a formally binding contract and data protocol for any GANT data for which they may be responsible.
- Their duty of care towards GANT assets in general.
- Their agreement to no-notice compliance audits by GANT or their appointed auditors.

Activity 6.12

Your description should include the following hierarchy:

- Policy – one-page document signed by the chief executive, requiring everyone to take information assurance seriously because it is everyone's responsibility.
- Standards – defining what is acceptable in terms of information assurance design and implementation.
- Procedures – the operating instructions to ensure compliance with policy and standards. The 'how-to' documents.
- Guidelines – documents that clarify any complex areas and processes, such as risk management.

Activity 6.13

The baseline documentation should consist of a specification for each type of workstation or server, listing at least:

- minimum required version and patch level of operating system;
- components of the operating system to be installed and their configuration;
- standard access control lists and privileges for default user accounts;
- lists of applications to install, the components of them to use and their configuration;
- any specific security countermeasures or configurations to be loaded and implemented;
- network addressing format and type (DNS, dynamic host configuration protocol (DHCP), static IP, etc.);
- backup strategy for data.

The documentation to be provided for each system should detail the physical and logical build state, hardware serial numbers and location. These documents will form the reference point for the change control process.

CHAPTER 7

Activity 7.1

Without detailed knowledge it is difficult to guess how long each section could survive, but it might be reasonable for the following.

1. Membership details: it is unlikely that any details about the membership would be required urgently and so it might be appropriate for this section to talk of a week being an acceptable period. However, there could be some events that make this timescale too long. For example, if the annual conference is coming up, then it might make the information much more important, and in the last few days before

the issue of the routine newsletter, fairly critical if the newsletter is to get out on time. There would of course also be the discussion around how important the newsletter was – would it matter if it was issued a few days later?

2. Natterjack toad breeding ground details: once again, this information might not be terribly important in general and so an interruption to availability of a few days or perhaps even weeks might not be a problem. If a high-profile, very significant planning application was coming to the final stages of a court action, though, with the final court case pending in a few days' time, this too might alter things. However, it might not be feasible or financially acceptable to have special measures in place for such an event and they would have to accept the higher risk to that information.
3. Forthcoming planning application where there was interest in the toads: it may be that the details of this application are held on paper, in which case there would be no problem, but if emails, copies of letters or other related documentation was held on the computer, once again this could be critical at certain times.
4. Financial information: the general running of a group like this would not entail significant financial work at any time except perhaps at the financial year-end when the Charities Commission or Her Majesty's Revenue and Customs require financial information to be presented for inspection. This might again alter the requirements enough to warrant some additional measures being put in place.

A full business impact analysis would determine the overall level of risk and specific areas where special measures might be needed. This would then provide the necessary justification for the senior staff to decide on the appropriate measures to be taken.

Activity 7.2

The first requirement would be to carry out a risk analysis. This would determine what the real risks to the information are likely to be by not having a clear desk. Considering threats such as those posed by visitors, cleaning staff and temporary staff members as potential information thieves, would help to determine the likelihood of the risk. This assessment would also need to consider who might be interested in obtaining the information and to what lengths they might be prepared to go in order to obtain it. It would also identify GANT's privacy responsibilities to protect information concerning people and whether the current approach is fulfilling this.

Considering how well the information is currently recorded (catalogued in some way), which would allow regular checks of the information to ensure there have been no losses or compromises, would provide information on the weaknesses and vulnerabilities. Keeping records of copies taken of documents and their distribution would be useful. An information asset register might be the starting point for a more orderly control of information. A retention schedule to destroy information that is no longer required would be a useful component of this. Locking drawers, cabinets and filing facilities would clearly be a good start, but the full risk analysis would be a better start and this must be done as a result of the business impact analysis, which would then provide the justification for such action and expenditure if required.

CHAPTER 8

Activity 8.1

Clearly, the bare minimum is to take a copy of the database onto an external storage device such as a memory stick, external hard drive or a CD/DVD. In more general terms, you should be advising on the frequency of the backup and the principle of GFS, depending upon the amount of updating that is done on the database. It could be that weekly is enough, but if the database changes daily then it might need to be more frequently. The business need should drive the decision with a clear understanding of the consequences of records being lost and the cost in effort and time that would be needed to recover the situation in the worst-case scenario.

As the organisation develops, then a more formal method of doing a BCP should be considered, perhaps using a cloud-based service of some form. There could be a case for the provision of a separate system for BCP/DR with a scheme for taking the media off site.

CHAPTER 9

Activity 9.1

Possible solutions for database protection could include:

- two-factor authentication;
- encryption of the database;
- audit trail of changes and the people who made them.

Possible uses of a commercial product would need to take into account:

- key management and distribution;
- key escrow;
- use of digital signatures for emails and other data;
- use of IPSec to create VPN tunnels;
- use of TLS to create temporary secure connections for occasional and third-party users.

SAMPLE QUESTION ANSWERS

CHAPTER 1

1. The correct answer is b.
2. The correct answer is c.
3. The correct answer is d.
4. The correct answer is c.
5. The correct answer is b.

CHAPTER 2

1. The correct answer is c.
2. The correct answer is a.
3. The correct answer is d.
4. The correct answer is b.
5. The correct answer is a.

CHAPTER 3

1. The correct answer is c.
2. The correct answer is c.
3. The correct answer is a.
4. The correct answer is b.
5. The correct answer is d.
6. The correct answer is d.
7. The correct answer is c.
8. The correct answer is b.
9. The correct answer is b.
10. The correct answer is c.
11. The correct answer is d.

