

FINANCIAL TIMES

# Enterprise Information Security

*Information security for  
non-technical decision  
makers*

PETER GREGORY

Executive  
Briefings

 [briefingsource.com](http://briefingsource.com)



Prentice Hall

FINANCIAL TIMES



In an increasingly competitive world, we believe it's quality of thinking that will give you the edge – an idea that opens new doors, a technique that solves a problem, or an insight that simply makes sense of it all. The more you know, the smarter and faster you can go.

That's why we work with the best minds in business and finance to bring cutting-edge thinking and best learning practice to a global market.

Under a range of leading imprints, including *Financial Times Prentice Hall*, we create world-class print publications and electronic products bringing our readers knowledge, skills and understanding which can be applied whether studying or at work.

To find out more about our business publications, or tell us about the books you'd like to find, you can visit us at [www.business-minds.com](http://www.business-minds.com)

For other Pearson Education publications, visit [www.pearsoned-ema.com](http://www.pearsoned-ema.com)



# Enterprise Information Security

*Information security for  
non-technical decision makers*

**PETER H. GREGORY**

**FT** Prentice Hall  
FINANCIAL TIMES

**IT BRIEFINGS' SERIES EDITOR: SEBASTIAN NOKES**

An imprint of Pearson Education

London ■ New York ■ Toronto ■ Sydney ■ Tokyo ■ Singapore ■ Hong Kong ■ Cape Town  
New Delhi ■ Madrid ■ Paris ■ Amsterdam ■ Munich ■ Milan ■ Stockholm

PEARSON EDUCATION LIMITED

Head Office:  
Edinburgh Gate  
Harlow CM20 2JE  
Tel: +44 (0)1279 623623  
Fax: +44 (0)1279 431059

London Office:  
128 Long Acre  
London WC2E 9AN  
Tel: +44 (0)20 7447 2000  
Fax: +44 (0)20 7447 2170  
Website: [www.briefingzone.com](http://www.briefingzone.com)

---

First published in Great Britain in 2003

© Pearson Education Limited 2003

The right of Peter H. Gregory to be identified as author  
of this work has been asserted by him in accordance  
with the Copyright, Designs and Patents Act 1988.

ISBN 0 273 66157 4

*British Library Cataloguing in Publication Data*

A CIP catalogue record for this book can be obtained from the British Library.

All rights reserved; no part of this publication may be reproduced, stored  
in a retrieval system, or transmitted in any form or by any means, electronic,  
mechanical, photocopying, recording, or otherwise without either the prior  
written permission of the Publishers or a licence permitting restricted copying  
in the United Kingdom issued by the Copyright Licensing Agency Ltd,  
90 Tottenham Court Road, London W1P 0LP. This book may not be lent,  
resold, hired out or otherwise disposed of by way of trade in any form  
of binding or cover other than that in which it is published, without the  
prior consent of the Publishers.

10 9 8 7 6 5 4 3 2 1

Typeset by Monolith – [www.monolith.uk.com](http://www.monolith.uk.com)  
Printed and bound in Great Britain by Ashford Colour Press Ltd, Gosport, Hants.

*The Publishers' policy is to use paper manufactured from sustainable forests.*

---

## About the series editor

Sebastian Nokes is Series Editor for the IT Management series within the Financial Times Prentice Hall Executive Briefings and has written and co-written several works on the subject. These Briefings are designed to provide concise, focused knowledge, concerning critical IT issues facing managers today. They deliver the information and insight needed to evaluate situations and make informed decisions.

Sebastian is a partner at Kimbell Evaluation Ltd, a leading consulting and analytical firm. Kimbell Evaluation was co-founded by Sebastian and helps clients measure and manage value added by information technology, both in stable organizations and in business or divisional turnarounds. His consulting clients include international financial institutions and commercial corporations.

Sebastian began his career in the IT and investment banking industries, and has been an employee of Credit Suisse First Boston and IBM. He was educated at London University, served in the 2nd KEO Goorkhas, and holds finance and engineering qualifications.

Sebastian may be contacted via [Nokes@ITVA.Net](mailto:Nokes@ITVA.Net)



---

## About the author

Peter H. Gregory, CISSP, is a Principal at the HartGregory Group, which provides business and information technology services including systems architecture planning, business development and infrastructure consulting, strategic marketing, communications and public relations. He is the Logical Security Strategist at a wireless telecommunications provider in Washington State, US, where he is responsible for developing security vision, strategy and policy for products, services and internal infrastructure. Prior to this, he was Manager of Enterprise Security Architecture where he developed a vision of centralized logical security mechanisms for the enterprise. His previous experience includes managing IT infrastructure planning, implementation and support, software design and engineering.

He is the author of *Solaris™ Security* (Sun Microsystems Press, 1999; published in English, Chinese and Japanese), *Sun Certified System Administrator for Solaris 8 Study Guide* (Sun Microsystems Press, 2002) and co-author of *CISSP for Dummies* (John Wiley & Sons, 2002). He is the creator and series editor for the *Defense in Depth* series of security books written for non-technical business decision makers (Prentice Hall PTR). He has tech edited or reviewed more than 20 computer science titles for Prentice Hall PTR and other publishers since 1994.

Peter Gregory is an active member of the Computer Security Institute, the New York Electronic Crimes Task Force, the IEEE Technical Committee on Security and Privacy, the Internet Society, the Microsoft Global Executive Roundtable Security Forum and the International Information Systems Security Certification Consortium. He is also a frequent conference speaker.

Peter can be reached at [peter.gregory@hartgregorygroup.com](mailto:peter.gregory@hartgregorygroup.com)



---

# Contents

List of tables	xii
List of figures	xiii
Executive summary	xv
Introduction	xix
<b>1 Security is on centre stage</b>	<b>1</b>
The priority of information security	4
Impact of 2001 events	5
Proliferation of extranets	6
Insiders: the real threat	7
Unprecedented dependence on information technology	7
Summary	8
<b>2 Threats and vulnerabilities</b>	<b>9</b>
Introduction	11
Threats	11
Vulnerabilities	17
Summary	25
<b>3 Security fundamentals – the principles and the mechanisms behind them</b>	<b>27</b>
Introduction	29
Identification and authentication	30
Authenticating other systems	33
Authorization	38
Access control	41
Encryption	48
Non-repudiation	60
Integrity	60
Audit	66
Availability	68

	Security mechanisms work together	68
	Summary	69
<b>4</b>	<b>Security policies and requirements – defining the standard of architecture and behaviour</b>	73
	Introduction	75
	What are information security policies?	76
	Who writes security policies?	76
	Audience	77
	Policy development	78
	Awareness	83
	Enforcement and effectiveness	85
	Summary	88
<b>5</b>	<b>Security is about people’s behaviour</b>	91
	Introduction	93
	Technology is not the solution	93
	The ‘people threat’	94
	Mitigating the threat	99
	Trust	100
	Summary	102
<b>6</b>	<b>Protecting corporate information beyond the corporate boundaries</b>	103
	Introduction	105
	The new world	105
	Regaining control	107
	Summary	112
<b>7</b>	<b>Privacy</b>	113
	Introduction	115
	What is personal information?	115
	It’s all about trust	117
	Privacy policy	118
	How security supports privacy	120
	Privacy certifications	121
	Summary	122

<b>8</b>	<b>Action items</b>	125
	Most important and urgent action items (Quadrant I)	128
	Most important but less urgent action items (Quadrant II)	130
	Important and urgent action items (Quadrant III)	132
	Important and less urgent action items (Quadrant IV)	133
	Epilogue	135
	<b>References/sources for additional information</b>	137

---

# Tables

3.1 Encryption algorithms	52
4.1 Summary of security policy statements	81

---

# Figures

2.1	Threats and vulnerabilities	11
3.1	Storing and using passwords	38
3.2	Network firewall	42
3.3	Host-based access control	46
3.4	Symmetric encryption	49
3.5	Public key encryption	51
3.6	VPN architectures	55
3.7	Change and configuration management	64
3.8	Interdependence of security mechanisms	69
4.1	Security policy hierarchy	80
6.1	The trading partner web	109
8.1	Enterprise security action item quadrants	127
8.2	Quadrant II	130
8.3	Quadrant III	132
8.4	Quadrant IV	133



---

# Executive summary

Information security has existed as a formal discipline since the mainframe era. Infosec, as it is sometimes called, has evolved over time and has spawned a number of formal methodologies. Universities and corporate research and development (R&D) have developed several security models that have been used as a design basis for access control and trust mechanisms.

Despite the advancement of processes, methodologies, tools, products and advanced thought from academia and R&D, infosec has through the years been maligned as a non-productive activity that lengthens project schedules, delays time-to-market and drains development resources. The infosec staff were frequently seen as paranoid and some wondered whether they got and kept their jobs using fear tactics instead of common business sense. Many organizations did not have infosec staff at all – they were adequately secure and did not need anyone dedicated to that fact.

This began to change in the mid 1990s as organizations began to connect their internal networks to the Internet. Companies soon realized that perimeter controls were necessary to protect the integrity of information systems. Firewalls become a common practice for perimeter security, but organizations did little else to protect their information assets and infrastructure.

The increased use of extranets in the late 1990s brought to light the revelation that most enterprise applications did not have sufficient access controls to admit trading partners properly and confine them to data that was relevant to them. This was seen as a lack of access control granularity, but some realized the problem for what it was: a security problem at the application level. For the first time, security was seen not just as perimeter issue but as an internal one as well. Security was beginning to earn legitimacy in the information and business world.

Still, security professionals struggled for legitimacy and met with resistance when they tried to integrate security into business processes, applications and products.

The great viruses and worms in the years 2000 and 2001 began to change attitudes, but the events that propelled security to full legitimacy were the terrorist attacks on the USA in September 2001.

No longer the domain of the information technology department, discussions about security take place in the CEO's office and the board room. Senior executives, who knew that there were firewalls and anti-virus technology, knew little else but were being called to describe how they were protecting their enterprises. The answer to the question, 'How is your company protecting its information assets?' demands an answer. 'Let's ask the infosec people' or 'anti-virus and firewalls' are no longer sufficient answers. The executive must be more knowledgeable about just how his or her company *does* protect its assets.

This book removes the veil from the executive's eyes and explains the whats and hows of information security in common language. This will enable the executive to understand the technologies, practices and business processes that are used to protect the organization's information assets and infrastructure, so that he or she can ask the tough questions of the infosec staff and truly understand what they are talking about. Executives understand how their company's finances, audits and taxes work; it's high time they understood how the business protects itself from the plethora of threats that exist today.

## **HOW THIS BOOK IS ORGANIZED**

---

This book provides a reference to the important issues in information security facing every organization. The topics covered are as follows.

- Chapter 1, *Security is on centre stage*, explains why information security has been propelled from the backburner to become one of the most important issues of the day. Trends on security spending, noteworthy events of the year 2001, and new business activities such as extranets are included.
- Chapter 2, *Threats and vulnerabilities*, describes in detail the business-process, behavioural and technical risks to the business. Threats and vulnerabilities are defined and discussed in detail. This chapter, more than any other, emphasizes the fact that security depends upon people far more than it depends upon technology.
- Chapter 3, *Security fundamentals – the principles and the mechanisms behind them*, introduces and discusses the fundamental concepts of identification, authentication, authorization, access control, encryption, non-repudiation, integrity, audit and availability. Each of these sections includes definitions, terminology and examples of the mechanisms used, including single sign-on, PKI, anti-virus, VPN, configuration and change management, and personal firewalls.
- Chapter 4, *Security policies and requirements – defining the standard of architecture and behaviour*, describes the role of business security policies and requirements in the enterprise, who writes and enforces them, and how they are written.
- Chapter 5, *Security is about people's behaviour*, asserts that technology alone cannot adequately protect the enterprise. In fact, without knowledgeable people doing the right thing, securing the enterprise with technology is just about worthless.
- Chapter 6, *Protecting corporate information beyond the corporate boundaries*, discusses the growing phenomenon of organizations that are connecting their networks, and even their applications, to business trading partners. This chapter outlines the risks and possible remedies associated with extranets.

- Chapter 7, *Privacy*, begins with a definition and discussion of personally identifiable information (PII) and then discusses identity theft, trust and privacy policy. The chapter continues with a discussion of privacy's need for security and concludes with a summary of third-party privacy certifications.
- Chapter 8, *Action items*, presents a Steven Covey-like quadrant with urgent and important axes, and provides a priority-order list of tasks that senior managers should perform in order to utilize their resources most efficiently.
- The final section, *References/sources for additional information*, contains recommended reading and over 70 websites, including security portals, government, privacy, corporate and personal security certifications, security intelligence and research, security policies, security trade groups and associations, security trade shows, conferences and seminars, periodicals and other information.



---

# Introduction

Hacking attacks. Defacements. The stealing of thousands of credit card numbers. State-sponsored cyberterrorism. Insider sabotage. Identity theft. Superviruses and worms. Some say that the end of innocence is over. Others say that we can't put the genie back into the bottle.

The readers of this book whose careers began in information technology can recall kinder, gentler days when information security was intrinsic, natural and real. Information security meant locking the data centre and the offices where computer printouts and backup tapes were located.

Where did the sense of security go?

We abandoned the sense of security as we evolved the science of information technology. The isolation of the batch-oriented, punch-card computer room had built-in security: the doors were locked, the building itself was locked, perhaps there were no windows, and all the data was inside and stayed there.

Bringing information and access to information out of the computer room brought with it the increased risk that someone with no justifiable need to see the data would see it anyway. In the batch mainframe environment, one had to have clearance to be in the building and to be in the computer room to have access to the information. But even then, there were userids (user identifications) and passwords.

When we strung communications lines from the computer room to people's offices in nearby buildings, we stripped away the stringent physical access requirements of the data centre, and the userid and password became the first and only defence. We lived with this for decades, thinking it was enough, and most of the time it was.

Local area and campus area networks exacerbated the access control problem. The data that once resided in the castle was now accessible to a growing number of people, and still the only defence was the userid and password. The openness and ubiquity of TCP/IP protocols soon enabled people to eavesdrop on network traffic: userids, passwords and every sort of sensitive data were easily readable as they traversed the network.

As if this wasn't enough, the next capability that facilitated access from a distance was remote access. This gave everyone with a terminal or personal computer and a modem the ability to reach the locked confines of a network from almost anywhere in the world. And still we relied solely upon userids and passwords, and still all traffic was transmitted 'in the clear'. We were extending the reach of the private network to a potentially greater number of people without balancing it with a stronger security mechanism. A few capabilities such as dial-back mitigated this exposure somewhat, but it was not widely adopted because people could only dial in from predefined locations.

But the Internet increased the exposure of the private network by several orders of magnitude at once. By the mid 1990s, there were hundreds of thousands of individuals in the world with detailed knowledge of TCP/IP, and thousands more being produced each year by the world's colleges and universities. Students are encouraged to experiment, and that prompting plus a relative immaturity of people of that age gave rise to vast numbers of curious and potentially malicious individuals.

But hackers are not solely to blame for security issues. In fact, they account for only a small minority of security incidents. The truth is, the majority of security incidents and issues are perpetrated by organizations' insiders. That's right – the people who are given the responsibility to be good stewards and are implicitly responsible for the integrity of corporate information are the very people who are altering it, stealing it, selling it, sabotaging it, corrupting it and destroying it.

The ubiquity of Internet protocols and personal computers has made it easier for people with poor judgement to access and manipulate the data with which they are entrusted.

Why has this all been allowed to happen? Time-to-market competitive pressures, together with the perception that security is an unnecessary obstacle to ease of use, made security of information a lower priority than it should have been. The technologies that facilitated easier access to information for those who have a need to work with it also facilitated easier access to information for those who had *no* need to work with it. While this may sound trite, the truth is that as we extended connectivity and usability further and further, we were increasing our risk of security incidents, most often without realizing it. We were so focused on network connectivity being 'cool' that many of us failed to consider the increased risk and to mitigate the risk with better controls.

To use an analogy: we made cars that went faster and faster, but failed to appreciate the increased risks of injury and death until we had literally years – or decades – of statistics that suggested that perhaps something had to be done. We are in similar straits with information technology. We let the genie out of the bottle without realizing that it could be used against us. At last we have awakened to the fact that enabling technologies are not risk-free.

## **THREATS AND DUE DILIGENCE**

Threats of mass technological disruption have moved from the storybook to the news headlines. Recently it has become painfully obvious that most companies have done a grossly inadequate job of protecting their information assets and infrastructure. Embarrassing hacking attacks and insider extortion jobs have forced senior executives to account for their failures to provide adequate protection.

Senior executives have found that they need quickly to come 'up to speed' on the methods and tools used to protect information assets. No, this does not mean

that the chief executive must learn how to operate the corporate firewall, but he or she does need to know that there is one, and also what other controls are in place to protect the company.

This book will bring the tools, technologies and practices into the light in a language that the senior executive can understand and work with. It will equip the reader with enough knowledge to have the confidence to ask the tough questions of IT managers, and understand the answers and what they mean.



# Security is on centre stage

- The priority of information security 4
- Impact of 2001 events 5
- Proliferation of extranets 6
- Insiders: the real threat 7
- Unprecedented dependence on information technology 7
- Summary 8



Senior executives in practically every industry in the world have been in a deep slumber. Only the security practitioners and their managers understood the issues of the day and how to protect the organization from harm. In their journals, newsgroups and technical conferences, they read and learned about hacking techniques, breakins and the resulting damage to their businesses. But rarely did these stories make the front page in the leading business newspapers and magazines. So in retrospect it is no surprise that security professionals were largely ignored when they tried to raise the issues familiar to them.

In the late 1990s, a new rash of threats and vulnerabilities entered the scene: there were new, rapidly spreading viruses and Trojan horses such as Melissa, ILoveYou and BackOrifice. Millions of dollars were stolen from a large US bank by a hacker in Russia. In February 2000, distributed denial of service attacks took Yahoo!, Amazon and other prominent websites off the air for several hours.

These events were sufficiently serious to rouse a few senior executives, but most pulled the covers over their heads and continued their blissful, innocent sleep, oblivious to the looming danger.

Two new worms wreaked havoc on the Internet in the autumn of 2001: Code Red and NIMDA crystallized the issues of software homogeneity, the difficulties associated with security patches and the challenges associated with keeping the business traffic flowing while screening out the bad stuff.

Through all of this, organizations continued their quest to increase speed to market, reducing total cost of ownership (TCO) and maximizing return on investment (ROI). They paid a bit more security tax (in other words, they accepted the need to add a little more security – many consider security a tax because they receive no tangible benefit from it), but most simply wanted the problems of security to go away. The black arts of cyber security were protecting organizations well enough, but most executives wanted nothing to do with the concepts, much less the details.

Despite the wolves that howled at the door, many ignored the din and slept on. Then September 11, 2001 happened.

The World Trade Center lay in ruins and the gash in the Pentagon in Washington DC continued burning. The FBI issued a grave warning that deadly cyber attacks could commence at any time.

They never did (on September 11 anyway), and in fact it was a slow virus day.