CHAPTER 4

1. The correct answer is b.
2. The correct answer is a.
3. The correct answer is d.
4. The correct answer is d.

CHAPTER 5

1. The correct answer is d.
2. The correct answer is c.
3. The correct answer is d.
4. The correct answer is c
5. The correct answer is b
6. The correct answer is a.

CHAPTER 6

1. The correct answer is b.
2. The correct answer is d.
3. The correct answer is a.
4. The correct answer is c.
5. The correct answer is c.

CHAPTER 7

1. The correct answer is b.
2. The correct answer is c.
3. The correct answer is d.
4. The correct answer is a.
5. The correct answer is c.

CHAPTER 8

1. The correct answer is d.
2. The correct answer is a.
3. The correct answer is d.

CHAPTER 9

1. The correct answer is b.
2. The correct answer is a.
3. The correct answer is c.

GLOSSARY

Acceptable use: A policy used to identify what personal use of company resources is acceptable

Accountability: The attribute of having to answer for one's actions

Accredited: Acknowledgement by an official body that an individual or entity has met predefined criteria

Active content: Content on a website that is either interactive, such as internet polls, or dynamic, such as animated pictures, JavaScript applications or ActiveX applications

Analysis: The detailed examination of the elements or structure of an entity

Anti-virus: Software designed to negate or destroy a computer virus

Assessment: An estimation of the nature or quality of an entity

Asset: Something that has a value to an organisation

Assurance: A positive acknowledgement designed to provide confidence

Asymmetric cryptography: A cryptographic system requiring two separate keys, one of which is secret and one of which is public

Audit: A formal inspection of an organisation's processes or procedures

Authentication: The assurance that a person or entity is who they claim to be

Authorisation: An official sanction that an individual is permitted to carry out a task or to have access to information

Availability: The property of being accessible where and when required by an authorised person, entity or process

Backdoor: A method of bypassing normal authentication methods, securing illegal remote access to a computer

Baseline controls: Standards that are used to define how systems should be configured and managed securely

Biometrics: Biometric identifiers are the distinctive, measurable characteristics used to label, describe and identify individuals

Bring your own device (BYOD): A scheme adopted by some organisations that permits staff to use their own desktop and laptop computers, tablets and smartphones instead of, or as well as, those provided by the organisation

Business continuity: The ability of an organisation to continue to function in order to deliver its products or services at an acceptable level following a business disruption

Business impact analysis: The process of analysing the consequences a business disruption might have upon the organisation's assets

Certification: A process confirming that a person has reached a predefined level of achievement

Classification: The arrangement of items into taxonomic groups – in the information security context, it labels information to identify any defined processing, handling, storage or transmission measures required to ensure appropriate security

Code of conduct: A policy that may apply to individuals to ensure that they behave in a certain way

Compliance: Acting in accordance with a set of rules or a policy

Confidentiality: The property that information is prevented from being available or disclosed to unauthorised persons, entities or processes

Corrective controls: A form of risk treatment, these are tactical controls applied after an event to prevent it recurring

Countermeasure: An action taken to counteract a threat

Cover time: The minimum time for which information must remain secret

Cryptanalysis: Used to breach cryptographic security systems and gain access to the contents of encrypted messages

Cryptography: Literally meaning hidden or secret writing, this is the practice and study of techniques for secure communication in the presence of third parties

Data leakage (also known as data loss prevention): Measures taken to prevent the unauthorised extraction of data from an organisation

Decryption: The process of taking encrypted information and returning it to a state of plaintext

Deming Cycle: The cycle of Plan, Do, Check, Act for any process or system to ensure continuous improvement

Denial of service (DoS) attack: The intentional paralysing of a computer network by flooding it with data

Detective controls: A form of risk treatment, these are tactical controls that identify events while they are taking place

Digital certificate: An electronic document that uses a digital signature to bind a public key with an identity – information such as the name of a person or an organisation, their address and so forth

Digital signature: A mathematical scheme for demonstrating the authenticity of a digital message or document

Directive controls: A form of risk treatment, these are tactical controls that provide instructions and can therefore only be procedural

Disaster recovery (DR): The activity of recovering telecommunications, IT or systems after a business disruption

Distributed denial of service (DDoS) attack: The intentional paralysing of a computer network by flooding it with data sent simultaneously from many individual computers

Domain: A common network grouping, under which a collection of network devices or addresses are organised

Encryption: The process of encoding messages (or information) in such a way that eavesdroppers or hackers cannot read it, but authorised parties can

Evaluation: The act of making a judgement about the amount, number or value of something

False positive: An indication that something has been detected or has happened when in fact it has not happened

Firewall: A technological barrier designed to prevent unauthorised or unwanted communications between computer networks or hosts

General Data Protection Regulation (GDPR): This EU legislation harmonises data privacy laws across Europe, protects and empowers all EU citizens' data privacy and reshapes the way organisations across the region approach data privacy

Governance: The action or manner of controlling a process

Hardening: The process of securing a system by reducing its surface of vulnerability

Hash digest or hash function: A derivation of data used to authenticate message integrity

Identification: The process of confirming the identity of an individual or entity

Identity: The fact of being who or what a person or entity is

Impact or consequence: The outcome of an incident that affects assets

Information security: The practice of defending information from unauthorised access, use, disclosure, disruption, modification, perusal, inspection, recording or destruction

Infrastructure as a service (IaaS): A form of cloud computing that provides virtualised computing resources over the internet

Integrity: The property of ensuring that information can only be altered by authorised persons, entities or processes

Interception or eavesdropping: The act of secretly listening to the private conversation of others without either their knowledge or consent

Intrusion: An unwanted or unauthorised access to an information system

Key-logger: A type of surveillance technology that tracks (or logs) the keys struck (or keystrokes) on a keyboard, typically in a covert manner so that the person using the keyboard is unaware that their actions are being monitored

Legal: Controlled on the basis of statutory law

Likelihood: The possibility that an event may happen

Malware: Any form of software designed to cause harm

Man-in-the-middle attacks: When an attacker infiltrates a network so that they can extract network traffic and then use, replace or corrupt it in some way to achieve nefarious aims

National Cyber Security Centre (NCSC): Part of GCHQ launched in October 2016 to be the single point of contact for all matters pertaining to cyber security for small and medium-sized enterprises, larger organisations, the general public, government departments and agencies

Network and Information Systems (NIS): An EU directive requiring member states to ensure their 'operators of essential services (OESs)' are well protected from cyber-attacks and can manage effectively any cyber incidents that may occur

Network segregation: A method of splitting a computer network into sub-networks, each being a network segment, in order to boost performance and improve security by helping to contain malware and other threats (see also partitioning)

Network sniffer: A hardware device or software program capable of logging information on a network

Non-repudiation: The ability to prove that a person, entity or process cannot deny having carried out an action

Operational risk treatment: Three types of risk control that can be used in combination to reduce risk – physical, procedural or technical

Partitioning: The division of a large network into a number of smaller sub-networks (see also network segregation)

Penetration testing: A method of evaluating the security of a computer system or network by simulating an attack from malicious outsiders and insiders

Personal data: Information relating to an individual who can be identified either from that data or from that and other data (see also personally identifiable information)

Personally identifiable information (PII): A type of data that identifies the unique identity of an individual and includes an individual's name, gender, address, telephone, email address or basic biometric data

Phishing: The act of attempting to acquire information such as usernames, passwords and credit card details (and sometimes, indirectly, money) by masquerading as a trustworthy entity in an electronic communication

Physical controls: Operational controls that consist of anything that places a physical barrier between an attacker and their target

Platform as a service (PaaS): A development and deployment environment in the cloud, with resources enabling the delivery of cloud-based applications

Policy: A principle or rule to guide decisions and achieve rational outcomes

Preventative controls: A form of risk treatment, these are tactical controls that stop things happening and they are therefore implemented before the event

Privacy: Implies personal control over personal information

Private key cryptography: A cryptographic system in which identical keys are used both to encrypt and decrypt information

Probability: The extent to which an event is likely to occur, measured by the ratio of the favourable instances to the whole number of possible instances

Procedural controls: Operational controls that consist of standards, guidelines, policies and procedures

Procedure: A list of steps that, taken together, constitute the instructions for doing or making something

Process: A sequence of events that result in an outcome, and that may consist of a number of procedures

Protocol: A set of rules that define how two entities communicate effectively

Public key cryptography: A cryptographic system in which non-identical keys are used to encrypt and decrypt information – one key is made public and the other is kept secret

Qualitative risk assessment: A subjective form of risk assessment that does not use specific values, but which may encompass a range of values

Quantitative risk assessment: An objective form of risk assessment based on numerical values

Reduce or modify a risk: A strategic risk treatment option that mitigates or reduces the threat, the likelihood or the impact of a risk

Redundancy: The inclusion of extra components that are not strictly necessary to functioning in case of failure in other components

Regulatory: Controlled on the basis of non-statutory rules

Residual risk: The risks that remain after all risk mitigation actions have been implemented

Resilience: The ability of an organisation to counter the effect of business disruptions

Reverse engineering: The ability to take a software patch, for example, and track back to determine why that patch was released, thereby exposing the vulnerability the patch was designed to close down – these vulnerabilities will then be exploited very quickly with the intention of attacking systems before the patch is installed

Risk: The combination of consequences of a threat exploiting a risk

Risk acceptance or tolerance: A strategic form of risk treatment involving an informed decision to undertake a risk when compared with the organisation's risk appetite

Risk appetite: The maximum level of risk that an organisation is prepared to accept

Risk assessment: The process of identifying, analysing and evaluating risks

Risk avoidance or termination: A strategic risk treatment involving an informed decision not to undertake, or to cease, an activity in order not to be susceptible to a risk

Risk matrix: A mechanism that allows risks to be plotted by impact and likelihood to illustrate the severity and to determine the priorities for risk treatment

Risk modification, treatment or reduction: A strategic form of risk treatment involving the reduction of the impact, or the likelihood, or both

Risk register: A database that records relevant information about risks, and can be used both for reporting purposes and to track risk treatment

Risk sharing or transfer: A strategic form of risk treatment involving the distribution of risk with other entities, for example insurance

Rootkit: A stealthy type of software, often malicious, designed to hide the existence of certain processes or programs from normal methods of detection and enable continued privileged access to a computer

Secrecy: The property that information is prevented from being available or disclosed to unauthorised persons, entities or processes

Segregation of duties: A procedural control in which one individual undertakes part of an activity and another individual undertakes the remainder

Semi-quantitative risk assessment: This combines the simplicity of a qualitative risk assessment with the more complex quantitative risk assessment by placing statements like 'low', 'medium' and 'high' in numerical bands

Sensitive personal data: Includes racial or ethnic origin; political opinions; religious beliefs; trade union affiliation; physical or mental health; sexual orientation; criminal record

Sniffer: A physical device or software designed to examine, but not stop, all traffic on a network segment to allow use by legitimate or, if unauthorised, illegal organisations

Social engineering: The act of obtaining confidential information by manipulating or deceiving people

Software as a service (SaaS): A cloud system in which software is licenced on a subscription basis, sometimes referred to as 'on-demand' software

Spyware: Software designed to gather information in a covert manner

Strategic risk treatment: Four control options of treat, terminate, tolerate or transfer (or the equivalent) a risk