But September 11 did accomplish one thing. Everybody woke up. *Terrible and catastrophic things could happen to us. Airplane-missiles, bombs, bioterrorism, cyberwarfare. If the symbolic pillars of world capitalism can be brought down in broad daylight while we stood around and helplessly watched, then certainly everything else is vulnerable, and any unspeakable, unimaginable act is possible.*

The slumber was over. It was replaced by a pounding headache and rising panic, but in between the throbs everyone was thinking the same thoughts: *must ... fix ...*

*vulnerabilities ... must ... be ... secure ... must ... defend ... against ... attacks ... must ... keep ... the ... enterprise ... running ...*

Executives in every industry are being asked the hard questions about information security:

- Where is the organization's data?
- Who has access to the organization's data?
- Who can change the organization's data?
- Will we know if/when we are attacked?
- Do we have firewalls?
- What can be found in the audit trails?
- Is the organization's data secure?
- How is the organization's data secured?

Executives and other non-technical business decision makers are being asked the hard questions. Many do not have the answers. They do not even presume to understand the concepts – or the technologies – associated with protecting information assets and infrastructure. But after Code Red, NIMDA and September 11, executives are painfully aware that they must learn and understand much more about information security.

## **THE PRIORITY OF INFORMATION SECURITY**

How important is information in the organization? How important is *securing* the information in the enterprise? In spite of its importance, security cannot be the number-one priority in the enterprise. If it were, the organization would never build new product or service capabilities, but only secure those it already has. Speed to market is nearly always the number-one priority. Speed to market and security are generally thought of as competing objectives – how can they be made to co-exist peacefully?

### **Security spending is rising**

Information security spending is rising. According to a Gartner Group survey, security spending increased 27 per cent in 2002 over 2001. In tight economic times, this is a matter of enormous significance. This is a bellwether that indicates that the events of 2001 got the attention of business decision makers, so much so that they significantly changed spending priorities. In most organizations, IT spending is flat, so the increase in security spending means that other items were reduced.

## IMPACT OF 2001 EVENTS

The security events of 2001 seized our attention either because they caused business disruption in our own organizations or because we believed that the events that occurred in other organizations could just as easily have happened to us.

### **World Trade Center attacks**

While it is unlikely that airplane-missiles will be used again in the near future, the annihilation of several enterprises' business operations jolted the business continuity planners of the corporate world. Few had accounted for the sudden loss of significant numbers of employees – many have since gone back to the drawing board to begin planning scenarios involving losses of large numbers of key employees.

Business continuity and risk management are beyond the scope of this book but are covered fully in the Financial Times Executive Briefing *Minimizing Enterprise Risk*.

### **Internet worms and viruses**

The two noteworthy worms of 2001, NIMDA and Code Red, caused widespread disruption and damage, and together cost over US\$1 billion in lost productivity. What contributed to these events being so significant?

- *Poor system administration practices.* NIMDA and Code Red exploited known weaknesses for which patches ('service packs' developed by manufacturers of operating system and application software to fix problems associated with functionality and security) had been in existence for a year or more. But NIMDA and Code Red spread rapidly because vast numbers of production web servers (both NIMDA and Code Red exploited weaknesses in Microsoft's Internet Information Server (IIS) web server product) were not patched and were therefore vulnerable.
- *Primitive software patch architecture.* While much of the blame for the failure to patch production systems rests in either ignorance or laziness, so too does much of the blame lie in Microsoft's woefully inadequate patch architecture. It is inadequate for three reasons. First, no tool exists that can be used to list which patches (security and otherwise) have been installed on a system. This makes it more difficult to identify systems that may have a specific vulnerability. Second, patches, once installed, cannot be 'backed out' (removed). With few exceptions, the only way to remove a patch is to delete and completely reinstall (and then reconfigure) the software product. For most

organizations this is completely unacceptable. Third, systems which were patched with security fixes are subject to those fixes being overwritten when other patches or upgrades are applied.

- *Software homogeneity.* NIMDA and Code Red ripped through the Internet because there were so many identically vulnerable hosts available. The problem with software homogeneity, where all systems in an enterprise have the same components, is that all of the systems have the same vulnerabilities. Should a worm, virus or other infection find a loophole in one of the systems, there is a great danger that the infection will spread to all systems in the enterprise. While standardization is more efficient from a TCO perspective, it does have its vulnerabilities.
- *Nonessential and vulnerable services active.* For an invasion from a worm or virus to be successful on any given system, the software service that is vulnerable must actually be running. If the vulnerable service is not running, the worm or virus cannot infect the system. The ease of software installation, together with the bundled installation of related software components, has led to scores of systems with nonessential and vulnerable software components installed and running. For instance, an administrator installing IIS might accept a 'default installation' configuration that may install many other software components that may never actually be needed or used. Weaknesses in any of these unnecessary components will nonetheless make the entire system vulnerable to an attack via those components.

## **PROLIFERATION OF EXTRANETS**

Prior to the 1990s, organizations were isolated islands. Their internal networks connected various parts of the organization, but rarely were organizations connected together.

In the 1990s, organizations began to connect themselves to the Internet. They installed firewalls and became connected islands with heavily defended entrances. Organizations' line of business applications remained isolated and protected by perimeter defences and only the organization's employees had access to them.

In the late 1990s, in order to build more efficient value chains, organizations began to open access to their applications to people in other organizations. While this practice can build many desirable efficiencies of production and customer service, it can also create unexpected problems. First, many legacy enterprise applications were not designed with multi-organizational access in mind: they do not have the granularity of access control that may be needed in this new setting. For instance, a customer with access to an enterprise application may be able to view records associated with *other* customers. Second, creating openings in

corporate firewalls for legitimate access increases the likelihood of unintended events such as worms, viruses or hacking attacks.

## **INSIDERS: THE REAL THREAT**

In the 1990s, the emphasis on securing the enterprise lay primarily in building a strong perimeter defence. But research by many organizations (including Gartner) indicates that the vast majority of security incidents occur within the perimeter, where security has traditionally been lax. Bill Cheswick, formerly of Bell Laboratories and co-author of *Firewalls and Internet Security: Repelling the Wily Hacker* (Addison-Wesley, 1994), has compared corporate networks to a well-known confectionary product: being hard and crunchy on the outside, but with a soft and chewy centre. By this he means that networks have traditionally had strong perimeter defences, with few security mechanisms protecting information assets within the enterprise.

This has led to the adoption of a ‘defence in depth’ paradigm wherein an organization will protect its data with a variety of mechanisms. The defence in depth protects in two ways: there is still protection in the event that one mechanism fails, and a combination of mechanisms is far more difficult to penetrate than a single one.

Firewalls and intrusion detection systems are powerless and largely irrelevant when it comes to protecting information assets within the organization. Instead, new security mechanisms must be integrated into enterprise applications and the rest of the infrastructure. These mechanisms include stronger authentication, access control, audit trails, host-based intrusion detection and encryption technologies, all of which serve to make information available to those who have a legitimate need to access it, while keeping others safely away.

## **UNPRECEDENTED DEPENDENCE ON INFORMATION TECHNOLOGY**

Businesses are more dependent than ever on information technology supporting day-to-day – even minute-to-minute – operations. Gone are the days when information technology made manually based business processes more efficient or cost-effective, where the manual versions of these processes could be used if the technology was temporarily unavailable.

Modern information technology is the enabler that allows business processes to take place. When technology fails, business operations cease. Should the technology be unavailable for more than a few hours, the business is at risk of serious disruption, revenue impacts and even failure.

This dependence on technology for core business operations heightens the need for information security. Business continuity plans now need to fully address information technology operations as a core business component.

## **SUMMARY**

Information security, while always a respected discipline, was a ‘backburner’ concern in most industries until 2000 and 2001, when a number of notorious events thrust it into the bright light of day. For some organizations it was one of the viruses or worms such as Melissa, ILoveYou, BackOrifice, Code Red or NIMDA; for others it was the terrorist attacks on New York and Washington DC on September 11, 2001. For nearly all technology-enabled companies, one of the above was a day of infamy.

Most senior executives enjoyed the ignorant bliss of business-as-usual regarding information security until this point, but many have been called by CEOs or boards of directors to explain whether their information assets were adequately protected in light of recent events. Many found it difficult to answer these questions succinctly, and many were hesitant to ask those in the organization responsible for information security about their protection because they did not know what questions to ask.

But it’s not just the sensational events described here that have given information security more visibility – changes in business practices that have further leveraged information technology have made securing information assets all the more important. Primarily this includes the explosion in the use of extranets to build ‘connectors’ (our casual term to mean any sort of electronic access or automation) between organizations. It is also becoming more apparent that hackers and other outsiders constitute the minority of threats on a global scale; instead it is the organization’s own workers who are precipitating most security events by betraying the organization’s trust in them.

Information security is also growing in importance because information technology is moving closer to the centre of businesses. Where at one time IT enabled the efficient or rapid execution of processes, in more and more instances IT *is* the process. Because IT plays such a central role in businesses’ operations, it is all the more important to ensure that information assets and infrastructure are protected from harm.

# Threats and vulnerabilities

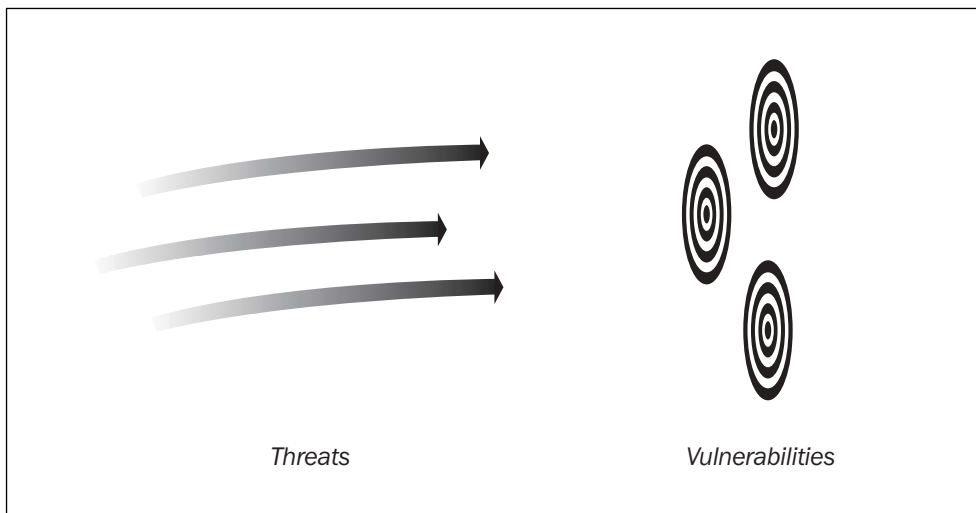
- Introduction 11
- Threats 11
- Vulnerabilities 17
- Summary 25



## INTRODUCTION

The risks associated with information security can be classified into two categories: *threats* and *vulnerabilities* (see Figure 2.1). Threats refer to the actions of people and nature that endanger an organization's information assets and infrastructure. Vulnerabilities are the weaknesses in the assets and infrastructure that are at risk of unintended and unwanted events.

**Fig. 2.1** Threats and vulnerabilities



Rather than being unrelated, threats and vulnerabilities are two sides of the same coin: threats are the potential actions that will follow the path of least resistance to the greatest vulnerabilities.

## THREATS

A security threat is the wilful intention on someone's part to inflict injury or damage to an individual's or organization's networks, computers, software or data. Threats come from people in the organization itself: employees, contractors and visitors. People outside the organization also threaten it.

The types of injury or damage that could occur are practically limitless. A few examples include:

- sabotage of computer hardware or software
- theft and subsequent disclosure of proprietary or personally sensitive information
- attacks on information infrastructure to render it unavailable for legitimate uses
- development and release of a virus or worm intended to cause widespread damage.

These and other threats are explored in more detail below.

## Employees and insiders

The people hired by the organization to carry out its business have access to its information assets and infrastructure. To a very large degree, the organization must simply trust its employees to handle its information properly.

Unauthorized disclosure of proprietary information is an age-old problem. From time to time, some individuals will betray the trust that organizations have placed in them and they will disclose proprietary information to an outsider. For instance, an employee might sell vital trade secrets to a competitor, leak sensitive information about the organization to the press or share company financial information with friends in advance of public announcements.

Advances in information technology have magnified this threat. In the 1960s, for instance, an employee walking out with the plans for a big project would have to physically carry large volumes of information or blueprints in boxes. Today, that same employee can put this information on a CD-ROM or e-mail it to someone outside the organization. Further, information theft is much harder to detect, since information that is stolen from a computer is still there, unlike boxes of files or blueprints. The employee leaking information can do so with the ease of sending e-mail; they could even do so anonymously by sending it via any of the free e-mail services available on the Internet.

In addition to disclosing information, the other significant threat is the wilful destruction of information and information infrastructure. There is almost no end to the variety of actions that disgruntled employees can take: they can alter information, possibly in subtle ways that will not be discovered for years (if ever) or modify software programs to change vital calculations or reporting mechanisms. They can sabotage computer systems or network devices in order to cripple operations, or plant a 'time bomb' or 'logic bomb' that will damage information or programs at some later time (presumably well after the employee has left the organization).

## Viruses, Trojan horses and worms

Viruses, Trojan horses and worms fall into a category called malicious code, or *malware*. They are so called because these programs are written for the sole purpose of destroying their intended targets. Growing in sophistication and speed, these destructive automatons spread throughout organizations and throughout the Internet like rabble-rousing street gangs throwing bricks through windows for the thrill of it.

- A *virus* is a piece of computer code that attaches itself to a program. When the program is run, the virus code is activated. Whatever instructions are contained in the virus are performed – from erasing or altering files to changing system configurations to displaying messages on the screen. In order to spread to other

computers, viruses replicate themselves by inserting themselves into other program files. Prior to the explosive growth of the Internet, viruses spread mostly through diskettes when one computer user shared files with a colleague. Beginning in the late 1990s, viruses rode the rails of the Internet by propagating through e-mail mechanisms and websites. Another innovation in viruses emerged in the mid 1990s in the form of a document macro virus – a virus that is embedded in a document such as a Microsoft Word, PowerPoint or Excel document. Merely opening such an infected document will cause the malicious code to execute and cause damage. Adobe PDF files, JPEG and Macromedia Flash have been similarly exploited.

- A *worm* is a program that is designed to propagate itself from system to system, while at the same time possibly wreaking some type of havoc, such as altering installed programs, changing system configurations, erasing files or installing a ‘bot’ (short for robot) in order to remotely control the computer at a later time. Some notable examples of worms are NIMDA and Code Red. Worms propagate themselves by taking advantage of known vulnerabilities in a program that is installed and running on the system. The programs that are most often exploited are Microsoft Internet Explorer, Microsoft Outlook and Microsoft IIS.
- A *Trojan horse* is a program that masquerades as another program. Trojan horses most often spread as attachments in e-mail messages that purport to be something they are not. For instance, an unsolicited e-mail from an unknown person might read, ‘Check out my holiday pictures!’ or ‘Here is that killer CV I was telling you about’. If the recipient is inattentive, they might double-click the attachment without noticing that the attachment is not a photograph or a CV but an executable program (an ‘EXE’) that might perform any number of malicious deeds on the victim’s computer. Typically, though, Trojan horses of this type will send copies of itself to all the recipients in the e-mail address book on the computer, or it might install a ‘bot’ on the computer that will permit the perpetrator to contact the PC at a later time and control its actions. Trojan horses propagate themselves relatively slowly because they depend upon a user to take some action, such as opening an e-mail attachment. A well-known example of a Trojan horse is BackOrifice, a program that permits a hacker to remotely control a victim’s system at any time in the future.

## Hacking

This is a very broad threat category, with a wide variety of activities and impacts to the organization. The common theme, though, is that of individuals or groups using personally owned computers to access an organization’s networks and

servers in order to disrupt, disclose or damage the organization's information or information infrastructure. Some of the methods used are discussed here.

- *Scanning* is a technique often used by hackers to search through the Internet – or through a particular organization – for systems possessing known exploitable weaknesses. There are a variety of scanning tools available that can quickly examine an individual system – or all systems on a particular network – and record the results for later use. Scanning can be thought of as the hacker's forward scouts on reconnaissance searching for easy targets.
- *Website defacements* result when hackers have discovered an exploitable weakness on an organization's web server. After discovering the weakness, the hackers can control the server and replace its contents with their own. Depending on the motivation, the hackers may install defamatory content that is intended to humiliate the web server's owner. Defacements are done for a variety of reasons. The hackers may simply be looking for an easy victim to add to their 'kill' list. Or they may wish publicly to embarrass or harm the organization that owns the website.
- *Stealing information* is a common motivation among hackers. The most easily marketable information to steal is that from credit cards as well as private information such as government ID numbers. Hackers can sell credit card numbers easily. Government ID information can also be sold to individuals or groups specializing in identity theft. Hackers with a more deliberate corporate or political espionage motivation may steal information of nearly every sort that can be found on a web server or on other systems in an organization.
- A *denial of service attack* (DOS) is a blockade committed in cyberspace. The attacker can flood a website or other network server with so much traffic that the attacked system(s) can become totally unreachable by legitimate users. At worst, the attacked system will 'roll over and die' after being completely overwhelmed by the attacking traffic.
- A *distributed denial of service attack* (DDOS) is a deadlier and more sophisticated version of the denial of service attack. A denial of service attack is generally perpetrated using a small number of computers, each specifically and manually instructed to send a flood of messages to the victim site. With a distributed denial of service attack, the perpetrator is using a master program to control hundreds or thousands of unsuspecting systems, instructing each of them to attack the same host. The unsuspecting systems are so described because they are the desktop and server systems in ordinary organizations – at some earlier time they were successfully attacked by a virus, worm or Trojan horse which implanted an ever-listening, ever-obedient attack program awaiting instructions from the master on who to attack.

- *Password attacks* are used to ascertain the password for a given user account. Password attacks take two forms: the hacker might take information known about the person (spouse, children, pet names, significant dates, etc.) and make educated guesses as to what the password might be. This is also called *password guessing*. The second form of password attack is one where the hacker has obtained the file(s) containing the encrypted account passwords on the system and will use a *password cracking* tool systematically to check each possible password until the correct one is found. This is also known as a *brute-force password attack*, which can take hours or days. An important distinction between password guessing and password cracking is that the hacker must use some means to break into the system, assume administrator (also known as root) privileges and obtain the encrypted password file(s) in order to engage in password cracking. If hackers cannot obtain the encrypted password file(s), they are limited to password guessing. It is for this reason that organizations insist that their employees do not use easily guessed words or phrases for their passwords.
- *Social engineering* is the term that describes methods used to obtain information needed to access an organization's computer systems. Generally, the hacker will make telephone calls to unsuspecting employees in order to get bits of information from them or to get them to perform certain tasks.

Social engineering is best illustrated by an example.

An intruder places a telephone call to an employee in the organization. 'Hello, Cynthia? This is Roger in marketing. I'm a new employee and one of my colleagues said you could help me. We're both on the road and do not have the telephone number for the remote access server. We need to dial in and get a mail message from Gordon Banks, the CFO, to finish up a presentation we are giving to a big customer in an hour. Can you help?'

Let us examine what 'Roger' has done. He has called Cynthia, possibly also a new employee. Roger sounds like an important person on an important mission – to land a new customer. Most people would be inclined to help Roger, to be part of the winning team. Because the presentation starts in less than an hour, there is not sufficient time to check things out through normal channels (for instance, asking Roger to call the helpdesk, which may have a long hold time, and besides, what harm can there be in giving out a phone number?).

Now our intruder places another phone call, this time to the helpdesk. 'Hello, my name is Gordon Banks [*yes, the CFO*]. I'm trying to dial in to pick up e-mail, but I cannot seem to log into the remote access server. I'm in an airport and don't have much time, and I need to get a vital message from the CEO before my next flight.'

After asking 'Gordon' a few questions, the helpdesk person resets the remote access password. Now our intruder has the remote access dial-in number, a userid (user identification) and a password. He might make a similar phone call to get a

password reset for an internal user, any of whom would have access to e-mail, internal websites and file servers. That phone call could go something like this.

‘Hello, my name is Phil McCann in Finance. I can’t seem to log into my NT account for e-mail and file servers, can you help me? My userid is ‘pmcann’. Can you watch me try to log in now? *[Since our intruder has remote access to the network, he can try this while the helpdesk person is watching. The helpdesk person probably thinks that Phil is in one of the campus buildings.]* Okay, you’ve set my password to ‘winter03’? Let me try to log in again ... yes, I’m in. Thanks a lot!’

Now our intruder has e-mail access and can browse the internal websites and trawl for information from the organization’s file servers. In just three seemingly innocent and legitimate phone calls, our intruder was given all the information he needed to infiltrate the organization’s network.

The secret to social engineering is to get small, seemingly harmless bits of information from unsuspecting individuals and to fit those pieces together in order to gain access to internal networks and information. Social engineers can also target other potential objectives. They might try to steal service from an Internet service provider or a mobile communications carrier, steal e-tickets from an airline or transfer funds from other people’s bank accounts.

## Hacker profiles

Hackers fall into two distinct categories (although individuals span the entire spectrum of expertise):

- *Script kiddies* are unsophisticated, non-technical people who uses pre-built tools to exploit weaknesses in software. They generally do not understand the technical details of the vulnerability they are attacking – only that their tool works. Script kiddies are completely dependent upon others for the knowledge to discover and understand weaknesses and to build the tools that exploit them. One should not underestimate the destructive potential of script kiddies, however. Some of the tools available to them can cause widespread destruction.
- *Professional hackers* have a deep understanding of TCP/IP network protocols and of one or more operating systems or software applications. They know how to experiment with a given software program in order to discover a weakness within it. These people, or others with whom they have shared the details of their discovery, will write a tool to exploit the weakness at will. It is this type of tool that is used by script kiddies.

Oversimplified, professional hackers and script kiddies can be compared to engineers who design and build easy-to-use weapons, and the large numbers of mercenaries who use them.

Thus far an important role has been left out – that of the academic information security researcher. These individuals, who possess a deep understanding of software and protocols, sometimes dedicate their energies to better understanding and experimenting with application software, operating system software and network protocols. If a security weakness is discovered, the researcher will publish a paper describing the specific details of the weakness.

One could argue that such research papers are reckless and contribute to the plague of hackers and script kiddies. However, we need to understand that research papers contribute to the improvement of software and protocols, and improvement is what we ultimately want. Were it not for the grad student researchers who discover and publicly describe their findings, the professional hackers would discover, too, but would not publicize their discovery. It is only by disclosing and publicizing weaknesses that all parties – researchers, developers and users – can learn these lessons and use them to build better software and protocols in the future.

## Threats summary

Threats are the potential acts perpetrated by those who wish to bring harm to an organization. Because of their access to key information (financial and customer records, strategic plans, policies and procedures), employees and contractors are in a position of being able to inflict serious damage on the organization's information and information infrastructure. Viruses, worms and Trojan horses are among the most visible threat. Organizations that are ill-prepared will experience significant business disruption during an outbreak.

Hackers, whether professionals or script kiddies, can cause damage through publicly visible (and embarrassing) defacements. They can steal information and then post it on Internet websites, sell it to competitors or use it to commit identity theft. They can launch denial of service attacks that can effectively shut down an organization's Internet-facing online business operations.

Hackers can carry out social engineering actions that can enable them to penetrate the organization's defences and get inside. Social engineering is the art of eliciting small amounts of information from several naïve company employees in order to build a composite 'picture' needed to penetrate the organization's defences.

## VULNERABILITIES

A vulnerability is any weakness in computer or network hardware or software that makes it open to attack or damage. A vulnerability can be the result of an imperfection in design, implementation or configuration. While a vulnerability is

generally thought of as an oversight, the existence of a vulnerability can be the result of a deliberate act. Some examples of vulnerabilities include:

- a flaw in a software program that permits an intruder to cause the program to malfunction, generally with the intention of breaking into the system running the program
- an operating system misconfiguration that permits an ordinary user to switch to privileged mode – which gives the user full administrative control over the system
- a flaw in a business process that permits an employee to log in using a new employee’s account by entering a well-known default/initial password
- a recently installed system with default administrative passwords, permitting anyone with knowledge of the password to gain full access to the system
- servers in the enterprise that individuals set up on their own that lack anti-virus protection and security patches.

These and other vulnerabilities will be discussed in more detail below.

### **Weak user account management processes**

As Chapter 5 will illustrate in more detail, the quality of information security is largely dependent upon people and processes, not technology. Nowhere is this more true than in the processes that support user account management (UAM). The organization’s first and best line of defence for protecting its information assets and infrastructure is each user’s login account userid and password, and secondarily the authorization and access control management that determines the capabilities of each user account in the enterprise.

If the processes supporting UAM are not well engineered or executed, then the enterprise will find itself in a situation where it does not know which employees have access to which specific information. If this is the case, there is no technology that can be put into place to protect information and infrastructure.

Specific examples of weak UAM processes include:

- *Lack of consistent userids on different platforms.* For instance, on one system, John Smith’s userid might be *jsmith*; on another system, Jane Smith’s userid might also be *jsmith*. If John and Jane later are issued user accounts on each other’s systems, they might be *josmith* and *jasmith*, respectively. This inconsistent use of userids will cause confusion for John and Jane, for system and network administrators who assign access rights for John and Jane, and for the UAM support staff.
- *Lack of consistent application of processes.* All too frequently, organizations are good at creating user accounts but lack the follow-up or procedures to

remove user accounts when employees or contractors leave the company. In a large organization, this can mean that many of its platforms will have thousands or tens of thousands of accounts that should have been removed.

- *Lack of authorization and access control processes.* In today's mobile workforce, employees move from function to function within an organization every few months to every few years. Seldom does an organization make the effort to make the necessary changes to an employee's access privileges when they have moved to a new function within the organization. As a result, an employee who has been with the organization for many years may find that they have the accumulation of access privileges acquired from each job function along the way.
- *Lack of adequate recordkeeping.* Over time, a UAM function in a large organization will have issued tens of thousands of user accounts in its portfolio of applications and systems, and made far more authorization and access control changes in the same period of time. Most organizations lack a rigorous recordkeeping capability that would allow them to quickly ascertain which employees have accounts on which systems and which have what accesses and authorizations.

The user account management function is labour intensive, tedious and time consuming, and yet the integrity of the organization's information and infrastructure depends to a great degree upon quality processes and procedures in the UAM function. Lacking this, an organization will have difficulty ascertaining the answers to even simple queries such as, 'To which systems does John Smith have access, and what access privileges does he have?'. The organization that cannot quickly and confidently answer such questions has lost control and cannot guarantee the integrity of its information.

## Software bugs

Software bugs represent a significant source of vulnerability in an organization. While most software bugs cause only improper functionality, some provide an intruder with the ability to break into the system running the software and gain a toehold that could lead to further penetration into the organization. Examples of software bugs include:

- *Buffer overruns.* This is an all-too-common vulnerability wherein a program that accepts input from any source fails to check the size of the input before storing or using it. If the chunk of data sent to the program is larger than the storage location that the program has prepared for it, the excess data will overrun the storage location and will overwrite other data, or even instructions, in the program. For example, a domain name service (DNS) server may be

designed to accept queries that are up to 1000 characters in length. Suppose that a query of 1200 characters is sent to the DNS server. If the DNS program does not first check to see how many characters are contained in the query, the excess 200 characters will overwrite some other program storage, or possibly even some of the program's instructions. This failure to check for the size of input data, and the subsequent storing of that information resulting in the corruption of other data or program instructions, is called a buffer overrun. This is because an input buffer has been overrun by a chunk of information that is larger than the storage location set up for it. The buffer overrun is the most common and preventable software bug. A large proportion of hacker attacks and Internet worms exploit buffer overruns.

- *Boundary conditions.* Programmers frequently fail fully to test the 'boundaries' of possible values when building their applications and tools. The most famous boundary condition is the two-digit year '99' and '00' – in other words, the Y2K problem. Boundary conditions come in all shapes and sizes and are almost limitless in the variety of unintended consequences.
- *Calculation errors.* Reduced to its fundamentals, computer software is about mathematics and calculations. Because they are invariably complex, software programs can be difficult to test and verify for correctness. Anyone who doubts this should be reminded of the likely cause of the failure of the Mars Polar Lander in 2000: one portion of the craft's software was calculating in imperial measurements and another portion used the metric system. The result of this overlooked mistake doomed the US\$165 million project from the start. Errors in operating system, database and application software can result in almost any imaginable consequence.

## Unpatched systems

The manufacturers of operating system and application software develop 'patches' (also known as 'hot fixes' or 'service packs') from time to time to fix problems associated with functionality and security. These patches can be downloaded from the vendor's websites and installed, usually fairly easily.

However, a surprisingly large proportion of organizations have not installed patches on their production systems, even when publicity about the vulnerability and the patch to fix it makes headline news. This widespread vulnerability directly contributed to the successful and rapid spread of worms such as Code Red and NIMDA in 2001. Moreover, while the worms in 2001 brought about a heightened awareness of the importance of staying current with security patches, in all probability there are still many organizations that continue running unpatched servers.

## Systems with no or outdated anti-virus protection

The proliferation of worms that propagate by scanning networks for fresh victims has raised the bar in terms of anti-virus protection. Even one system lacking anti-virus protection can wreak havoc on a network should such a worm infect that system – that one system will begin scanning the organization’s network for more victims.

The risks associated with the lack of anti-virus protection are similar to those of systems without security patches: these vulnerabilities leave the system open to attack by worms and viruses. An attack from any of these can result in disclosed, altered or destroyed information, as well as interrupted availability to people or other systems that depend on the system.

## Configuration errors

Errors in operating system and application configuration can leave a system open to attack. System, network or database administrators perform most configuration changes on systems manually. A lack of attentiveness or knowledge can turn what is supposed to be a simple change into one that exposes a new vulnerability to future exploitation or results in an immediate degradation in availability. Some examples of configuration errors include:

- incorrect permissions assigned to programs, files or directories – these can expose information or functionality to unauthorized individuals
- mis-spelled administrative commands – ‘typos’ can result in the wrong commands being entered, incorrect options or the wrong files or directories affected by commands
- unnecessary network services – these can needlessly expose a system to attack should a vulnerability be discovered in an activated, but unnecessary, network service.

## Unauthorized modems

Modems that can be easily set up and connected to analogue lines can spell real trouble for an organization. A modem connected to a workstation is an unprotected point of entry for a would-be intruder. These modems can be easily detected with ‘war diallers’ – software tools that dial every phone number in a given area and record all modems that they find. Intruders can then dial into the discovered modems manually to see whether any of them are configured or connected in a way that provides easy access to the organization’s network.

Because of these vulnerabilities, many organizations have established policies prohibiting the use of unauthorized modems, as well as analogue lines themselves. Organizations can use war diallers to search for unauthorized modems on their own networks.

## Wireless networks

Wireless networks in organizations are proliferating because of the unprecedented convenience and performance they provide. Wireless network access points, as well as the client hardware, are easy to set up. The access points themselves can be plugged into any live network jack in an office, conference room, data centre or lab. The ease with which they can be set up, together with identified vulnerabilities in the encryption protocols, have made wireless networks, most notably 802.11x, a source of significant risk to the enterprise.

The risks associated with 802.11x cannot be overstated: the common practice of using default settings, together with the fact that the WEP (Wired Equivalent Privacy) protocol has been broken, has led to a new phenomenon called 'war driving' where individuals with laptops equipped with 802.11x devices and specialized tools such as Aircrack-ng can literally drive down a city street and discover all the wireless networks in office buildings. The tools they use are able to determine which networks have vulnerabilities that would permit the individual easily to penetrate the network.

Organizations, as well as the companies that develop and sell wireless network equipment, are becoming aware of the risks associated with wireless networks. Poorly designed or configured, they can become a new target for individuals who are determined to break into an organization's network.

## Rogue servers

A rogue server in the enterprise is one that exists, typically under someone's desk, outside the control and support of the IT operations group. Rogue servers result from business groups or departments that purchase and install operating system and application software on their own, usually for testing or for running small, unsupported applications. Because these systems are set up by non-IT professionals, they will rarely have anti-virus or security patches. When these servers have frequently-attacked software such as Microsoft IIS installed and running, they are vulnerable to attack from worms such as NIMDA and Code Red.

Many organizations discovered their rogue servers with the onslaught of the NIMDA and Code Red worms in 2001. These servers were attacked successfully and then became *attacking systems* within the enterprise. Organizations that had

up-to-date anti-virus and security patches installed found that the bulk of their problems lay in these rogue servers.

## Vulnerability to social engineering

The social engineering example in this chapter occurred because the employees lacked security awareness training in this particular area. Employees who receive calls from people requesting access information such as dial-in phone numbers should transfer the call to the helpdesk, which should be trained to identify employees.

Making employees aware of the techniques of the social engineer will help to give them pause when they receive such a call. Employees need to think before giving up even small titbits of information and overcome the natural tendency to help a colleague in distress and be a hero. The helpdesk staff should be well trained in the procedures of identifying employees when they call in. These procedures should include questions to callers that only the employee would know. While it is true that a careless employee could, at one time or another, share some of these ‘secret’ facts, good screening techniques will filter out a significant number of ‘bad actors’.

## Weak perimeter security

A typical enterprise network has a wide variety of ingresses, any of which can be a path of easy entry for a would-be intruder. While some of these vulnerabilities have been mentioned elsewhere in this section, it is useful to list them all here.

- *Internet connections.* In most organizations, this is the well-marked and lighted grand entrance. Most organizations are continuously bombarded by intruders’ tools that are compiling lists of vulnerable systems and sites.
- *Back-door Internet connections.* Larger organizations have other connections to the Internet, used for lab, testing, demo or remote access purposes.
- *Lab connections.* Development, testing or demo labs may have their own Internet connections, which may or may not have the same degree of protection as the main corporate Internet connection. From the perspective of the enterprise network, most labs should be considered ‘outside’ networks and the perimeter defined as the lab-to-enterprise network connection point.
- *Modems.* Whether authorized or not, modems present an opportunity for an individual to try to penetrate the enterprise network. A single individual can war dial even a large organization in their sleep – literally, using commercially available war dialling tools that not only identify all modems in the enterprise but also use a variety of means to identify and even penetrate the systems or networks to which they are connected.

- *Wireless LANs.* Access points for 802.11x local area networks (LANs) are inexpensive, preconfigured and can be plugged into any network jack in the enterprise. These access points, in effect, set up broadcast radio stations on the enterprise network, inviting all to bypass the enterprise firewalls and intrusion detection systems and connect directly to the inside. There are war driving tools available to scan for and identify all vulnerable networks. Wireless LANs are similar to – but a greater threat than – modems, because it is not necessary to dial phone numbers or get inside the building. All one needs to do is drive down the street with a laptop and war driving tools.
- *Extranet connections.* Network connections to partners, suppliers and large customers introduce significant vulnerabilities to the organization. These connections – at best – increase the number of ingress points that an intruder or worm can use to penetrate the enterprise. At worst, these connections are a literal extension of the enterprise network, but without the enterprise's control or defences. The level of security in the enterprise is effectively lowered to that of the weakest entities to which it is connected.

All paths of entry must be well protected. Any determined intruder is going to follow the path of least resistance – so don't expect a frontal assault on the main Internet firewall, but rather a sneak attack via an unguarded back door.

### **Vulnerabilities summary**

A vulnerability is a weakness in hardware, software or processes that exposes it to attack or damage. It can be a result of design, implementation or configuration imperfections. In addition, vulnerabilities can be planted deliberately or can be the consequence of carelessness.

Weaknesses in enterprise user account management processes can result in inconsistencies with userids, user accounts that are not deleted when employees leave the organization, user accounts with too many privileges, and the lack of or poor recordkeeping that hinders effective user account management and user account investigations.

Software bugs will cause a nearly infinite variety of unintended consequences, ranging from functional errors to vulnerabilities that permit an intruder to attack and take over the application or the entire system. Buffer overflows are the most common software bug that can expose a system to a hacker attack.

Systems that do not have security patches installed are vulnerable to mechanized attacks – worms, viruses and Trojan horses – and attacks from script kiddies and hackers looking for random or targeted victims. Systems with outdated virus protection (or none at all) are similarly vulnerable to attack and subsequent damage. Rogue servers – those that are unauthorized and unmanaged by the organization's IT group – are seldom patched and also may lack virus protection.

Systems and software configuration errors can sometimes result in a system being vulnerable to intruders. Configuration errors include incorrect file or directory permissions that permit access to data, and network services that permit inappropriate and unauthorized access. Unauthorized modems provide undocumented and seldom protected and unwatched access to the enterprise. An intruder with a war dialler can quickly find back doors to the enterprise network. Similarly, wireless networks present an opportunity for an intruder to war drive and subsequently discover and attack wireless networks.

Without adequate awareness and training, an organization can be vulnerable to social engineers, who take advantage of employees' willingness to help 'colleagues' in distress. Social engineers can gain full access to an enterprise network, right through its 'front door', if they can dupe just a few people into giving out titbits of information.

The enterprise's perimeter is growing, thanks to Internet connections, extranet connections, wireless LANs and labs. This presents a growing opportunity for intruders to penetrate the enterprise, should one or more perimeter defences be inadequate.

## **SUMMARY**

An understanding of threats and vulnerabilities will enable the business decision maker to make informed risk decisions to protect the enterprise. Knowledge of specific threats and vulnerabilities is the first step to objectivity when making security decisions.

The proper holistic approach to a better understanding of threats and vulnerabilities in the enterprise is a detailed risk assessment. A risk assessment will identify an organization's weaknesses in its business processes and its technology, and assign probabilities for the occurrence of events that pierce the vulnerabilities and disrupt business operations. For a more detailed discussion on risk assessment, refer to the Financial Times Executive Briefing *Enterprise Risk Management and Planning for Business Continuity* by Corinne A. Gregory.



# Security fundamentals – the principles and the mechanisms behind them

- Introduction 29
- Identification and authentication 30
- Authenticating other systems 33
- Authorization 38
- Access control 41
- Encryption 48
- Non-repudiation 60
- Integrity 60
- Audit 66
- Availability 68
- Security mechanisms work together 68
- Summary 69



## INTRODUCTION

The goal of information security is to protect an organization's information assets and infrastructure from accidental or malicious disclosure, modification, erasure and misuse. While people – especially the trusted people inside the organization – are the most important factor in information integrity and protection, the technology of security also plays a vital role. Technical controls protect information assets and infrastructure primarily from people *outside* the organization – those who are *not* trusted.

Security technology plays an important role for insiders, too, through access controls and audit capabilities. These help to reinforce accountability and also provide valuable information during investigations.