Symmetric cryptography: A cryptographic system in which one key is used both to encrypt and decrypt information

Tactical risk treatment: Four ways in which risk treatment controls can be used – detective, preventative, corrective or directive

Technical controls: Operational controls that are used to restrict access to sensitive electronic information

Threat or hazard: A source of potential disruption that has the potential to cause a risk

Transfer or share a risk: A strategic risk control option used to mitigate the impact of a risk by sharing with a third party

Trojan horse: A non-self-replicating type of malware that appears to perform a desirable function but instead facilitates unauthorised access to the user's computer system

Virtual private network (VPN): Enables a host computer to send and receive data across shared or public networks as if it were a private network with all the functionality, security and management policies of the private network

Virus: A piece of software that can replicate itself and spread from one computer to another

Vulnerability: The property of something that results in susceptibility to a threat or hazard, which can result in business disruption with a consequential detrimental outcome

War dialling: Systematic dialling of every telephone number an organisation has (or may not have) to try and discover, for example, unauthorised modems, or higher levels of access to the main telephone exchange

Worm: A standalone malware computer program that replicates itself in order to spread to other computers

Zero-day exploit: An attack that exploits a previously unknown vulnerability in a computer application, meaning that the attack occurs on 'day zero' of awareness of the vulnerability

INDEX

4G networks 133
5G networks 133, 138

A5 encryption 195
access control lists (ACLs) 135, 162
accidents 14, 20, 23, 69, 95, 97–8, 112–15
ACPO guidelines for computer-based evidence *see* National Police Chiefs' Council (NPCC)
actions
  avoid 4
  mitigate 4
  prevent 4
  share 5
  terminate 4
  treat 4
active content 128
Active Cyber Defence strategy 205
activity solution pointers 215–29
Adobe Acrobat Reader 185
Advanced Encryption Standard (AES) 195–6
algorithms 132, 138, 196, 199, 201
American National Standards Institute (ANSI) 74
anti-virus software 23, 30, 32, 43, 128–9, 131, 162, 174
Apple® 74
Asimov, Isaac 12–13
assurance controls 41, 44, 53–7, 61–2, 69, 71–2
asymmetric digital subscriber line (ADSL) 225
asymmetric model 197–8
asynchronous replication 180
attempted extortion 192

backdoors 72, 127
Bayesian statistical analysis 136, 140
Big Data 202
biometric identification 111, 134
Bishop, Sir Michael 22
Black Hat 204
blackmail 67, 71, 71–2, 192
Blancco Technology Group 172
block ciphers 195–6
blockchain services 201
Blowfish 195
Bluetooth ports 129
'bots' 127
Breach of Confidence 74
bring your own device (BYOD) policy 12, 21, 133, 139, 172
Bromium Inc. 72
brown envelopes 183–4
bugs 94, 98, 131
Buncefield oil storage depot disaster 20, 66, 145, 161, 181, 185
business continuity plans (BCPs) 180–1, 183–8
business fraud 71–2
business impact analysis (BIA) 24, 28, 33, 66, 178, 182, 186
business-to-business transactions 12

Canadian Trusted Computer Product Evaluation Criteria (CTCPEC) 83
Capability Maturity Model (CMM) 97, 221
caveats 116–17
CDs 15, 128–9, 171
Centre for the Protection of National Infrastructure (CPNI) 125, 192
Certificate in Information Security Management Principles (BCS) 8, 188
certification authorities (CAs) 77, 199–200
Certified Assisted Products (CAPS) 84
Certified Cyber Professional (CCP) 6–7
Certified Ethical Hacker (CEH) 145
Charities Commission 228
Chartered Institute of Information Security (CIISec) 6
chief finance officers (CFOs) 41
chief information officers (CIOs) 41
chief information security officers (CISO) 40–1
chief risk officers (CROs) 41
child pornography 44, 73, 192, 220
China State Council directive 273 78
Chinese walls 116
ciphertext 195–6
Civil Evidence Act (UK) 76
cloud access security brokers (CASB) 136
cloud computing
  cloud-based services 21, 97, 132, 134, 185–6
  introduction 153–4
  legal implications 154

nature of 226
selecting a supplier 155–6
supplier commercial and purchaser risk 157–8
Cloud Security Alliance (CSA) 158
code 98–9, 126
code of connection (CoCo) 62, 139
codes of ethics 56
Commercial Licensed Evaluation Facilities (CLEFs) 98
commercial off-the-shelf products (COTS) 94, 100
Commercial Product Assurance system (CPA) 84, 99
Committee of Sponsoring Organizations of the Treadway Commission (COSO) 57
Common Criteria certificates (CC) 84, 99
Common Criteria for Information Technology Security Evaluation Criteria (CC ITSEC) 83
Common Criteria Recognition Arrangement (CCRA) 84
Common Vulnerabilities and Exposures database (CVE) 204
Communications-Electronics Security Groups (CESG) 83–4
CESG Tailored Assurance Service (CTAS) 83–4
Companies Act (UK) 14, 42
Companies Audit, Investigations and Community Enterprise Act (2004) 56
computer aided instruction (CAI) *see* training delivery, computer-based (CBT)
computer emergency response teams (CERTs) 67, 192
Computer Fraud and Abuse Act 1984 (USA) 71
Computer Misuse Act 1990 71
contingency plans 15, 45
contractual threats 20
controls
corrective 30
detective 30
directive 30
operational 31
preventative 30
Copyright, Designs and Patents Act 1988 (UK) 94
copyright law 74
corporate governance 14–15, 191
Council of Registered Ethical Security Testers (CREST) 46, 145
countermeasures 4, 13, 15, 22, 96, 104, 115, 127–30, 149
cover time 196
crime
blackmail 67, 71, 71–2, 192
business fraud 71–2
child pornography 44, 73, 192, 220
computer fraud 72
cyber stalking 73
cyber terrorism 20, 64, 71, 126
deception 71–2, 75, 202
espionage 20
extortion 67, 192
hacking 20, 71–2
identity theft 5, 20
and information security 16
information theft 20
internet 16
invoice fraud 16
obtaining information by deception 72
piracy 73–4
sexual grooming 73
theft 71
website defacement 72
cryptanalysis attacks 196
cryptography
basic theory 194–5
cyber threat intelligence 202
hash functions 201
policies for cryptographic use 201–2
Pretty Good Privacy (PGP) 200–1
and privacy 78
public key (PKI) 197–200
regulation of controls 68
role of 194
secret, or symmetric, key 195–7
securing data exchange 149
threat intelligence *see* threat intelligence
cyber-attacks 12, 142, 180
Cyber Essentials scheme 84
Cyber Security Breaches Survey 2019 12
Cyber Security Information Sharing Partnership (CiSP) 204–5
cyber stalking 73
cyber terrorism 20, 64, 71, 126

dark web 203
data acquisition systems (SCADA) 145
Data Encryption Standard (DES) 196–7
Data Protection Act (DPA) 20, 42, 69, 111, 117, 151, 160, 185, 225–6
deception 71–2, 75, 202
defence in depth concept 62
demilitarised zones (DMZs) 115, 139, 149
Deming Cycle *see* Plan–Do–Check–Act (PDCA)
denial of service attacks (DoS) 20, 64, 110
Department for Digital, Culture, Media and Sports (DCMS) 12
Department for Homeland Security (DHS) 192
digital certificates 138
Digital Shadows 203
digital signatures 198
Directive on Computer Misuse (EU) 71
disaster recovery (DR)
business continuity plans (BCP) 177–9
compliance with standards 188
documentation, maintenance and testing 182–4
incident management 187
managed service provision 184
off-site storage 185–6
personnel, suppliers and IT system providers 186–7
resilience and redundancy 179–80
writing and implementing plans 180–2
distributed control systems (DCSs) 145
distributed denial of service attacks (DDoS) 20, 67, 72
domain name systems (DNSs) 137

EAL 1 83, 99
EAL 4 84, 99
EAL 5–7 83, 99

eavesdropping 20, 71, 148–9, 236
eBay 172
electronic data interchanges (EDIs) 138, 149
Electronic Identification, Authentication and Trust Services (eIDAS) 77
electronic signatures 77
emails 16, 107, 129–30, 162
encryption
  benefits of 224
  GSM (2G) mobile 195
end-user code of practice *see* people security
enterprise IT 145
enterprise resource planning systems (ERP) 112, 145
espionage 20
ethics 6–7
European Commission 70, 84
European Free Trade Association (EFTA) 84
European Patent Convention (EPC) 75
European Telecommunications Standards Institute (ETSI) 84–5
European Union Agency for Network and Information Security (ENISA) 85, 122, 158
European Union (EU) 42, 44, 69–70, 74–5, 77–8, 84, 160
external services
  other organisations 151
  protection of web services 149–50
  real-time 147–8
  securing data exchange 149
  service management 152
extortion 67, 192

Fair and Accurate Credit Transaction Act 2003 (US) 73
Federal Bureau of Investigation (FBI) 126
Federal Information Processing Standards Publications (FIPS PUBS) 84
Federal Rules of Evidence (US) 76
Federal Trusted Computer System Evaluation Criteria (TCSECUS) 83
Financial Conduct Authority (FCA) 45, 82
Financial Services Act (FSA) 151
fingerprint identification 134, 187
firewalls 99, 115, 129–31, 133, 135–6, 139, 224
forensic investigation 16, 187
Forum for Incident Response and Security Teams (FIRST) 192
France 78, 84
Freedom of Information Act (FoIA) 71, 151

GATT, Trade Related Aspects of Intellectual Property Rights 1993 (GATT TRIPS) 74
GCHQ, Cheltenham 83, 197
General Data Protection Regulation (GDPR) 20, 42, 44–5, 67, 69–70, 73, 133, 135, 143, 151, 185, 225–6
  Argentina 70
  Canada 69–70
  New Zealand 70
  Switzerland 70
  United Kingdom 70
Get Safe On-line (website) 122
Global Information Assurance Certification (GIAC) 145
Gnu PrivacyGuard 200
Good Practice Guide for Computer Based Electronic Evidence 2012 (ACPO) 67, 76–7
good practice guidelines (Business Continuity Institute) 188
GOVCERT (emergency response team) 192
Gramm-Leach-Bliley Act (GLBA) 69,143
Grandfather-Father-Son system (GFS) 130, 160–1, 229
GSM (2G) mobile encryption 195

hacking 20, 25, 71–2, 115, 145
hashing 198–200
Health Insurance Portability and Accountability Act (HIPAA) 69, 135, 143
health and safety 14
HMRC 228
'hole-in-the-wall' cash dispensers 22
host intrusion detection systems (HIDSs) 161
https protocol 114
Human Rights Act (HRA) 70, 151