This chapter discusses the concepts, mechanisms and technologies used to protect information assets and infrastructure. The topics discussed are:

- authentication – identifying people who are allowed to access information
- authorization – controlling what information and functions a person is allowed to access
- access control – managing access to information and infrastructure
- confidentiality – obscuring information from prying eyes
- integrity – assurance of the accuracy of information
- availability – assurance that information is accessible when needed
- non-repudiation – proving the authenticity of a transaction and of its originator
- audit – recording events for possible subsequent problem solving, fact finding or investigation.

While these topics may be dreary to business decision makers, these are the domains in which an organization's security practitioners spend most of their time. The protection of the organization's information assets and infrastructure – the viability and integrity of the business itself – depends upon whether the focus and priorities of the security practitioners are sufficient.

### Depth of defence

Current best practices call for the protection of information with more than one mechanism. If information is protected with only one mechanism, a weakness, misconfiguration or successful attack on the sole mechanism will compromise the information. If, on the other hand, more than one mechanism defends the information, it will be more difficult for an attacker to reach the information because there will be multiple hurdles to overcome.

The analogy of the castle is used frequently to illustrate this point. Imagine the crown jewels are placed within a castle. The jewels can be defended in one of two ways: just the castle (with really high walls) or the castle plus several moats. The castle with moats will usually do a better job of protecting the information than a castle with high walls alone. If attackers can work out how to scale (or dig under) the high castle walls, they will reach the jewels. But a castle surrounded by multiple moats (each with monsters, of course) presents a better defence on account of its variety of defences working together to protect the jewels. This topic will be addressed again at the end of this chapter.

## **IDENTIFICATION AND AUTHENTICATION**

In its most familiar form, authentication is the term that describes how a person logs on to a computer or a network. However, the term ‘authentication’ is frequently misused and confused with other terms. This section will discuss authentication and identification and describe their differences.

These two terms are defined as follows:

- identification is the *assertion of identity*
- authentication is the *proof of identity*.

With identification, one’s identity is asserted and accepted without further proof. Apart from anonymity, where one’s identity is not known at all, identification is the lowest form of recognition. Identification is a weak and generally unreliable way of relating an asserted name to an individual. This is because anyone knowing someone else’s identity can assert himself or herself as that individual. It would be a poor practice for a bank to accept an individual’s assertion of identity, based only on their word. But that is what identification is. This is why there is a way to prove one’s identity: with authentication.

Authentication involves not just the *assertion* of identity but also *proof* of identity. When a person is authenticated, their identity is not accepted until they can prove their identity by also providing something else that they know, have or do. In the context of information systems, authentication is most often accepted with a userid or user name and a password or pass phrase. It is assumed that, while many individuals may know a person’s userid or user name, only the person associated with the userid or user name will know the password. When the person furnishes his or her userid *and* password, the system to which they are identifying themselves knows that this person is in fact who they claim to be.

The difference between identification and authentication can be further illustrated with an example. An individual seeks to enter a museum where he/she is a member,

and can enter the museum at any time by simply giving his/her name at the door. This is identification. In the strictest sense of identification, any individual can give that person's name and be granted entrance to the museum. The fact that someone else can assert the member's identity by merely speaking their name may be acceptable in this situation. But what if, upon entrance, the member also gets to shop in the gift shop and have the bill sent to their house? It would be easy to argue that identification is insufficient.

Suppose that the museum requires *proof* of the member's identity before admitting him or her. Perhaps the museum will ask to see an ID card or membership card, or perhaps the museum will ask the member for the member number or something else that only the member would be expected to know. When the member supplies this additional information, they are said to be *authenticated*. And if the member has shopping privileges, both the museum and the member will have much better assurance that only the member – and not an imposter – will be able to select items from the gift shop, given that the member has proven their identity to the museum's satisfaction.

Authentication is said to be *stronger* than identification. In the context of identification and authentication, the terms *stronger* and *weaker* are used to compare the relative strengths afforded by different methods used to identify and authenticate.

## Passwords: complexity and quality

The usual method for authentication involves a person supplying a userid (or user name) and a password. A password is a collection of characters, chosen either by the user or by the system. A vital characteristic of a password is its *complexity*. The complexity of a password refers to:

- its length (the minimum number of characters it must contain)
- the types of characters it requires (lower case letters, upper case letters, numbers and special characters).

There are other characteristics of a password that are as important as its complexity. The other characteristics are:

- number of days before the password expires
- minimum number of days allowed before a password can be changed
- whether previously used passwords may be reused
- the number of times an incorrect password can be entered before the account is barred from further login attempts.

These together define the quality of a password. The complexity of a password also contributes to its overall quality.

Why does password quality matter? There are two main reasons. First, a higher-quality password is harder for a colleague or intruder to guess. Suppose that a system's low-quality password permits as few as four ordinary characters, as well as no account 'lock out' after unsuccessful attempts. A worker who wishes to access a system using a colleague's userid can log in and guess the password, trying things such as the user's spouse or children's names.

The other reason that password quality matters is that a higher-quality password is less vulnerable to attack by hackers' tools called *password crackers*. A password cracking tool employs a *dictionary attack* (by using common dictionary words) in an attempt to guess an account's correct password.

An important consideration regarding password quality is attaining the balance between adequate quality and user acceptance. Low password quality standards permit passwords that are easily guessed, while password quality standards that are too high will result in cranky users and passwords appearing on 'sticky notes' pasted to monitors because they are too difficult to remember.

### It's what you know

Over the years, new authentication technologies have been developed that have gone beyond the traditional userid and password. This has led to the development of a fundamental concept illustrating how one gains access to a computer or application. A person logs into a system by proving that:

- they *know* something – this refers to a userid and password: a person logs into a system because they know the userid and password
- they *have* something – a person has in their possession a token, smart card or some other physical device that is somehow required to complete the login process
- they *are* something – this refers to some sort of biometric authentication that inspects a fingerprint, palm print, retina scan, pupil scan, voice print or something else that demonstrates that it is *they* who are authenticating and not someone else
- they can *do* something – for instance, they can sign their name or repeat a word spoken to them.

The technologies behind all of the authentication methods mentioned in this list will be explained in more detail later in this section.

### How many factors

The advent of these alternative authentication technologies has led to a new terminology that refers to the *strength* of the authentication. The buzzwords that have been around for many years are:

- *one-factor authentication* – the process of authenticating by providing one piece of information, or factor, usually a password or personal identification number (PIN)
- *two-factor authentication* – the process of authenticating with more than one factor such as password; logging in requires not only something the person knows but also something they *have* or *are* or *do*.

Some examples of two-factor authentication are as follows:

- the user must enter a userid, password and a numeric value displayed on a token card
- the user must enter a userid and password and must insert a smart card into the smart card reader attached to the computer or workstation
- the user must enter a userid and password and must place their hand in a palm reader.

Two-factor authentication is also called *strong authentication* because it employs a far more reliable way of confirming the identity of the person logging in.

## Burden of proof

The strength of authentication should be commensurate to the value of the information or function being protected. Generally speaking, an organization will employ two-factor authentication in the portions of its business that require the additional protection.

## AUTHENTICATING OTHER SYSTEMS

The discussion to this point has focused on authenticating people who wish to access information. More frequent than people-to-system authentication is system-to-system authentication. Distributed and layered application architectures have led to environments where a single person interacting with a server could invoke dozens of system-to-system transactions. In many cases, these systems are not all contained within a single organization.

Distributed and layered architectures have led to the necessity that systems authenticate to other systems. For instance, a front-end application that needs to perform a transaction with a back-end database should be required to authenticate with the database. Otherwise, it would be far easier for someone to build an unauthorized front-end that could generate forged transactions.

For readers who are not convinced that system-to-system authentication is necessary, consider a hypothetical situation where a person with a little inside

knowledge builds a system that sends forged credit card transactions to a bank or merchant system. Without adequate controls such as system-level authentication, these transactions could be difficult to trace.

## **Authentication mechanisms**

The mechanisms available for userid and password authentication consist of two types: host-based and service-based.

### ***Host-based authentication***

A host-based authentication mechanism refers to the components present on a single host used to log in to or access information or functions present on that host. With host-based authentication, each host has its own local userid and password database. Host-based authentication can be useful in very small organizations; however, as an organization grows and adds more systems, it will quickly find that user accounts are difficult and time-consuming to maintain. Keeping the authentication databases in perfect sync is an unattainable and resource-intensive endeavour. This is in keeping with the principle that states that several 'identical' copies of the same actively used database will not be identical.

In an environment with several different kinds of systems, it is likely that a person's userid and password will not always be the same on all of the systems he or she uses. This can be due to technical differences between operating systems – some may require longer or shorter userids than others, for instance. Also, the password quality requirements on different systems and system types will almost guarantee that a person will have many different passwords. To further confuse and frustrate the users their passwords will probably expire on different systems at different times.

### ***Service-based authentication***

A service-based authentication mechanism is designed to be used by several hosts – dozens to thousands – in an organization. In such an environment, the authentication authority function is centralized. Instead of having an authentication database on each host, each host is configured to use a centralized authentication server.

Service-based authentication is highly desirable in a large organization with many computing systems. Maintaining central user accounts that are used on multiple systems is more efficient than maintaining separate user accounts on all hosts. With service-based authentication, a person's userid and password will be the same for all hosts and applications participating in the authorization service. With fewer userids and passwords to remember, users are less likely to use poor passwords or write down their userids and passwords somewhere where they could be found and exploited by others.

## Authentication technologies

This section describes some common authentication technologies.

### *Host-based authentication technologies*

When host-based authentication is used, basic user account and authentication tools are included as part of the operating system. Those that are well known are discussed here.

- *UNIX/Linux*. Local accounts are stored in the ‘password’ and ‘shadow’ files and can be administered using a text editor or a graphical user interface (GUI) tool provided by the UNIX vendor.
- *Windows NT/2000/XP*. Local accounts are stored in the SAM (security accounts manager) file and are administered using a GUI tool provided with the operating system.

### *Service-based authentication technologies*

There are several service-based authentication technologies available. Some of these are built into operating systems, and there is a plethora of third-party products providing authentication services. The more common offerings are discussed here.

- *LDAP*, or Lightweight Directory Access Protocol, is an increasingly popular enterprise-wide authentication tool. LDAP is both a data model and a communications protocol: the LDAP specification describes the organization of stored data as well as the format for communicating over networks. LDAP can be the basis for an organization’s ‘employee directory’ as well as its authentication service, since it can be used to store passwords and other authentication-specific information. LDAP is available as public domain tools, as well as being a part of several commercial products.
- *LDAP metadirectories* allow an organization to integrate several LDAP instances into a virtual LDAP repository.
- *Active Directory™* is Microsoft’s implementation of LDAP. Microsoft has added some features and functions to the standard LDAP specification – organizations considering using Active Directory need to explore fully compatibility issues with any other LDAP mechanisms in the enterprise.
- *NT Domains* is an older Microsoft authentication mechanism. It is entirely proprietary, although Microsoft has published some of the interfaces, thereby allowing third-party products to utilize or extend its capabilities and interoperability.
- *NIS*, or Network Information Service, is a centralized authentication and system information tool originally developed by Sun Microsystems. NIS, and its more

recent successor NIS+, has been adopted by other vendors. For purely authentication purposes, NIS is losing market share to LDAP.

- *Token* authentication is one of the two-factor or strong authentication tools, employing the *something you know* and *something you have* concepts. A token is a credit-card-sized or key-fob-shaped device with a character display and perhaps a button. To authenticate, a user with a token enters their userid, a password or PIN, and the characters displayed on their token at that moment. It is infeasible for an imposter to pose as the owner of the token, unless they actually have the token in their possession, along with knowledge of the owner's userid and password. An organization using token authentication also has a server that stores information about all of its tokens. This server is queried every time a user attempts to log in.
- *Smart card* authentication, another two-factor or strong authentication tool, utilizes *something you know* and *something you have*. Like token authentication, smart card authentication requires that you possess information you know (the userid and password) as well as a physical object (the smart card) in order to log in and access information. A smart card is a credit-card-sized object that has a small chip containing a microprocessor and memory. Stored in the memory of the smart card is information about the user. To log in, the user must insert the smart card into a smart card reader. The reader, usually connected to a workstation, communicates to an organization's central server to authenticate the user.
- *Certificate* authentication is another two-factor authentication tool. A certificate, also known as a digital certificate, is issued by a centralized certificate authority to a specific individual. A digital certificate contains a user's *private* and *public encryption keys* and other identifying information. The public key allows a user to authenticate (to a server or application), as well as encrypt, decrypt and create digital signatures. A large organization using certificates assigned to individuals must also implement a PKI (public key infrastructure).
- *Certificate authority (CA)* is an organization that issues certificates. Some of the root CAs include Verisign, GT Cybertrust, Globalsign, Thawte and American Express.
- *PKI*, or public key infrastructure, is not really an authentication mechanism but a network service that is required in an organization that issues certificates to its personnel. A PKI is a repository of all individuals' public and private encryption keys. There is no single accepted standard for the storage of PKI information or the transmission of PKI queries – all PKI implementations are vendor-specific and proprietary.
- *Biometric* authentication refers to several similar technologies whereby a device attached to a workstation scans a fingerprint, palm print, pupil, retina or voice

in order to confirm the identity of a person who wishes to log in. Biometrics measures *what you are*. There are no biometrics standards – all solutions are proprietary. Most biometrics mechanisms attempt to integrate with existing one-factor authentication mechanisms.

- *Signature* scanners record a handwritten signature and measure the movements used to create one. Signature scanning is a way of measuring *what you do*. Like the other two-factor authentication tools, signature scanning integrates to some degree into an organization's one-factor authentication environment.
- *Single sign-on (SSO)* is a mechanism that several applications in an organization utilize, so that a user who logs into one of those applications can quickly switch to another participating application without having to log in again. In a single sign-on environment, all participating applications communicate with a central SSO server which keeps track of each user's sign-on status and controls whether a user must first sign on in order to access an application.
- *Internet-based authentication* refers to an emerging class of decentralized authentication services accessible from anywhere on the Internet. A notable example is Microsoft .Net Passport. Any organization considering such a managed authentication service needs to understand the issues resulting from surrendering and outsourcing such an important function.

Of the methods listed here, LDAP is the standard in widest use. NT Domains authentication, while not an open standard, is also broadly used. As organizations migrate from NT Domains to Active Directory, this will further the convergence towards LDAP. The other authentication methods listed here are all somewhat more proprietary and not as ubiquitous.

### ***How are passwords stored and verified?***

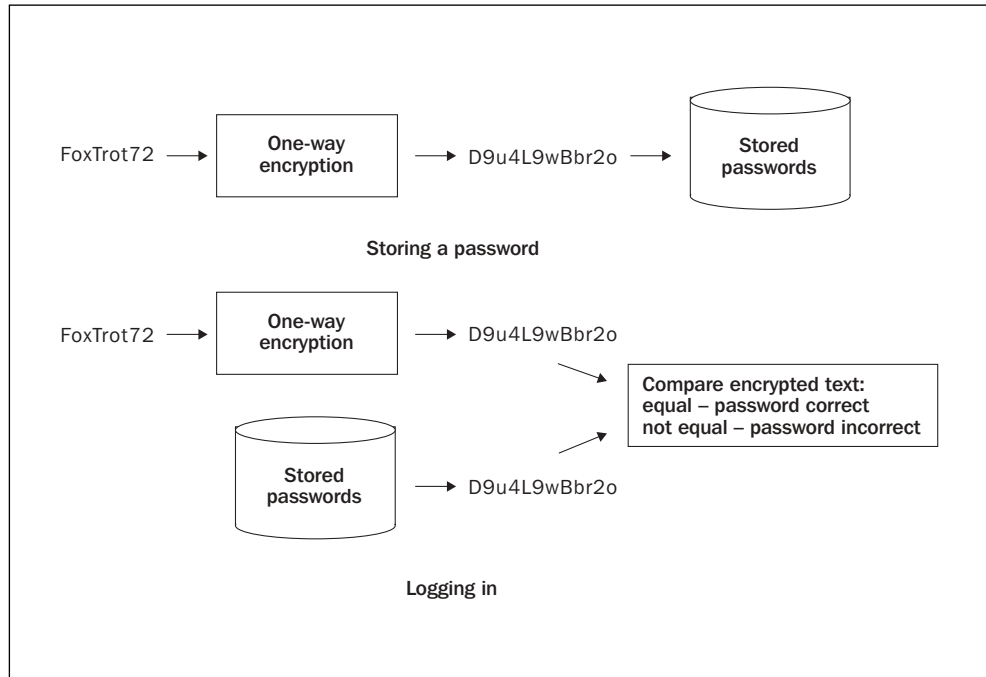
Passwords are stored as encrypted text (*see* Figure 3.1). The current method employs a 'one-way hash' or 'one-way encryption' mechanism. These 'one-way' mathematical algorithms have no known 'undo' whereby a stored encrypted password can be transformed back to its 'clear text' format. This is why a password cannot be recovered. Even a system administrator or engineer with access to a system's inner workings cannot discover or recover a password. When a user forgets his or her password, all that can be done is to assign a new one.

Some web-based services send users their passwords should they forget them: these sites have a 'Forgot your password?' button or link. Such sites are not storing user passwords using one-way encryption; in fact, they might not be encrypted at all.

If a password cannot be decrypted, how does a system know whether a user logging in has entered the correct password? The password that the user entered is encrypted, using the same algorithm used when storing the password in the first

place. The stored encrypted password is compared with the encrypted password the user entered – if the encrypted passwords are the same, the user entered the correct password and is logged into the system.

**Fig. 3.1** Storing and using passwords



## Authentication summary

Authentication is the cornerstone of information security. Access to information and information functions is granted based upon the proven identity of the person (or thing) requesting access.

## AUTHORIZATION

Authorization describes the concepts, mechanisms and technologies used to control the specific privileges and access rights of individual users.

## Why use authorization?

Authorization provides the essential control that governs which users are allowed to see (or modify) which data elements and to perform which functions. Left to their own means, people – either through mistake or misdeed – will see more and do more than they are allowed to.

Authorization and access control are the technological implementation of *separation of duties*. Nearly all business processes employ methods where people in different roles perform tasks associated with those roles. Organization data is (or needs to be!) classified and released to individuals on a need-to-know basis. Authentication is part of the means for controlling who can see what information. A practical example concerns the various roles associated with a modern corporate purchasing function, which are:

- user – uses the product or service
- requestor – requests a product or service
- approver – approves purchase; often approvers have *signing limits*, meaning they are permitted to approve purchases up to a specific amount
- buyer – purchases the product; this person issues the corporate purchase order and sometimes also chooses the vendor and specific product to be purchased
- receiver – indicates that the organization has received the product or service
- payer – authorizes a payment be made to the producer of the product or service.

In this example, each userid in the organization's financial system will be associated with only one of these roles. A mechanism inside of (or external to) the financial application will control which functions each user is permitted to use. This separation of duties is required to ensure the integrity of the purchasing process which, if not properly managed and audited, is likely to be abused.

Authorization and access control depend upon the integrity and strength of authentication. Strong authorization coupled with weak or flawed authentication is easily circumvented if users are able to assume the identities of others whose privileges they wish to exploit.

## Roles

For ease of administration, especially in larger organizations, authorization is managed through the use of *roles*. The role is the identity that is permitted to access certain information or perform certain functions. People access information or perform functions because their userids are associated with one or more roles.