IAM Roadsmart 181
ID&A *see* user access controls, authentication and authorisation mechanisms
identity theft 5, 20
IdenTrust 200–1
incident investigation
  common processes 190–1
  cryptography *see* cryptography
  forensic services/third parties 190, 192–4
  legal and regulatory guidelines 191
  relations with law enforcement 191–2
Independent Evaluation for Assured Services (CAS) 84
indicators of compromise (IoCs) 203
industrial control systems (ICSs) 11, 145
information
  retrieval 5
  security 6–7
information assurance (IA)
  business models 10–12
  controls 4–5
  international/national standards 8
  management of 4
  middle management 14–15
  not an 'add-on' 10
  policy 14
  for the whole organisation 10
Information Commissioner's Office (ICO) 45, 122
information and communications technology (ICT) 84
information risk
  accepting or tolerating 30
  analysis 26–8
  assessing in business terms 34
  assets 22, 33
  avoidance or termination 29
  calculating 23–4
  classification policies 33–4
  communication and consultation 28
  context establishment 25
  cost against potential losses 34–5
  identification 25–6

identifying value 33
impact 22–3
likelihood 23
management process 25
matrix 24, 26–8
monitoring and review 28–9
other controls 30–1
reducing 29
role of management 35
threat categorisation 20
threats 3, 19
transferring or sharing 29–30
treatment 28
vulnerabilities 21–2
information security framework
audits and reviews 54–5
board/director responsibilities 41–3
compliance checks 55–6
compliance status 56–7
defence, depth and breadth 51
end-user code of practice 51–2
good information security practice 46–7
governance 53–4
legal *see* legal framework
organisational responsibilities 43–4
organisation's management of 39
physical, procedural, technical controls 50–1
placement 40–1
policies, standards, procedures 47–50
policy violation 52
programme implementation *see* programme implementation
roles 39–40
security incident management *see* security incident management
security standards and procedures *see* security standards and procedures
specialist information 45–6
statutory, regulatory, advisory requirements 44–5
Information Security Management System (ISMS) 7, 17, 140
information security managers 40
Information Technology Security Evaluation Criteria (ITSEC) 83
information theft 20
infrared ports 129
infrastructure as a service (IaaS) 21, 153–4, 155, 158
instant messaging (IM) 147–8
Institute of Advanced Motorists *see* IAM Roadsmart
Institute of Information Security Professionals (IISP) *see* Chartered Institute of Information Security (CIISec)
integrated services digital networks (ISDNs) 148
intellectual property (IP) 68, 72
International Electrotechnical Commission (IEC) 80
International Organization for Standardization (ISO) 74, 80–3
International Telecommunication Union (ITU) 80
internet
activity solutions 224
and business models 11–12
and crime 16
cyber threat intelligence 203
disaster recovery 185–6
information security framework 72–3, 75, 84
management principles 172–3
procedural/people 110
and risk 30
security life cycles 89
technical security controls 132–9, 145–6, 148–50, 153–4, 162
Internet Engineering Task Force (IETF) 84
Internet Key Exchange (IKE) 197
internet protocol security (IPSec) 132, 136, 138, 196, 197, 229
Internet of Things (IoT) 41
internet-crime 16
Into the Web of Profit (McGuire) 72
intrusion detection system (IDS) 225
invitation to tender (ITT) 226
invoice fraud 16
IoT devices 136
IPS solutions 161–2
iris identification 134
ISA/IEC 62443 145
ISF Standard of Good Practice 40, 41, 49–50, 158
ISO 12812-2 7
ISO 15489-1:2016 (Record Management Standards) 74, 81
ISO 15638-15 6
ISO 19092 10
ISO 22300 (2018) 4
ISO Guide 73 4–5
ISO/IEC 2000 ITIL (management of information technology services) 81
ISO/IEC 9001 (quality assurance) 81, 221
ISO/IEC 13335 2
ISO/IEC 15408-1 (2009) 83
ISO/IEC 15944-6 5
ISO/IEC 21827 6
ISO/IEC 22301:2019 (business continuity requirements) 81, 188
ISO/IEC 22313:2014 (business continuity requirements, additional guidance) 188
ISO/IEC 24760-1 5
ISO/IEC 24762:2008 (DR services) 188
ISO/IEC 27000 series 1–3, 40, 49–50, 68, 78–80, 95–6, 99, 143, 221
ISO/IEC 27001 18, 57, 61, 80–2, 139, 143, 159, 226
ISO/IEC 27002 80–1, 159
ISO/IEC 27005 (risk management) 80
ISO/IEC 27006 82
ISO/IEC 27017 (using cloud resources) 80–1
ISO/IEC 27031:2011 (guidance for IT readiness/business continuity) 188
ISO/IEC 27033 (network security) 80
ISO/IEC 27799 (managing information security in health sector) 80
ISO/TR 19591 6
ISO/TR 22100-4 5
IT infrastructure
access control lists and roles 159–60
configuration management 163

conformance with security policy 159
input correctness 160
installation of baseline controls 162
intrusion monitoring and detection methods 161–2
recovery capability 160–1
security documentation 163
separation of systems 158
IT Infrastructure Library (ITIL) 143, 188

just-in-time operations 12

Kegworth air crash 22–3
Kerberos 134, 201
key AES lengths 197
key exchange problem 197
key-logging 16, 114, 127
kopf unten (head-down generation) 139

legal framework
collection of admissible evidence 76–7
common concepts 71–3
contractual safeguards 75–6
copyright law 74
employment issues/employee rights 70–1
intellectual property rights (IPR) 74–5
introduction 67–8
protection of personal data 69–70
records retention 73–4
restrictions on purchase 78–9
securing digital signatures 77–8
line of business managers (LOB) 42
local area network (LAN) 136
local security co-ordinators 44
lottery win notifications 16