The use of roles simplifies the administration of authorization. For example, a large organization increases the signing authority of all of its directors from £10 000 to £20 000. In an organization with 2000 directors, making this change can be a daunting task in the absence of roles. But in an organization using roles, the signing limit of all of its directors can perhaps be changed simply by increasing the signing limit for the role called director. Then, all individuals who are directors (and where their userid is associated with the director role) will automatically have their signing limit increased.

This kind of administrative control must, of course, be tightly protected. Only a small number of people in an organization will be permitted to change the privileges of roles and who is assigned to them.

## Authorization mechanisms

Enterprise-wide authorization is still an emerging phenomenon. Applications that implement authorization frequently do so within the confines of the application; few have provided interfaces to permit 'plugging in' to an enterprise-wide authorization service.

There is a critical mass of authorization service development forming around web-based application servers. The ubiquity of the web client and the movement towards web-based applications (away from client-server applications) have provided a natural location for the authorization control: on the web server platform itself, or on a nearby server.

Because controlling access to data and functions is so dependent upon knowing the identity of the user, authorization mechanisms are frequently integrated with authentication mechanisms. User credentials (their roles and what they are allowed to perform) are frequently stored in the same database as the user's account information (userid and password).

## Authorization technologies

Unlike authentication, where a rich set of standards is gaining acceptance and maturing, authorization technologies are still proprietary, not yet compatible with most applications, and more costly to integrate into the enterprise. Still, authorization is converging around some major technologies.

- *LDAP*. Quickly becoming the de facto standard for authentication, LDAP databases can also store (or point to) authorization credentials. Several of the major authentication products such as Netegrity and Oblix integrate with (or include) LDAP authentication.
- *Kerberos*. Developed many years ago at MIT (Massachusetts Institute of Technology), Kerberos still enjoys wide use as an authorization mechanism. It works on the principle of issuing tickets to authorized users who wish to perform certain functions or access particular resources.
- *NT Domains*. The older Microsoft proprietary NT Domain architecture provides a mechanism that controls which users are permitted to access which network resources.
- *Active Directory*. Microsoft's newer directory service is built on LDAP and Kerberos.

- *UNIX*. Host-based UNIX user and group permissions can control which users (and groups of users) are permitted to execute which applications and data.
- *NIS/NIS+*. These network service extensions of UNIX-based user and group permissions provide enterprise-wide control of application and data access permissions.

## Authorization summary

Authorization provides the identity and role-based access control that most organizations require. In large, complex environments, such as enterprise resource management (ERM) and human resource (HR) applications, authorization provides the separation of duties required to assure the integrity of information and processes.

Authorization technologies are largely proprietary – there are no standards in wide use. As a result, most enterprise applications manage their own authorization internally. Organizations with several large applications will find themselves managing multiple, unsynchronized authorization mechanisms.

## ACCESS CONTROL

Access control refers to a variety of technical means for limiting access to information and infrastructure. This is a vital part of information security, especially for organizations that are connected to other organizations, directly or via the Internet.

### Network access control

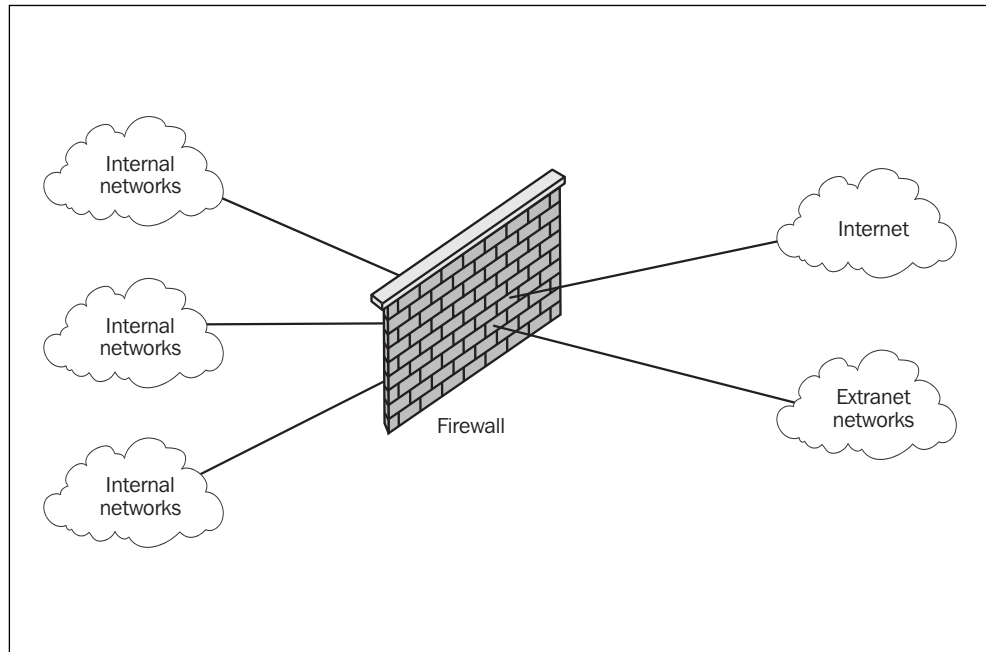
Every organization has controls in place to limit access to its physical facilities using locked doors and perhaps card-key technology or security guards at its entrances. Likewise, an organization with an information infrastructure must erect controls to limit access to its network. Every organization with public access (e.g. the Internet) must protect itself at the perimeter. Larger organizations, or organizations with several geographic locations, may also need to add protection within their borders.

### *The firewall*

The firewall is the best known – and still the best current practice – for achieving effective perimeter protection. A firewall is a coarse-grained filter designed to block or admit various kinds of network traffic (*see* Figure 3.2). For instance, a firewall can block e-mail to or from certain Internet sites, but it cannot block messages containing specific content. Also, a firewall can block access to certain

websites based on their domain name or IP address, but not on the content that the sites contain.

**Fig. 3.2** Network firewall



Firewalls are placed at the ingress and egress points in a network. In order to effectively protect the organization, all network traffic that enters and exits the network must pass through the firewall so that the firewall can examine each packet of information and decide whether that packet should be allowed to pass through or not. This is very similar to an international border checkpoint or a toll road's tollbooth.

It is important to understand what a firewall can and cannot do. The general capabilities are discussed here.

A firewall:

- blocks or admits specific *types* of network traffic; for instance, it may admit e-mail and World Wide Web traffic but block all other types of traffic
- blocks or admits traffic *to* specific hosts and networks and *from* specific hosts and networks; for example, it may admit traffic only to a web server and from an e-mail server but block everything else
- contains a set of *rules* that specify precisely which types of network traffic are permitted to and from which hosts and networks
- blocks all traffic except that which is explicitly permitted
- can keep a log of which traffic it admits and/or which traffic it blocks; it can specifically log successes or failures for just certain hosts or networks.

A firewall cannot:

- block or admit specific content
- block e-mail from certain individuals
- block viruses
- block SPAM (unsolicited junk e-mail).

Mechanisms for the above are available, but firewalls are not the place where these functions are performed.

An example set of rules for a firewall could be (in English):

- permit World Wide Web traffic from anywhere on the Internet only to the organization's Internet web server
- permit World Wide Web traffic from anywhere inside the organization to anywhere on the Internet
- permit inbound and outbound e-mail to and from anywhere on the Internet to/from the organization's e-mail server
- permit FTP (file transfer protocol, a tool used mostly by system and network engineers) from anywhere inside the organization to anywhere on the Internet
- log all unsuccessful attempts to enter the organization's network from the Internet
- block all traffic except that which is explicitly permitted above.

While the above list is typical for an organization, an actual ruleset would be longer than this, since there are additional network services required for an organization to communicate with the Internet.

### ***Firewall technology***

For many years, firewalls were built using a UNIX system that had two separate network connections – one to the Internet and one to the organization's internal network. Some add-on software was placed on the UNIX system that blocked or admitted traffic based upon rules defined by the administrator system. Today, many firewall products are still based on UNIX systems, but some are based on Windows NT, while others are taking on an 'appliance' architecture where there is no visible operating system but only an interface for configuring rules and logs.

The factors that influence the selection of a firewall include:

- performance – it is important that the firewall does not become a 'choke point' that significantly slows down traffic between the organization and the Internet
- security – the reputation of the firewall itself is important; it is vital to select a firewall that does not have a history of exploitation by hackers and that does exactly what it is supposed to do

- redundancy – for organizations with heavy reliance on Internet communications, the use of load-balancing and failover firewalls may be justified
- training and expertise – it is important to choose a firewall that is well known so that a selection of training classes, books and technologists is available; there is little value in a firewall that few know how to configure and operate
- price – organizations always seek value according to their specific needs.

## Intrusion detection

Intrusion detection is a passive component in an organization's network infrastructure. While it does not actually *do* anything, it is nonetheless vital to the overall network access control architecture. Intrusion detection refers to the task of passively listening to network traffic, looking for anomalous traffic that could be an indication of misuse or attack. When an intrusion detection system detects suspicious traffic, it sends out an alarm of some type so that someone responsible for network security will be alerted.

There are two primary types of intrusion detection: network-based and host-based. Each is explained here.

- *Network intrusion detection.* A network intrusion device connected to a network listens to all traffic on the network. The nature of network traffic permits the intrusion detection system to observe each communication packet on the network, including the source of the packet, the destination of the packet, the type of communication (e.g. e-mail, World Wide Web, domain name service, file transfer, etc.) as well as the contents of the packet.
- *Host intrusion detection.* Host-based intrusion detection takes the form of a software add-on installed on a particular host. Host-based intrusion detection is designed to detect only the attempts to penetrate the host on which the software is installed.

Organizations will frequently employ both network-based and host-based intrusion detection. While there may be some redundancy between the two, arguably they are also complementary. The integrity of some hosts may be so important as to justify the use of host-based intrusion detection on them. Examples include an organization's Internet web server, application server or e-mail server.

In order to be effective, intrusion detection systems must be updated frequently so that they have a current listing of 'attack signatures' (hackers are among the most innovative people in society today). The intrusion detection system (IDS)

administrator must download a copy of the new attack signature file from the IDS vendor and then upload it to each IDS system in the enterprise.

The location of network-based intrusion detection in the organization depends upon what the organization wishes to monitor. For instance, a network-based intrusion detection system can be placed outside an organization's firewall in order to observe attempts to penetrate the organization's network. But network intrusion can also be placed just inside the firewall in order to measure malicious traffic that has penetrated the firewall. Network intrusion can also be placed well within the organization to detect malicious traffic deep within (and possibly originating within) the organization.

A common complaint about intrusion detection systems is the incidence of 'false positives'. Intrusion detection systems tend to be a little more sensitive than most organizations would like in the long run. A rigorous process of 'tuning out' the false positives is required after acquiring an IDS in order to eliminate the noise. A well-tuned intrusion detection system will have few false positives (but rarely none), lending to its integrity and value.

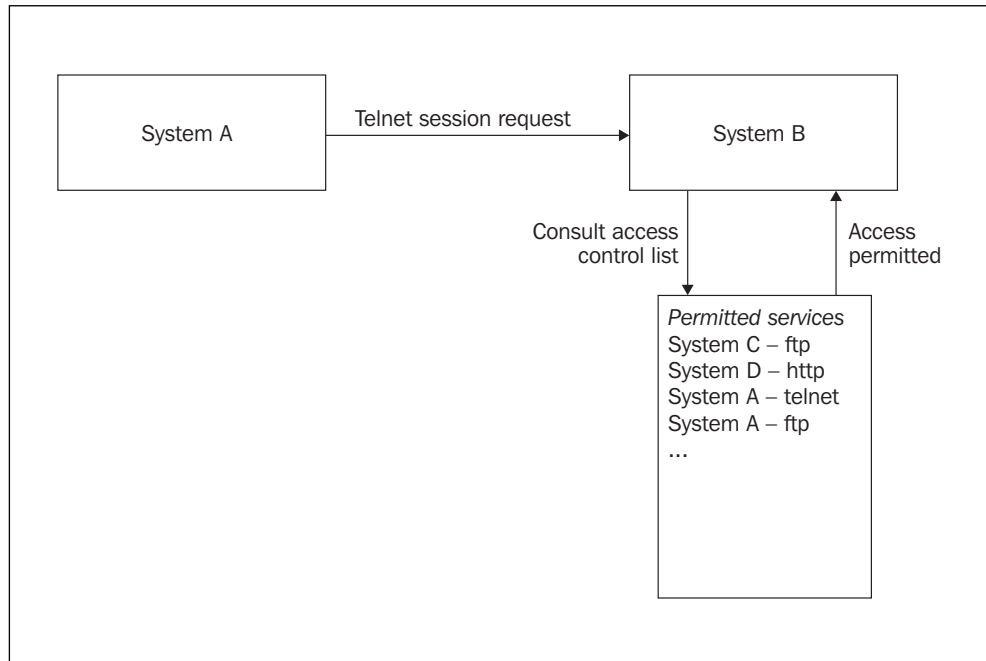
Firewalls and intrusion detection are complementary in two ways. First, both log malicious traffic. Also, intrusion detection systems are useful for detecting events within a network protected by a firewall, whether or not they penetrated the firewall or originated behind it.

Intrusion detection technology is considerably newer than firewalls. It is important to choose a vendor that has leadership and vision in the market so that over time the intrusion detection systems that an organization implements will be able to 'keep up' with innovations.

## Host-based access control

Host-based access control is used to restrict access to a system from other systems. A system with host-based access control can permit or deny other systems' requests to connect, based upon a combination of their network address and the service to which the other system wants to connect. For example, System A wishes to make a telnet connection to System B. System B's access control will check to see whether System A is permitted to connect via that service. If so, the connection will be permitted; if not, the connection is denied. This example is illustrated in Figure 3.3.

Host-based access control is used to limit *inbound* connections only. Generally, a system may attempt to connect to any other system via any network service – it is the responsibility of the destination system to determine whether the connection should be permitted or not.

**Fig. 3.3** Host-based access control

### Personal firewalls

The term ‘personal firewall’ refers to the class of software products that provide firewall functionality on a PC. These products mitigate a genuine risk: increasingly, PCs are connected to the Internet via high-speed connections (e.g. DSL, ISDN, cable modem) from residences and hotels, etc. without the protection of a corporate firewall. The threat of penetration of a non-protected PC is very real, as anyone who has purchased a personal firewall product will attest: the number of scans and attempted intrusions can typically exceed a dozen for each hour that the PC is connected to the Internet.

Most personal firewall products also have an intrusion detection capability that informs the user when scans or attempted break ins occur. With most products the level of alerts is configurable so that users can avoid constant interruptions if desired. Some products also include an enterprise logging capability wherein the client will – in real-time or in batches – forward alarms to a central monitoring system. While it may be useful to gather statistics about the frequency of scans and attempted intrusions, the frequency of these events makes it infeasible for most organizations to take any action when the alarms occur.

Personal firewall products work by inserting themselves in the PC’s network drivers and listening to all inbound (and, for some products, outbound) network traffic. All

types of traffic that match a list are blocked; all other traffic is permitted to pass. These products typically introduce negligible overhead and are barely noticeable from a performance perspective.

## Logging events and taking action

Many organizations that implement firewall logging and intrusion detection assume that they will respond to scans and attempted intrusions with investigative work, countermeasures or reporting to law enforcement. They will quickly discover that the sheer number of logged events exceeds even the largest organization's capability for chasing down each one.

There are only two types of instances where an organization need take action:

- a scan has revealed an actual weakness or bug that can be or has been exploited
- an attempted intrusion was successful.

Either of these types of events would (should) trigger alarms in a network-based intrusion detection system that is *behind* the organization firewall – in other words, a scan or intrusion that successfully penetrated the firewall. An organization that detects malicious activity of this nature should:

- quickly investigate all potentially compromised systems and devices and look for signs of tampering
- take all compromised systems offline in order to protect information from disclosure or tampering
- gather forensic information if criminal activity is suspected
- consider notifying local or national law enforcement authorities
- consider notifying the Computer Emergency Response Team (CERT) at [www.cert.org](http://www.cert.org).

No organization should undertake these activities without first developing a disaster recovery plan. A security event should, in most circumstances, be considered a disaster, since the impact of a security event is similar to a natural or man-made disaster: one or more information systems are offline. In the case of a security event, the organization deliberately takes the system offline if the intruder hasn't done this for them. In any case, steps need to be taken to restore capabilities to normal.

Network-based intrusion detection alarms *outside* the firewall require no action unless it can be determined that the firewall did not block the traffic that the intrusion detection system observed.

## **ENCRYPTION**

Encryption refers to the process of transforming data into a secret code that can be read only by someone (or *something*) that has possession of a secret key. Encryption is used to ensure confidentiality by making it *extremely difficult* for an unauthorized third party to read a block of information. ‘Impossible’ is a term that cryptographers avoid using because a few of the ‘impossible to break’ encryption algorithms have been broken or weakened.

Cryptography is the black art of using mathematical algorithms to encrypt and decrypt information. While it is not necessary to know precisely how encryption works, it is imperative that the reader be familiar with the basic concepts. There are two primary venues for encrypting data: data that is stored on a system and data that is transmitted over a network from one system to another.

### **Encrypting stored data**

Data in storage can be encrypted for the purpose of providing an additional layer of protection from unauthorized disclosure. An intruder who breaks into a system with the intent of obtaining specific data and who is able to penetrate authentication and access control mechanisms will have the daunting (if not impossible) task of decrypting any data he or she may find.

### **Encrypting transmitted data**

With increasing regularity, organizations find themselves sending confidential information over the Internet to external locations, individuals and organizations. Examples include the following:

- *Remote access over public networks.* Increasingly, organizations are using VPN (virtual private network) products and technologies to provide remote access capabilities for their personnel. Also, many organizations with branch offices or retail locations communicate between headquarters and these other locations via the Internet, using tools that automatically encrypt all traffic that traverses the Internet. This is usually more economical than leasing private circuits from telecommunications carriers.
- *Electronic transactions.* Purchasing goods and services, trading financial securities and accessing personal information databases all need to be encrypted in order to protect sensitive information such as credit card numbers, bank and investment account numbers and personal medical records from disclosure to unauthorized third parties. These transactions take place over the public Internet.

- *Business-to-business transactions.* Businesses frequently exchange sensitive information and perform transactions over the Internet. Examples of the kinds of information that are transmitted include e-mail, quotations or proposals, purchase orders, invoices, product specifications, engineering designs, test data, contracts, payroll and personnel information, marketing plans and sales forecasts. In the wrong hands, information such as this can damage the organization.

Someone who is able to monitor and record transmitted data (not a trivial task) faces an extreme challenge when that transmitted information is found to be encrypted.