McGuire, Mike 72–3, 75
Maersk (shipping line) 142
malevolent cookies 127
malicious code *see* malicious software
Malicious Communications Act (UK) 73
malicious software
code of conduct 107
and common concepts of computer misuse 72
and disgruntled employees 175
and law enforcement 67
malware countermeasures 129–30
methods of control 130–2
and monitoring 91
now very sophisticated 51
and off-the-shelf products 94–5
and risks of unwanted code 98
routes of infection 128–39
security breaches 161
and security incident reporting 64
sheepdip scanners 130
subverting stored data 160
technical controls 126, 128–9
types 20, 126–8
malware *see* malicious software
management reviews 54
Markets in Financial Instruments Directive (MiFID) 151
masquerade attacks 198
message digest 198–9
Microsoft Office 365® 153
Microsoft® 74, 180
multi-factor authentication (MFA) 135
multi-protocol layer switching (MPLS) 136

National Counterintelligence and Security Center (US NCSC) 122
National Crime Agency (NCA) 67, 192
National Cyber Security Centre (NCSC) 6–7, 46, 67, 83, 84, 99, 192, 205
National Institute of Standards and Technology (NIST) 57
NIST Cybersecurity Framework 66, 84, 122, 125, 158, 204
National Police Chiefs' Council (NPCC) 191, 220
Nessus (scanning tool) 145
network cabling 168–9
network and communication links
control of third-party access 138–9
cryptography 136–8
the DMZ 136
entry points 132–3
external services *see* external services
firewalls 141
identification and authentication 134–5
intrusion monitoring 140–1
network usage policy 140
operational technology 144–7
partitioning networks 135–6
secure management 142–4
wireless 133–4
Network and Information Systems Regulations 2018 (NIS) 42, 143
network intrusion detection systems (NIDSs) 161
network sniffers 114, 135
Next Generation (firewall) 141
Nmap (scanning tool) 145
non-disclosure agreements (NDAs) 7, 55, 193
NotPetya attack 142

Obscene Publications Act (UK) 73
Offensive Security Certified Professional (OSCP) 145
Official Secrets Act (OSA) 151
OKTA (cloud-based services) 134
onion model 103
Open Source Intelligence (OSI) 203
OpenIOC 204
OpenPGP 200
operational technology (OT) 145
operational types of control *see* information security framework, physical, procedural, technical controls
operators of essential services (OESs) 42
Orange book *see* Federal Trusted Computer System Evaluation Criteria (TCSECUS)
out of band authentication (OOB) 134
outages 20

PACE/ACPO guidelines 193
paedophilia *see* crime
Pakistan 78
Pakistan Telecom Authority 78
'passing off' 74–5
passwords
  developing standards, guidelines 48–52
  hard-coded 162
  one-time (OTPs) 111
  and partitioning networks 135
  protection 195
  sharing 120
patches 98–9, 131, 162
patents 75
Payment Card Industry Data Security Standard requirements (PCI DSS) 57, 82, 135
payroll systems 112
Peabody, Jane 8, 123
PenTests 145–6
people security
  acceptable use policies 107
  codes of conduct 106–7
  contracts of employment 104–5
  organisational security culture 104
  security awareness 105
  segregation of duties 108
  service contracts 106
  third-party obligations 108–9
Personal Information Protection and Electronic Documentation Act (Canada) 69
personally identifiable information (PII) 69, 106, 154, 160, 172, 224
Peter, Laurence J. 170
phishing 16, 72
photocopiers 15
physical security
  clear screen and desk policy 169–70
  delivery and loading areas 173
  description 103
  introduction 166–7
  moving property on and off site 170–2
  protection of equipment 167–9
  secure disposal procedures 172–3
  tactical controls 174–5
PIN numbers 52, 110, 135
piracy 73–4
plaintext 195
Plan–Do–Check–Act Cycle (PDCA) 57, 91, 143–4
platform as a service (PaaS) 153–4, 155, 158
Police and Criminal Evidence Act (1984) 76, 190
Ponemon Institute report 142
privacy 69–70
Privacy Shield framework (US) 70
private automatic branch exchange systems (PABX) 148
procedures and people
  general controls 102–4
  introduction 102
  people security *see* people security
  security 103
  training and awareness *see* training and awareness
  user access controls *see* user access controls
professionalism 6–7
programme implementation
  planning 58–9, 63
  presenting positive benefits 60–1
  security strategy and architecture 61–2
Protection from Harassment Act (UK) 73
ProtMon tools 140–1
public key infrastructure (PKI) 134–5, 197–200
Public Order Act (UK) 73
Public Records Acts 1957, 1967 (UK) 70
Publicly Available Specification 77 (PAS 77) 81, 188

Radius 134
ransomware 20, 67, 72–3, 127, 142, *see also* malicious software
Ratner, Gerald 22
Ratners jewellers 22
Recorded Future 203
Regulation of Investigatory Powers Act 2000 (RIPA) 70
relevant digital service providers (RDSPs) 42
Request for Comments (RFCs) 84
retina identification 134
return on investment (ROI) 60
risk
  acceptance 5
  appetite 27
  and corporate governance 14–15
  countermeasures 13
  description 3–4
  do nothing option 5
  during an audit or review 55
  and life 13
  and organisations 13
  reduction 4
  registers 27, 30, 35–6, 56
  toleration 5
  transfer 4–5
risk assessments
  and BIA 178
  as part of design/development life cycle 95
  qualitative 31
  quantitative 32
  questionnaires 32–3
  results of 34
  semi-quantitative 32
Risk Guidance 2014 (Financial Reporting Council) 35
road warriors 138
role-based access 113
Rolls-Royce 139
rootkits 127
RSA SecurID device 110, 199, 204