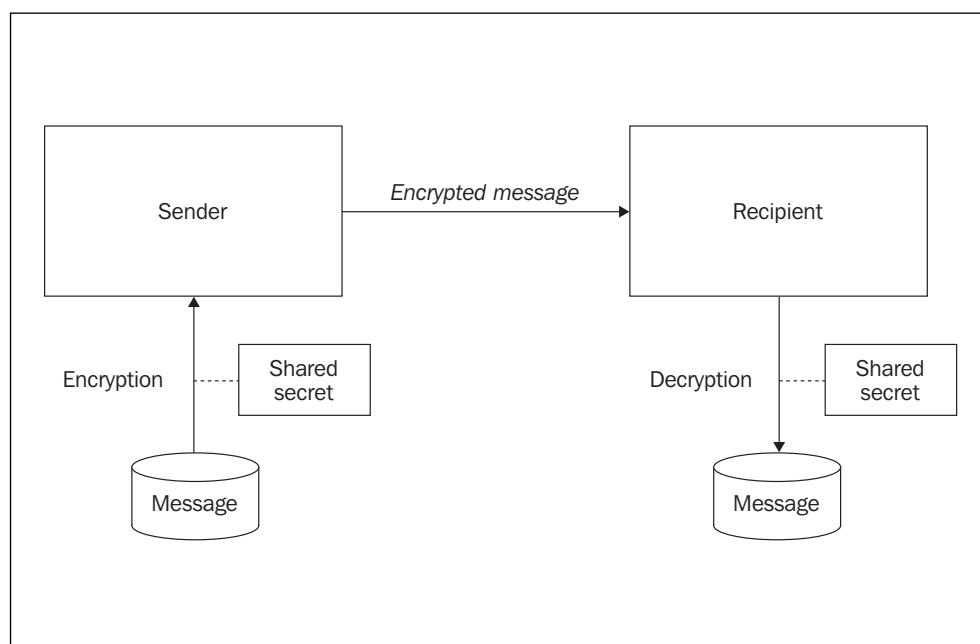
## Encryption technology

There are five types of encryption technology in use today: symmetric encryption, public key encryption, key exchange, digital signature and message digest.

### *Symmetric encryption*

Symmetric encryption refers to encryption wherein both parties must have the same encryption key, called the *shared secret* (see Figure 3.4). For instance, user Bob wishes to encrypt a file using a shared-secret encryption algorithm and send it to user Judy. Bob encrypts his file using a software program and he enters a word or phrase – the secret key. After sending the encrypted file to Judy, Bob also has to tell Judy the secret key so that she can decrypt the file to read or use it.

**Fig. 3.4** Symmetric encryption



The challenging part of shared-secret encryption is that the two parties must share the secret by a means other than the usual communication path. In other words, if Bob and Judy are sending each other shared-secret encrypted files via e-mail, it would be imprudent for them to send the shared secret also via e-mail. If they did, and if an intruder was recording their transmissions, the intruder would intercept the secret key and be able to decrypt all subsequently encrypted messages, thereby defeating the purpose of encryption.

Two parties exchanging data encrypted with a shared secret must transmit the shared key *out of band* – that is, by some other method. For instance, Bob and Judy, who send their encrypted files via e-mail, should exchange keys via telephone, fax or surface mail. Doing so greatly reduces the likelihood that their encrypted messages can be decrypted by a third party, since it is unlikely that the third party will also be able to intercept the secret key exchange.

### **Public key encryption**

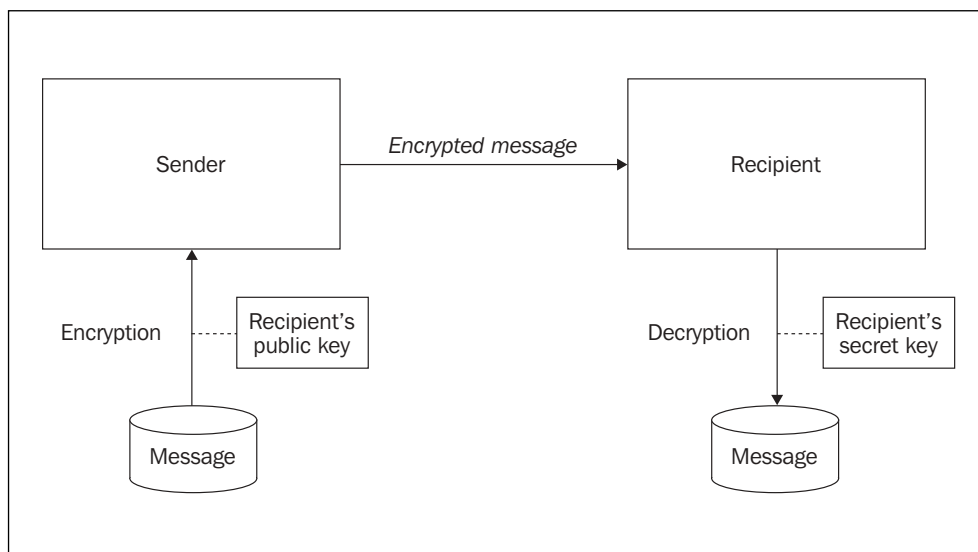
A newer encryption algorithm, called public key encryption, eliminates the problem of how to get the shared secret from one party to another. In public key encryption, each participant creates a *key pair* that consists of a *secret key* and a *public key*. A member of a community who wishes to communicate using public key encryption will send their public key to all the members of the community. For example, user Bob will send his public key to everyone in his group, and everyone else in his group will send their public keys to Bob and to all the other members. In a very large community, this process of key exchange becomes cumbersome, so the practical solution is to store everyone's public keys in a public key infrastructure. This will be discussed later in this chapter.

When Bob wishes to send an encrypted file to Judy, Bob will encrypt the file with Judy's public key and send it to her. Judy will decrypt the file with her secret key. If Judy sends a file back to Bob, she will encrypt it with Bob's public key and Bob will decrypt it with his secret key. Public keys are used to encrypt files and private keys are used to decrypt files (*see* Figure 3.5).

### **Key exchange**

Key exchange is a method whereby two users generate a secret key without having to resort to an out-of-band communications channel. The algorithm behind key exchange is fairly simple and yet two users can calculate and agree on a shared secret even if others are monitoring all their communications.

The most common key exchange algorithm in use is Diffie-Hellman. Developed way back in 1976, it is still in wide use today.

**Fig. 3.5** Public key encryption

### Digital signature

A digital signature is a block of data that is attached to a message or transaction. A digital signature serves several functions:

- It contains and asserts the *identity* of the signer. When the originator signs a document or transaction, he or she must have access to their encryption key and password.
- It guarantees the *integrity* of the message. A message signed with a digital signature is unalterable. Rather, it is more accurate to say that the digital signature will fail to verify if the message has been altered.
- The digital signature *cannot be reused*. A digital signature associated with a document or message cannot be applied to any other document or message, even one originated by the same person.
- The digital signature *cannot be repudiated*. Short of the originator claiming the loss of his or her private key *and* password, a digital signature on a message or document irrefutably ties the originator to the message.

### Message digest

A message digest is a block of text that is created and associated with a message, transaction or computer file. The purpose of a message digest is to create a compact 'fingerprint' of a larger message or file.

The power of this technology is that it is hard to find two messages that will create the same message digest. Thus, a message digest is used to verify the integrity of a message, transaction or file. For instance, it is common for a software developer or packager to create a message digest of a software package archive. This will give the user of the software package the assurance that the entire software package has not been altered.

## Encryption algorithms

There are many encryption algorithms, as they are called in the trade, used in encryption products today. An encryption algorithm is a mathematical formula that is used to encrypt or decrypt data. Some have been developed in and for the public domain, while others are strictly commercial ventures. Some in each category are patented. A discussion of each of these follows.

- *Public domain algorithms.* There are several high-quality encryption algorithms that have been released into the public domain. Companies building products with encryption are free to use and sell these products without having to pay royalties to the developers/owners of the algorithm. The best public domain algorithms are considered to be very high quality since they have been scrutinized by cryptography researchers all over the world.
- *Commercial algorithms.* There are several encryption algorithms that are commercially available; the owners of these algorithms sell or license their use.
- *Proprietary algorithms.* Some commercial products use proprietary encryption algorithms rather than one of the public domain algorithms. Generally these algorithms will be of lower quality than any of the public domain or commercial algorithms because they have not been scrutinized by large numbers of researchers like the public domain algorithms.

Table 3.1 lists information about some well-known encryption algorithms, listed in order by type of encryption technology.

**Table 3.1** Encryption algorithms

<i>Encryption algorithm</i>	<i>Domain</i>	<i>Type</i>	<i>First year of use</i>	<i>Notes/status</i>
DES	Public	Symmetric encryption	1976	Broken in 1989. No longer in widespread use.
Triple DES	Public	Symmetric encryption	1990	Viable and in widespread use.

<i>Encryption algorithm</i>	<i>Domain</i>	<i>Type</i>	<i>First year of use</i>	<i>Notes/status</i>
AES	Public	Symmetric encryption	2001	Brand new, but heavily scrutinized and believed to be secure. Uses Rijndael algorithm.
Rijndael	Public	Symmetric encryption	2001	Chosen as the worldwide Advanced Encryption Standard (AES).
Blowfish	Public	Symmetric encryption	1994	Still considered viable.
RC2	Commercial, not patented	Symmetric encryption	1994	Commercial product developed by RSA. In widespread use.
One-Time Pad	Public	Symmetric encryption	1917	Low tech, but still considered the only perfect encryption algorithm.
SAFER	Public	Symmetric encryption	1994	Not in widespread use – not considered viable.
CAST	Patent pending	Symmetric encryption	1994	Still considered viable.
IDEA	Commercial, patented	Symmetric encryption	1992	Very strong and viable.
SKIPJACK	Patented, but held in secret	Symmetric encryption	1994	Developed for tamperproof hardware (e.g. Clipper chip).
RC4	Commercial, not patented	Symmetric encryption	1987	Commercial product developed by RSA. In widespread use. Anonymously posted to the Internet in 1994.
Diffie-Hellman	Commercial, patented	Key exchange	1976	In widespread use. Patent expired in 1997.
Station-to-Station		Key exchange	1992	Similar to Diffie-Hellman, but not vulnerable to a man-in-the-middle attack. Not widely used.
RSA	Commercial, not patented	Public key encryption, digital signature	1979	Still the most popular public key encryption algorithm.
ELGamal	Not patented	Public key encryption, digital signature	1985	Still considered viable.

<i>Encryption algorithm</i>	<i>Domain</i>	<i>Type</i>	<i>First year of use</i>	<i>Notes/status</i>
DSA	Public	Digital signature	1994	Not widely used.
MD4	Commercial, not patented	Message digest	1990	Moderate viability.
MD5	Commercial, not patented	Message digest	1992	Improved version of MD4. Moderate viability.
MD2	Commercial, not patented	Message digest	1992	Very strong, but not very fast.
SHA	Public	Message digest	1993	Very strong and viable.

## Encryption applications

Encryption is found in a wide variety of applications. Four general categories are discussed here: virtual private networks (VPN), file, directory and database encryption, digital signatures and certificates.

### *Virtual private network (VPN)*

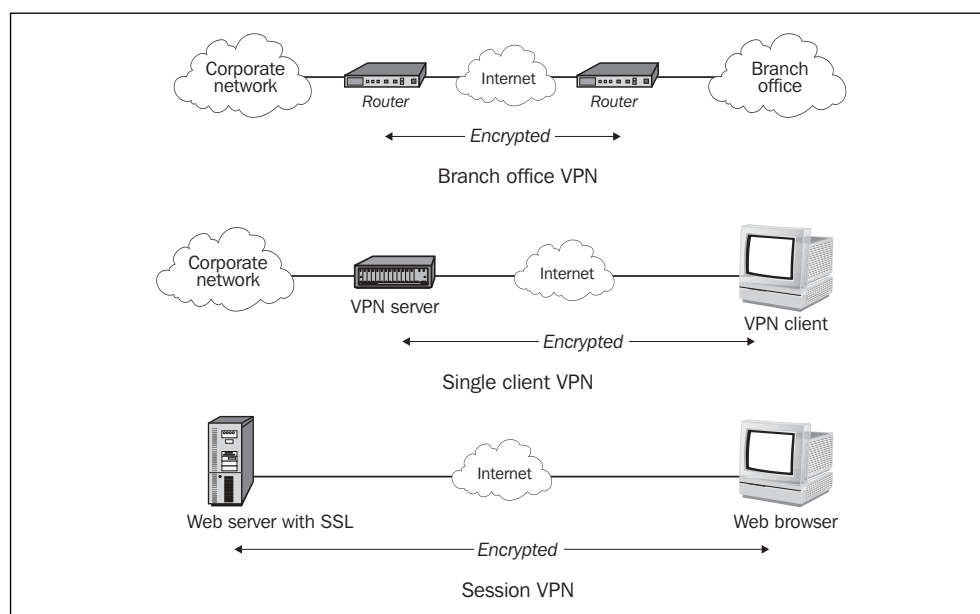
Virtual private network technology provides a means for automatically encrypting all network traffic between two endpoints. Application programs and tools that communicate over the network function normally and are completely unaware of the presence of the VPN.

There are three primary types of VPN technologies: ‘session’ VPN, ‘single client’ VPN and ‘branch office’ VPN (see Figure 3.6). Each is explained here.

- *Session VPN* refers to a single client-to-server secure communications session. The session is secured through encryption and may also be secured through authentication. The most well-known example of session VPN is SSL, or Secure Socket Layer. This is the technology that performs two functions: first, the communications session between a client (usually a web browser) and server is encrypted; second, the client examines the server’s certificate to determine whether it is genuine, thereby assuring the user that the server is authentic and not an imposter. Another well-known example of session encryption is SSH, or Secure Shell. SSH is used by system and network administrators to establish encrypted session to servers and network elements. SSH is used as a replacement for the telnet, rsh and FTP tools, all of which communicate over networks without using encryption.

- *Single client VPN* consists of a client system (usually a PC) with VPN client software installed. This software inserts itself into the PC's network drivers to encrypt automatically all network traffic. The VPN software establishes a 'tunnel' to a VPN server; the user must authenticate to the VPN server by providing a password or token. Once the tunnel from the client PC and the VPN server is established, all communications between the PC and the organization's network are encrypted. The single most popular reason to use VPN technology is to shift the burden of dial-in remote access from the organization to an Internet Service Provider (ISP). The person's traffic may safely pass over the public Internet via a dial-up connection to the ISP and the path over the Internet from the ISP to the organization.
- *Branch office VPN* is a network-to-network encrypted tunnel that is established using existing network equipment. A branch office will usually have a network device called a 'router', which is used to send packets to their appropriate far-off destination over the Internet and/or lease lines. The 'headquarters' network will likewise have a router. A network engineer will configure the router on each end to encrypt automatically all information that passes from the router in the branch office to the router at headquarters, and vice versa. The primary advantage of the branch office VPN is that expensive leased lines, possibly spanning hundreds or thousands of kilometres, can be avoided. Instead, each participant in a branch office VPN can simply just connect to the Internet via the nearest ISP, which in most locales will be just a few kilometers away and, hence, far less expensive.

**Fig. 3.6** VPN architectures



### ***File, directory and database encryption***

Many situations call for the encryption of individual files, entire directories or volumes, and databases stored on computers. Generally these situations arise from the risk of physical or network access to information that compels its owner to use additional means to protect it. Some examples include:

- data stored on laptop computers and other portable equipment
- data stored on systems that are accessed by many people, for instance a department or organization-wide file server
- data that needs to be transmitted to other systems where there are no other facilities in place to ensure that the data cannot be viewed by eavesdroppers while in transit.

In each of these cases, software tools can be used to provide encryption using shared secrets or public key encryption. Generally speaking, it is only necessary for all the people accessing these encrypted files to exchange public keys or shared secrets and to use the same encryption and decryption tools.

File encryption tools generally require that the user deliberately encrypts each file that needs to be encrypted. On the other hand, the directory encryption tools are more automated: when a file is written to an encrypted directory, it is automatically encrypted; when a file is retrieved from an encrypted directory, it is automatically decrypted. The method used to protect the contents of these encrypted directories is usually a password challenge that is required to unlock it. Then, for a given period of time, the user can continue reading from and writing to the encrypted directory before having to re-authenticate. After a predetermined amount of idle time, the directory will lock itself and the user must be challenged to be able to access it again.

The value of directory and volume encryption is two-fold: first, since all files in the directory or volume are automatically encrypted, the user does not have to remember to encrypt them manually. Second, if an unauthorized person accesses the computer, they will not be able to read the contents of the encrypted directory or volume.

### ***Digital signatures***

The term *digital signature* refers to a method for ‘signing’ a document or transaction as a way of making it genuine. While the digital signatures method is different from encryption, digital signatures are associated with encryption because they have the same technical roots: both are applications of complex mathematical formulas – in some cases even the same formulas. In some products, encryption and digital signatures are performed using the same keys.

A digital signature is created by a program that reads the contents of a file, message or transaction and mathematically combines the file, message or transaction with the

encryption key of the person signing it. The digital signature is then appended to the original data and accompanies it so that other people can verify the genuineness of the original data. When the digital signature is verified, the recipient knows:

- that the contents of the original data have not been altered
- that the claimed originator of the original data is, in fact, the originator.

Like encryption, the choice of algorithm for digital signatures is important so that the organization can depend upon the strength and integrity of the digital signatures it uses. If an organization uses digital signatures in transactions it performs with other parties, it needs to make sure that the other parties are able to create and/or verify digital signatures. The software and algorithms that various parties use need to be compatible.

### **Certificates**

A certificate is a block of information that is associated with a person, an application or server. The certificate contains the following information:

- the person's (or application's) public encryption key
- information about the person (or application).

The above is signed by a certificate authority's private encryption key.

Certificates are most often encountered by individuals who visit websites that utilize web-server-to-web-browser encryption. But there is more than encryption going on: the site that the person is visiting also asserts its certificate, including an electronic signature from a certificate authority (CA). The web browser verifies the server's certificate by checking it against the copy of the CA's certificate stored in the browser. The browser performs a digital signature verification function, and it also compares the domain name (e.g. `www.ft.com`) with the current URL. If the comparisons match, the web browser is satisfied that the server's identity is genuine. If the comparisons do not match, the browser displays a warning message informing the user that it has not been able to verify the identity of the website.

Server certificates have an expiration date. This forces the site periodically to reassert its identity with the CA. Web browsers visiting a site with an expired certificate will display a warning message, giving the person an opportunity to decide whether to proceed.

In addition to server certificates used to verify the identity of the server, a certificate can be issued to an individual. This certificate is known as a *client-side certificate* and is used to verify the identity of the person visiting the website. Like server certificates, client certificates expire, forcing the person to reassert his or her identity to the CA.

Client-side certificates are not in wide use today, primarily due to the costs involved in issuing certificates – organizations that want to issue client certificates

must pay a hefty sum to one of several CA software vendors. It is even more costly to require that employees obtain their own certificates – the organization will be paying the highest unit price. For these reasons, nearly all websites use a somewhat less secure method of identifying and authenticating users, typically with a userid and password.

### ***Public key infrastructure (PKI)***

An organization using client-side certificates or public keys may find that manual distribution methods are not scaling well. It can centralize the storage of its certificates and public keys in one or more servers and reconfigure client systems to reference the server(s) instead of local storage when they need to verify or retrieve certificates and public keys. The environment containing one or more central servers, as well as client systems, is known as a public key infrastructure (PKI). PKI servers store all public encryption keys and certificates, thereby relieving individuals from having to distribute keys and certificates to everyone in the organization.

An organization that intends to build a PKI must form a relationship with one of several CAs. The CA's role is to guarantee that the holder of the certificate is who they assert themselves to be. This is based upon a rigorous process used to issue certificates to individuals within the organization. This process assures the greatest possible reliability in determining that people are who they say they are. The CA is the entity that issues certificates to individuals in the organization. Certificates can be issued directly from and by the CA, or indirectly: a certificate administrator in the organization can issue certificates on behalf of the CA.

Another important function performed in a PKI is *certificate revocation*. For instance, when an employee leaves the organization, the employee's certificate is added to the certificate revocation list (CRL), which is propagated throughout the PKI. Applications that verify the identities of individuals in an organization query the PKI for the person's certificate and also search the CRL to see whether the certificate was recently revoked. A common issue regarding CRLs is the question of how frequently the CRL should be updated in order to avoid situations where an individual whose certificate has been revoked can attempt to use it before his or her certificate appears in the CRL. Updating the CRL too frequently will have network and PKI performance implications, while too seldom will increase the likelihood that individuals with revoked certificates will be able to continue using them.

## **Encryption issues**

### ***Encryption algorithm strength***

An organization shopping for applications or tools that use encryption technology needs to know which encryption algorithms are being used and how they are being

used. While it is generally unnecessary for an organization to have a cryptologist or mathematician on hand when making the decision, it is important that a security expert weighs in on the strength, merit and viability of the algorithms that vendors use in their products. The expert also needs to understand the business functions involved and the associated qualitative and quantitative risks.

For this reason, this book cannot prescribe specific encryption algorithms for any given setting. This is because every decision to use encryption is a risk decision that takes into account the value of the information being encrypted, the business impact if the subject information were to be disclosed, and the cost and effort required to perform the encryption. Many situations fit neatly into areas where off-the-shelf products provide adequate solutions, but there are also circumstances where the canned solutions are either inadequate or too expensive.

### ***Key escrow and recovery***

Suppose that an employee follows the rules and encrypts important files on his or her PC. But unexpectedly the employee is fired, and now it is someone else's job to retrieve some of those files so that the task can continue. A colleague discovers that the files they need are encrypted with the former employee's private key. Presuming the former employee will not be co-operative, how can the files be decrypted and recovered?

This kind of situation has led to the key recovery capability. When encrypting files (or entire directories), some encryption tools will encrypt with not only the user's private key but also a special administrative key. This permits an organization to recover materials encrypted by its employees by decrypting them with the administrative key.

Termination is by no means the only situation warranting key escrow and recovery. Sometimes people forget their passwords, or they can lose their private key (for instance, if the only copy of the key was on a PC that had a hard disk failure). And unpleasant as this may be, employees can die unexpectedly. All of these situations justify the use of key escrow and recovery.

There is another side to key escrow and recovery. Some governments may assert 'master key' or 'key escrow' requirements to support law enforcement purposes. This would permit a law enforcement agency to recover any unencrypted data for investigative purposes. A notable example is the 'Clipper chip' initiative in the USA in the 1990s, in which the master key for each encryption chip would be placed in escrow and could be recovered by law enforcement agencies for investigative purposes.

A key escrow and recovery capability on this scale introduces significant risks, as there would be a high probability of abuse in the form of even a small number of law enforcement or government officials inappropriately using their key escrow capabilities to eavesdrop on private communications.

## Encryption summary

Encryption increases costs through the cost and use of additional tools and through additional computer resources required to perform encryption. Using encryption everywhere will present the best economy of scale, but not the best use of funds and effort. Analysis that identifies specific encryption needs will make the best use of the organization's time and resources.

Encryption also reduces risk from accidental or deliberate disclosure of proprietary or private information by making it infeasible for an unauthorized party to read and use such information.

There are a wide variety of encryption mechanisms and technologies available that make it possible for an organization to purchase and use encryption where it is needed. Many of the tools available make encryption of data automatic, thereby taking decisions and additional work out of people's hands in many situations.

## NON-REPUDIATION

Non-repudiation refers to the ability to *prove the authenticity* of an object of data that could be a file, a transaction or a message. The object's authenticity is proven in two ways: first, that the data object was created by (or approved by, whatever the case may be) the stated sender, and second, that the content of the data has not been altered.

There are no software tools or products that, by themselves, are used to create non-repudiation. Rather, the term refers to a concept or requirement that some organizations assert is necessary in certain situations or settings, for instance in electronic commerce or online banking. Generally speaking, non-repudiation is achieved by implementing digital signatures, which provide the means for providing non-repudiation through their originator and non-alteration guarantees.

### Why could non-repudiation be important?

The value of non-repudiation is most often defined by an organization's legal requirements. The ultimate test of a transaction's validity is the courtroom, and as such a transaction may require a strong chain of evidence. The business and financial value of transactions and the need to be able to prove their authenticity will determine the need for non-repudiation.

## INTEGRITY

Information integrity refers to the protection of the accuracy of information and the correct operation of systems, operating systems and networks. This assures the

individuals using or providing information that the information is genuine and that it has not been altered by unauthorized people.

This section explores integrity from the inside out, beginning with the data itself and including systems, networks, configuration management and change management. There is also a lengthy discussion of viruses, which remain a significant threat to data and system integrity.

## Data integrity

Data integrity is achieved as a result of other security mechanisms and processes:

- authentication – knowing who is accessing or modifying data
- authorization – controlling who can view or change which data and perform which functions
- access control – mechanisms controlling how data can be changed and from where
- audit – maintaining a record of all changes
- configuration management – recordkeeping for all system changes
- change control – the organization’s approval process for all system changes.

Data is at the core of the layered information model. It is tightly controlled through all these mechanisms and processes to ensure that the organization’s data is as safe and accurate as possible.

## System integrity

There are no simple operating systems. The UNIX, NT, Windows and mainframe operating systems that power most of today’s computers have a plethora of dials, knobs, switches, shortcuts and undocumented features. Strict adherence to the best system-level configuration management system can still leave stones unturned. Legacy tools on these operating systems permit knowledgeable administrators to make changes by ‘going around’ the configuration tools, which sometimes allows changes to be made without audit entries.

All of this complexity addresses the issues related to those people who are *allowed and qualified* to administer these systems. What about the people in the organization who are *not* adequately trained, overstep their bounds, cut corners or decide to wreak havoc, and those in the organization who are not authorized to access systems at all?

Combine operating system complexity, honest mistakes and the nearly continuous discovery of new ‘holes’ in operating system software, and it should be self-evident that there is a real risk of ‘things going wrong’. The complexity of

operating systems can make it extremely difficult even to know when a change has taken place. For this reason it is suggested that regular audits of the operating system be performed.

### ***Auditing the operating system***

Perhaps the most effective method for detecting changes (those that are proper and authorized, and those that are not) is to perform regular comprehensive audits of the operating system. Some operating systems have auditing tools (Sun's Solaris has ASET) and there are third-party tools available such as Tripwire. These tools detect the most minute changes; if even a single character is changed in a configuration file or in an executable program, these tools will detect and report the change.

### ***Deactivating unnecessary services***

One of the most common errors committed by system administrators lies in the failure to deactivate unnecessary functions on a system. This is best illustrated by an example. Suppose that an application is installed on an NT server and that the NT server also has IIS (Internet Information Server, the Microsoft web server software), SQL Server and Index Server installed and running. In this example, neither IIS, SQL Server nor Index Server is necessary for the application to run. Nonetheless, any security hole in the IIS, SQL Server or Index Server products will make this application server vulnerable to attack.

The best practice is to disable and remove all unnecessary services and packages from a server. If the service or software package is not present on the server, then any security bulletin describing a new security weakness is irrelevant on this server. If the service is not on the server, it cannot be attacked.

It is important, then, that all systems – particularly servers – be stripped of all unnecessary services. They only add to each system's complexity, overhead and – most important of all – vulnerability.

### **Network integrity**

An organization's network fabric itself – its routers, switches, hubs, etc. (hereafter referred to as *network elements*) – must be in correct working order and be defended against attacks. Weaknesses in the network can expose the organization to events that can bring it to a virtual standstill.

Just as with servers, network elements must utilize the best access management available. Only authorized users should be able to access a network element and all changes should be recorded. Network elements need to be configured to prevent malicious network traffic from entering and propagating throughout the organization. While network elements generally cannot detect or block viruses and worms, there are plenty of 'bad things' that are preventable, such as an SYN

flood (a type of denial of service attack that drains a server's resources until it hangs or crashes), ping of death (a type of attack that causes a server to crash), vulnerability scans, source-routed packets and routing table poison. Some of these and other threats are discussed at length in Chapter 2. It is unnecessary to delve into the details of these types of attacks – the organization's network engineers and administrators must be aware of them and configure the network elements to resist these and others to the greatest extent possible.

## Configuration management

Configuration management is central to the theme of integrity. It is the accurate *recordkeeping* of every configuration change made to a system's hardware, operating system, application or content. Configuration management allows a system's owner to track what changes took place at what time.

With configuration management, a system can be reconstructed – using the configuration management records as a guide – to any desired point in time. Configuration management records provide information that permits someone to rebuild a server's hardware, operating system, application or content in the event of a catastrophic hardware failure, a site disaster or a security incident. This capability is especially important if an organization encounters problems with a server. Frequently an organization will recover a server from the most recent backup, but sometimes the situation calls for a recovery from an earlier backup (for instance, software or content corruption discovered after the fact would compel an organization to rebuild from a previous, known good point in time). With accurate configuration management records, it could revert the system to an earlier, reliable point in time.

## Change management

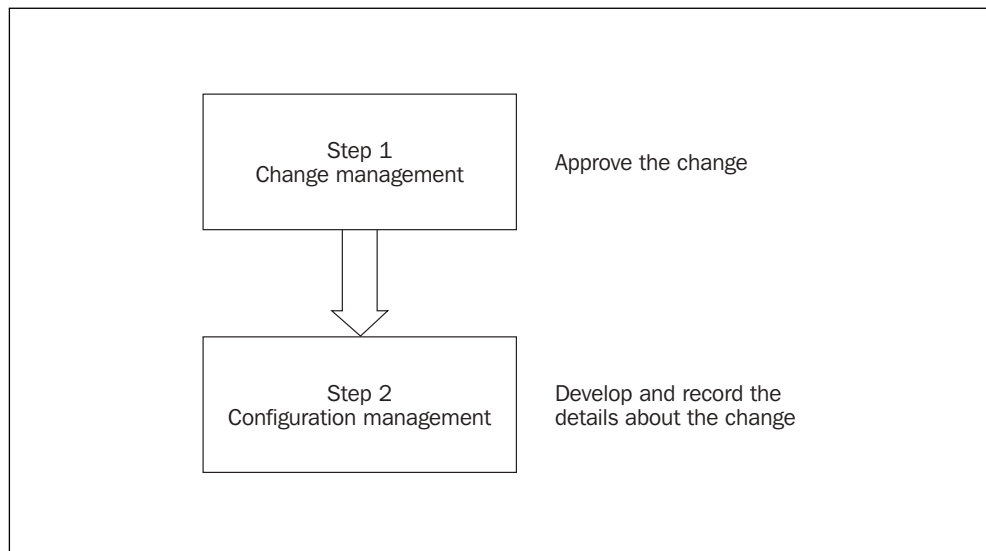
Change management refers to the *process* of making changes to a system. It is not enough for a system administrator to make arbitrary changes to a system and record them in configuration management logs. Proposed changes should go through an *approval process* to ensure that all stakeholders agree to the proposed change (see Figure 3.7).

A change to a system begins with a *change request*, which should contain:

- a description of the change
- why the change will take place
- when the change will take place
- who will be affected (and how) during and after the change
- a communication plan to include any required announcements or notifications

- who will perform the change
- who will verify the change
- a description of the backout plan should something go wrong with the change
- any documentation changes required to reflect the changed system.

**Fig. 3.7** Change and configuration management



A *review board* consisting of organization stakeholders should examine, understand and approve the change request. After the change takes place, a review of the change should be performed to discuss the change and to capture any lessons learned to be incorporated into future changes on the platform. Part of the reason for a change management process is to prevent an individual, operating alone, from performing unauthorized changes to the organization's information infrastructure.

## Malware

Perhaps the greatest threat to the integrity of systems and data is the continuous assault by viruses, worms and Trojan horses. Together these are sometimes called *malware*, a contraction of *malicious software*. Malware spreads in many ways. The best-known methods are e-mail, document macros, and web clients and servers. Some forms spread when individuals open e-mail attachments or view certain websites, others propagate automatically without any human intervention.

### *Defending against viruses*

The traditional defence against viruses is anti-virus software. This made its debut on desktop systems many years ago and has changed little since then. It works by

inserting itself into the mechanisms used to read and store files and execute programs. Whenever a file is being read or stored, or when a program is being executed, the anti-virus software examines its extensive list of virus signatures to see whether there is a match (between a signature in the anti-virus signature file and the contents of the file being read or stored) and, if so, the anti-virus software intervenes and prevents the operation from completing. It will alert the user that a virus has been found.

The anti-virus mechanism described here works well, provided that the list of virus signatures is kept up to date and that the anti-virus software is actually running and configured correctly. While these may seem like trivialities, in a large organization with thousands or tens of thousands of workstations, this can be a daunting task. A large organization often has the most to lose when a virus runs rampant within it, and this can happen if even only a handful of workstations are unprotected or if their anti-virus signature files are out of date.

Larger organizations need to choose a desktop anti-virus product that provides automatic 'push' updates. This refers to a mechanism where a central server, upon receiving a new anti-virus signature file, will 'push' it out to all workstations in the organization. The anti-virus software on each workstation will be configured to accept automatically the new signature file, thereby providing protection from the latest viruses and worms. This can occur without the user even being aware of the virus update taking place.

Another desirable characteristic of workstation anti-virus software is the enterprise alert capability. What this means is that any workstation, upon detecting a virus, will send an alert to a central console, thereby providing warning that a virus has been detected within the organization. But in order to work, a mechanism for getting new signature files to all workstations needs to be in place. If a new virus is released and it reaches a workstation without the latest anti-virus signatures, the virus will infect that workstation and will attempt to propagate to other systems in the organization. The workstation's anti-virus program will not send an alarm because without the latest signature file, it will be unaware that an infection has taken place.

Anti-virus software should also be installed on all file servers, print servers, application servers and e-mail servers. On e-mail servers in particular, specialized anti-virus software will examine each e-mail message in transit and strip infected files from any message found with a virus in it. One may wonder why these additional layers of anti-virus protection are necessary. Think for a moment about the hundreds or thousands of workstations in use: are fully 100 per cent of them running a current anti-virus program? E-mail anti-virus software will also serve as an organization's first line of defence by catching viruses coming in from outside the organization.

### ***Keep patches current***

The single most effective defence against viruses and other malware is to keep operating system and application patches up to date. The infamous NIMDA and Code Red Internet worms in 2001 exploited security weaknesses for which patches had been available for nearly a year. Only those organizations that were diligent with installing patches (which, unfortunately, were in the minority) were spared the full wrath of these two assaults.

To keep operating system and application patches current, it is necessary for one or more people in the organization to subscribe to a variety of security alert e-mail distributions (the best known are CERT, CIAC and Bugtraq) that provide useful information about the latest known security holes. Then, system administrators need to read the detailed information available about a particular security hole to see which of the organization's systems might be vulnerable. Next, the system administrator will download a patch and install it on a test system. Testing is an essential step, to ensure that the patch does not inadvertently break or change some other function of the component being patched.

Patching a server generally involves some downtime for that server, usually just several minutes while the system is rebooted. However, the frequency of patches released on a given platform can significantly contribute to its TCO (total cost of ownership). Availability statistics will also reflect the downtime required to install patches and reboot systems.

Many organizations provide financial incentives for their operations groups to maximize availability of their servers. This incentive, however, is in direct conflict with keeping systems as secure as possible, since this involves brief periods of downtime for patch installation. Organizations need to re-think their uptime-centric incentives to provide balance between uptime and security. The lessons of NIMDA and Code Red illustrate that saving minutes by avoiding or deferring security patch installations can cost hours or days in downtime later on.

An organization contemplating a new application platform should do some research or consult any of the well-known IT intelligence and research companies such as Gartner, Meta, CIO, Giga or Forrester for guidance on this topic.

## **AUDIT**

Events at every level, from operating system through to application software, need to be recorded and logged. An audit trail is the independent record of events that can be used to reconstruct a series of events or to provide valuable information in a root-cause analysis or security investigation.

Modern operating systems contain a rich audit-logging capability, although for most organizations the most extensive audit logging will provide too much detail. Applications vary widely in their audit-logging capabilities, from none to extensive. This characteristic of any application should rank high on the list of selection criteria.

### **Is anybody watching?**

While it is important for systems and applications to log events, it is all in vain if no one ever reviews the logs to look for suspicious events. Logs can contain important clues to people trying to break into networks, systems and applications, people who are snooping around in an organization's customer data (perhaps recording information to sell to identity theft rings) and people who are attempting to perform functions beyond their ability (for instance, attempting to create phony purchase or remittance orders).

### **Reduction, rollup, and correlation**

No organization wants to pay people to read tens of megabytes of log data every day, and equally unsavoury is the prospect of being that person. This is perhaps the greatest example of information overload and noise-to-signal ratio. The events that one would be looking for are like needles in a haystack, and equally hard to find.

Fortunately, there are tools available that can consume the avalanche of log data and parse it for significant events, as well as create and store useful information. Knowing the frequency and location of events is generally far more useful than the lists of events themselves.

Other tools can correlate events that occur on multiple systems or applications that by themselves appear insignificant but together give new meaning. An example of this is a single failed login attempt on one system, versus scores of failed login attempts on dozens or hundreds of systems. One failed login attempt is hardly significant, but dozens or hundreds could indicate someone trying to break into systems. The same can be said about accesses to customer data. A lone query by a customer service representative into a customer record is insignificant, but high numbers of queries could indicate that the representative is harvesting information and selling it to outsiders.

### **Investigation support**

Without adequate audit trails, an organization will have a difficult time piecing together complex events. Lacking a strong chain of evidence, an organization cannot consider bringing civil or criminal charges against a wrongdoer.

## AVAILABILITY

Availability refers to data and functions being accessible when needed. While availability is not *solely* a security issue (for instance, availability and performance are closely related), there are several scenarios that associate availability and security:

- an intruder could make changes to the system in order to prevent legitimate access to the information
- an intruder could launch a denial of service attack on a system in order to block legitimate access
- an intruder could sabotage communications, power generation or heating and cooling equipment, effectively taking information offline by disabling the underlying computing equipment.

While the term ‘intruder’ most often conjures up an image of a masked criminal prowling the space (or cyberspace) around the organization, much of the time the intruder is someone who is legitimately in or on the premises, because he or she is an employee. This fact cannot be overlooked, as most security incidents are related not to unknown external strangers but to the organization’s employees and contractors.

Unlike authentication, authorization or audit, there are no availability *mechanisms*. Rather, availability is achieved through a number of measures, including:

- configuration of systems and firewalls to resist denial of service attacks
- intrusion detection systems to alert the organization of a denial of service attack and other anomalies
- physical controls to protect data centres, power and communications facilities
- a distributed, redundant or fault-tolerant architecture that provides the greatest likelihood that information and functions will be available in the event that an attack or disaster cripples a system or an entire data centre.