sabotage 20
SABSA matrix 143
sample questions 17–18, 36–7, 85–7, 100–1, 123–4, 164–5, 175–6, 206–7
SANS Institute 123
Sarbanes–Oxley Act (2002) 14, 42, 56, 108
screen-scraping 16
secure sockets layer (SSL) 114
security
  analytics 202
  champions 42

forums 42
technical 103, 125
security incident management
corporate systems 66
introduction 63–4
law enforcement 66–7
reporting, incident response teams/procedures 65–6
reporting, recording 64–5
security information and event management tools (SIEM) 92, 140–1, 203
security life cycles
information 88–90
monitoring system principles 91–2
systems development and support *see* systems development and support
testing, audit and review 90–1
testing, links between IT and clerical processes 91
security operations centres (SOCs) 92, 136
Security Operations Maturity Architecture (SOMA) 57
security standards and procedures
certification of information security management systems 81–2
introduction 79
national and international standards 79–81
product certification 82–4
production of key technical standards 84–5
security training programmes 119–20
service level agreements (SLAs) 92
sexual grooming 73
SHA-256 199
shadow IT 138
single sign on systems (SSO) 134
slammer worm 127, 142
smart energy metering networks 196
smartphone zombies 139
smartphones 129
Snort (tool) 161
social engineering 20
Social Media Today 89
software as a service (SaaS) 21, 153, 155
sponsorship 119
spyware 127, 239
SSL cryptography 149
stand-by power generators 168–9
standby systems 179–80
Stateful Inspection (firewall) 141
steering committees 42
strategic controls 28
stream ciphers 195
Subject Access Requests 70
supply chains 187
symmetric encryption keys 138, 196, 199
synchronous replication 180
sysadmins 112–13
system misuse 108
systems development and support
accreditation 96
change control 96–7
commercial products 94
links with all business areas 94
preventing covert channels 98
security of acceptance processes 95–6
security issues from outsourcing 97
security patching 98–9
security requirement specification 93
separation of development/live systems 95
system and product assessment 93–4
use of escrow 99–100

target of evaluation (ToE) 83
Target (retailer) 139
Telecommunications Directive (UK) 70
terrorist activity 192
theft 71
threat intelligence
co-operative 204–5
cyber (CTI) 203–4
introduction 202–3
and vulnerability data 202
Tiger Scheme certifications 145
TLS cryptography 137, 149, 196–7, 229
training and awareness
approaches to 119–20
available materials 120–2
information sources 122–3
introduction 117–18
purpose 118–19
training delivery
computer-based (CBT) 121
electronic formats 121
escape rooms 121
external 120
face-to-face 120
videos 120–1
transport layer security (TLS) 132
Triple-Data Encryption Standard (Triple-DES) 195
Trojans 127–9, 135
Turnbull Report (1999) 14, 35, 42
two-factor authentication (2FA) 110

United States Department of Commerce 70
USB memory sticks 15, 128–9
user access controls
access points 114–15
administration 113–14
authentication and authorisation mechanisms 109–11, 114, 116, 135, 139, 150
data protection 115
effective use of 111–13

video-conferencing 148, 225
virtual private networks (VPNs) 114, 132–3, 136, 139, 149–50, 229
viruses 20, 32, 126–7
voice over internet protocol (VOIP) 148
vulnerabilities 3, 90–1, 146, 215

Wannacry ransomware virus 127, 180
war chests 185
war dialling 148

Wassenaar Arrangement 1996 (WA) 78
website defacement 72
websites 127–31
wide area network (WAN) 136
Wi-Fi protected access protocol (WPA) 148
WikiLeaks 142
wired equivalent privacy protocol (WEP) 148
wireless access points (WAP) 134
wireless networks 114–15, 133, 136, 148
Wireshark (software) 135, 145
World Standards Cooperation 80
World Trade Center incident, New York 181, 185
worms 127, 129
WPA3 134

YouTube 122

zero-day exploits 128

INFORMATION SECURITY MANAGEMENT PRINCIPLES

Third edition

Andy Taylor (editor), David Alexander, Amanda Finch, David Sutton

Information is one of the currencies of today's society. As access to fast, reliable data at work and at home becomes increasingly essential to day to day operations, new risks emerge which threaten the very information that enables businesses and helps society to function.

By focusing on the three main areas of information assurance – confidentiality, integrity and availability – this book gives you the skills to identify information security threats and protect yourself and your business against them.

- **Understand information threats and vulnerabilities and implement countermeasures against these**
- **Manage emerging risks to your data**
- **Learn information assurance best practice from experienced authors**
- **Supports BCS certification in Information Security Management Principles**

ABOUT THE AUTHORS

The authors are at the forefront of information security and are instrumental in shaping policy and implementing best practice. They have gained considerable experience across a wide range of public and private sector bodies including the Home Office, MoD, RAF, Royal Navy, British Airways, Marks & Spencer and O2.

Fantastic for those studying information security management and as a desk-side reference...refreshingly understandable.

Helen Mary Jones, Group Information Security Manager, The Jockey Club

An excellent introduction to information security. Highly recommended.

John Hughes, InfoSec Skills
Review of previous edition

You might also be interested in:

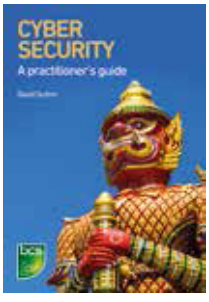

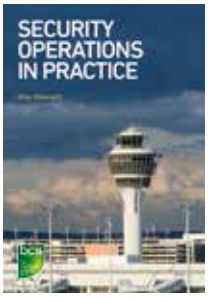

Information Technology, Management

Cover photo: iStock © SteveMcsweeny

bcs The Chartered Institute for IT

Paperback available

ISBN 978-1-78017-518-8

9 781780 175188