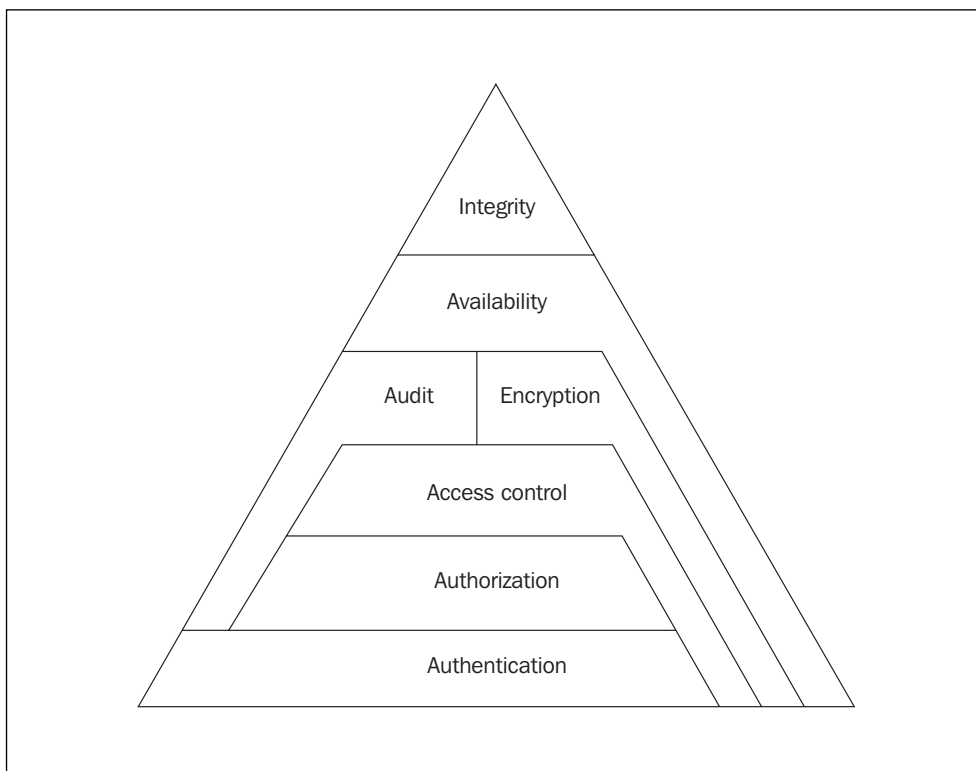
In each of these examples, the organization is ‘hardening’ the environment that surrounds the information systems in order to resist actions and events that make them inaccessible to legitimate users.

## SECURITY MECHANISMS WORK TOGETHER

The security mechanisms described in this chapter do not work entirely on their own. Rather, some of the mechanisms are dependent upon the existence and operation of others. Some examples of this interdependency include the following (*see* Figure 3.8):

- any authorization mechanism is heavily dependent upon its corresponding *authentication* mechanism: before a user can be granted access to specific data or a particular function, it is imperative that the system knows to whom this access is being granted
- audit entries also require that the system properly authenticates its users: audit trail entries contain ‘who’ (in addition to ‘what’, ‘when’ and ‘where’), so the system had better have a pretty good idea of who the ‘who’ really is
- access control mechanisms, in part, depend upon authorization to carry out authorization rules
- integrity depends upon all of the other security mechanisms.

**Fig. 3.8** Interdependence of security mechanisms



## SUMMARY

The *identification* and *authentication* of people is the first line of defence protecting an organization’s information assets and infrastructure. The strength of authentication used determines its reliability: one-factor authentication (usually

just a userid and password) is one of the weakest forms of authentication, but is still the most frequently used. It is important, then, that organizations pay close attention to password quality, in order to thwart would-be intruders. Two-factor authentication (e.g. tokens, smart cards) provides far stronger verification, as does biometrics. There are several technical standards that an organization can use to implement an organization-wide authentication service.

*Authorization* provides the means for controlling which individuals and roles may access which data and perform which functions. Authorization provides segregation of duties control that is necessary for many organization functions. There are no established standards for authorization – organizations that need to implement authorization must do so via proprietary mechanisms, or separately within each enterprise application.

*Access control* refers to mechanisms used to limit access to networks and systems. Network-based access control is used to control traffic between networks, most notably between the Internet and an organization's internal network. Host-based access control is used to limit traffic to an individual system. Intrusion detection devices are used to identify unauthorized traffic. The personal firewall refers to a class of software products used to protect desktop computers that are connected to the Internet.

*Encryption* is used to scramble data so that only the intended recipient(s) can read it. There are a variety of encryption algorithms available from which the developers of encryption products can choose. Encryption algorithms are also used to create digital signatures, which are used to verify the identity of the originator of data, as well as to determine whether the data has been altered in transit.

A *digital certificate* is a block of data that is issued to a system or to an individual. A certificate contains the individual's public encryption key as well as identification information. World Wide Web servers utilizing SSL (secure socket layer) encryption use digital certificates to guarantee their identity to any individual connecting to such a server. An organization using client certificates (a certificate issued to each person in the organization) must build a PKI (public key infrastructure) to manage the certificates and their revocation.

*Non-repudiation* refers to the ability to prove the authenticity of a transaction – irrefutably associating it with its originator and guaranteeing that it has not been altered. *Digital signatures* are most often used to assure non-repudiation.

The *integrity* of an organization's information and information infrastructure (that is, its computers, software and networks) is critical to its viability. Controls must be in place to prevent and detect unauthorized changes to computer and network hardware, operating system software and configurations, application software and information. Anti-virus mechanisms are counted among the methods used to protect the integrity of an organization's desktop systems, e-mail servers, web servers and other systems.

*Audit trails* provide the means for reconstructing events, and they aid investigators who are piecing together complex events. For their auditing capabilities to be effective, organizations need to consider audit log aggregation and correlation tools.

*Availability* refers to the ability of a system to keep running and be available under almost any conditions, even during a hostile attack. Systems and networks can be designed and configured to resist attacks. This characteristic shares common ground with systems architecture and network-level access control.



# Security policies and requirements – defining the standard of architecture and behaviour

- Introduction 75
- What are information security policies? 76
- Who writes security policies? 76
- Audience 77
- Policy development 78
- Awareness 83
- Enforcement and effectiveness 85
- Summary 88



## INTRODUCTION

Security policies define the limits of acceptable behaviour on the part of people and the operating characteristics on the part of hardware and software. There are multitudes of best practices – for instance, the concepts of separation of responsibility or password quality. However, each organization, with its unique mission and ways of doing business, will have unique practices and needs that call for specific policies and requirements.

### What is important

Before one can embark on a quest to develop security policies, one must first ask this question: what is important in the organization? What things need to be protected?

- *Information assets.* For organizations that are based upon or reliant upon information, that information needs to be protected. For instance, its access should be restricted to those persons who have a need to know. Meticulous records must be maintained that keep track of any changes made to the information. If the information is transmitted from one place to another, steps must be taken to ensure that no eavesdroppers along the way can read and intercept it. Information must be kept out of the reach of intruders and others who would steal or sabotage it.
- *Information infrastructure.* If the organization's information assets are important, so is the information infrastructure: computer systems, operating system, databases, application software and networks. The integrity of the organization's information relies heavily upon the integrity of the infrastructure that contains it.
- *Physical facilities.* The preceding items require building space with the proper utilities (power, water, heating or cooling) and physical protection (access by authorized organization personnel).
- *Information availability.* Information must be made available to those who want it (and have the authorization to read and/or modify it), when they want it. Steps need to be taken to ensure that intruders and others are not able to alter or sabotage the information infrastructure – or the information itself – to make it unavailable when needed.
- *People.* No organization can operate for very long without people. This refers not just to the people who manage the business and its information assets but also to the information they need, so that they know how to manage the organization's information and operate its systems and networks.

## **WHAT ARE INFORMATION SECURITY POLICIES?**

Security policies are formal statements that describe how corporate assets should be protected. Policies describe the limits of behaviour for the organization's personnel and the limits of functionality for its systems. A complete information security policy includes the following:

- responsibilities of the organization's personnel
- description of information assets to be protected, including a classification model that specifies levels of protection required
- description of user accounts: how they are obtained, used and maintained
- description of users' authorization to access information and perform functions
- description of access control functions that protect information assets
- measures to be taken to ensure the availability of information
- measures to be taken to ensure confidentiality of information
- measures to be taken to ensure the integrity of information
- description of audit events and audit logging
- description of security assessments and how they are to be performed
- incident response and incident management
- measures taken to protect physical information assets and physical access to those assets.

The role of security policies is to describe 'what to do' to protect information assets. Security requirements, discussed later in this chapter, describe 'how' to carry out security policies.

## **WHO WRITES SECURITY POLICIES?**

A business security or information security group usually writes security policies. In a small organization, policies may be written by an individual security analyst responsible for policy.

Critical to the success and effectiveness of the security policy is a formal written statement from a high-ranking official in the organization stating the critical nature of security policy and the need to abide by and enforce it. This official might be the CEO, CIO, CSO or CISO. This statement, at a minimum, must include:

- a declaration of the existence of a formal security policy
- the department or individual responsible for writing and approving security policy

- the requirement that all personnel follow the security policy
- the consequences for failing to follow security policy.

Without this high-level support, security policy will only be as effective as the individuals who write it and the extent to which they can enforce it on their own. Employees will not have incentives to follow policies if there is no support from senior management.

## AUDIENCE

Up to this point, this book has implied that security policies are written for the organization's employees and contractors. However, the intended audience is generally much wider and should include several more entities:

- *Vendors.* Any outside company that supplies the enterprise with hardware products, software products and services must be aware of and adhere to your policies and requirements. For instance, if the organization has developed an enterprise-wide LDAP-based authentication function, any prospective vendor from whom an application would be purchased needs to know that its application must be able to externalize authentication through LDAP.
- *Business trading partners.* Because they are supplying or purchasing your goods and/or services, organizations participating in the enterprise extranet must follow your security policies and requirements in order to ensure that your information is secure. An organization's data is only as secure as its weakest link, and frequently the weak link is the extranet partner. This will be explored more fully in Chapter 6.
- *Developers and integrators.* Those who develop and integrate software and applications used in the enterprise must be fully attentive to all downstream aspects of the organization's security policies and requirements. Not only must developers and integrators build applications that are directly compliant with security policies (for instance, by conforming to password quality, audit trail and authorization policies), but the *users* of those applications must be able to comply with security policies as they use the application. Further, those who operate the applications (e.g. computer operators, system administrators and database administrators) must be able to comply with security policies as they perform their care-and-feeding duties.

These audiences all need to know and comply with the organization's security policies. While all audiences must comply with all policies, each audience may find that it spends most of its time in just certain sections of the policy collection.

## POLICY DEVELOPMENT

Now that the various audiences for security policies have been discussed, what will the security policies say? Is there a format or structure that other organizations use that works well?

### **Security policy best practices**

Information from large companies and from intelligence sources such as Gartner and SANS<sup>1</sup> suggests that organizations adopt a structured top-down model for organizing their collection of policies and requirements. This will make it easier for different audiences to find the information they need. The different types of security policy statements, from the top down, are as follows:

- *Security policy charter or mission statement.* This is usually a short statement in the form of a memo from the organization's CEO or CSO that declares the existence of official enterprise security policy and that compliance to security policy is not optional. The statement will usually specify who is responsible for security in the organization – which establishes and enforces policy. The consequences of deliberate violations of policy (generally speaking, reprimands or termination) are usually stated here.
- *Security guiding principles.* These are high-level statements describing the primary objectives regarding the protection of the organization's information assets and infrastructure. In highly abbreviated form, some examples of guiding principles include the following:
  - All company documents will be classified according to their sensitivity and handled accordingly.
  - Users will use good passwords and will not share their passwords with others. Users will not use other users' accounts.
  - Users will only have access to information, equipment and functions that they need to perform their job functions.
  - Users will practise behaviours that will protect the confidentiality of information: they will not discuss confidential matters in crowded public places such as restaurants or airports, and they will make sure that others in public will not be able to view confidential information on their laptop or hand-held computer screens.
  - Proprietary information will always be encrypted when it is sent over public networks such as the Internet. Proprietary information stored on portable media such as laptop computers, backup tapes, diskettes and CD-ROMs will be encrypted.

- Customers' private information such as credit card numbers, driving licence numbers and other highly sensitive information will be encrypted in storage and in transit, and will be viewable by few employees and only on a need-to-know basis.
- All employees in the organization will handle and process information in a manner that conforms to all local laws.
- All access control mechanisms will explicitly deny access to all, except to those that are explicitly permitted.

The preceding examples do not intend to represent a complete set of guiding principles, nor are they necessarily a precise reflection of today's best practices. Instead, they are abridged examples of high-level principles that an organization might adopt to set the tone for its larger collection of security policies and requirements.

- *Security policies.* These statements describe – using more detail than in the guiding principles – what the organization will do to protect its information assets and infrastructure. Like guiding principles, security policies are general in nature and are not technology specific. Policies do, however, get closer to 'how security is done' in the organization. Some examples of security policies include the following:

- Systems and applications will utilize and enforce best practice password complexity to include letters, numbers and special characters. Passwords will expire every 90 days or more frequently, and cannot be changed more frequently than every seven days. Old passwords cannot be re-used. Passwords cannot be the same as the user account name, nor can they be a common name or dictionary word.
- All Internet-facing applications will be implemented using a three-tier DMZ (de-militarized zone) architecture. Web servers will reside behind a firewall in the outermost DMZ. Application logic will reside behind a second firewall in the second DMZ. Application data will reside behind a third firewall in a third DMZ. In no cases will connections be allowed to cross two firewalls; for instance, the web servers will not have direct access to application, nor will business logic servers be accessible from the Internet.
- Intrusion detection systems will be placed on all DMZ networks, as well as on all internal networks, in order to detect intrusions and other anomalous events.

The preceding examples illustrate the general and technology-neutral nature of security policies. This makes security policies durable, flexible and adaptable to new technologies and products.

- *Security requirements.* These statements give precise guidance specifying how security policies are to be implemented. Unlike guiding principles and policies,

security requirements *are* technology-specific. However, requirements that specify which products are to be used may infringe on the role of an IT standards function. This is just one example that illustrates the necessity for security policies and requirements to be developed not in a vacuum but in collaboration with several other functions in the organization. Some examples of security requirements include the following:

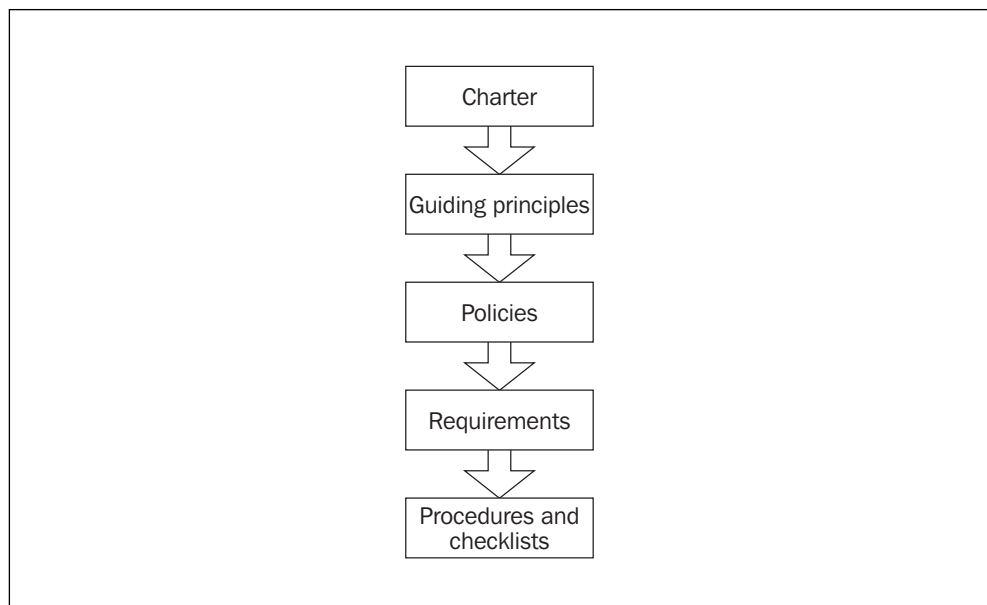
- Applications will externalize authentication to a central LDAP v2 server.
- SSH v2 will be used for all remote administration to servers. Telnet, rsh and rcp will not be used.
- Servers will not have a `/.rhosts` or `/etc/hosts.equiv` file.
- Administrative passwords will be a minimum of ten characters in length and will include at least one upper case letter, two numerals and one special character.

These examples illustrate the granularity and detail typical of requirements. Requirements describe how security policies are to be carried out – with which tools and protocols.

- *Security procedures and checklists.* Security procedures are the step-by-step instructions used by practitioners to implement security requirements. Checklists permit a system or network administrator to work quickly through a list of detailed steps or configurations to ensure that a hardware or software component is security compliant.

Table 4.1 contains a summary of the four types of security policy statements and their use. Figure 4.1 illustrates their hierarchical relationship.

**Fig. 4.1** Security policy hierarchy



**Table 4.1** Summary of security policy statements

	<i>Guiding principles</i>	<i>Policies</i>	<i>Requirements</i>	<i>Procedures, checklists</i>
<i>Created by</i>	Security department	Security department	Security department	Security department and operations
<i>Focus</i>	Highest-level priorities and principles guiding everyone on the protection of information assets	General non-technology-specific statements describing how information assets are to be protected	Technology-specific statements describing what components will be used – and how they will be used – to protect information assets	Step-by-step instructions on how to carry out security requirements
<i>Audience</i>	Everyone in the organization	Everyone in the organization. Some sections may apply to certain job classifications such as technology architects, developers, integrators and administrators; also may apply to extranet business partners	Developers, integrators, administrators, hardware and software vendors, extranet business partners	System, network and data administrators; software developers
<i>Collaborators</i>	Executives, business and technology strategists, and principal system architects	Business and technology strategists, architects, principal system designers	Architects, designers, operations personnel	Designers, operations personnel

### Sources for policy content

Now that it is clear that security policies and requirements are the essential guide for defining acceptable behaviour and uses of technology, how does the organization go about writing them? Would it be reasonable for the security organization to create a

committee or working group to begin drafting security policies and requirements? This would certainly get the job done, but it would take a great deal of time and effort.

A viable alternative is to examine existing collections of security policies as sources for content. Some of these sources are described here.

- *ISO 17799*. Formerly known as British Standard (BS) 7799, this is a comprehensive and up-to-date collection of security policies as well as guidance on the implementation of security policies. This entire collection is copyrighted, but organizations may purchase copies and consider adopting it in whole or in part. It is available from [www.iso17799software.com](http://www.iso17799software.com).
- *Information Security Policies Made Easy*. This vast collection of security policies by Charles Cresson Wood is another rich source of security policy language. This collection can be purchased from online booksellers.
- *RUSecure Information Security Policy Suite*. This is another extensive commercial collection of security policies, available at [www.rusecure.co.uk](http://www.rusecure.co.uk) or [www.information-security-policies.com](http://www.information-security-policies.com).
- *Generally Accepted System Security Principles (GASSP)*. This is an excellent collection of high-level security principles published in the early 1990s by the Organization for Economic Cooperation and Development (OECD). GASSP is available from [web.mit.edu/security/www/gassp1.html](http://web.mit.edu/security/www/gassp1.html).

### **Publishing security policy**

The security department must consider all of the audiences of security policy and decide how to make the content available to them. One or more of the following formats are among the obvious choices:

- *Online HTML*. Most intranet, and possibly extranet, users are accustomed to accessing reference material online. HTML lends itself well to online content.
- *Online PDF*. Content written for Adobe Acrobat Reader works well online, with new browsers' ability to display PDF documents within the browser. Printed policies' format will be better than HTML, since there will be headers and footers on each page.
- *CD-ROM*. External entities such as vendors may not have online access to the organization's online security policy content. CD-ROM may be a nice alternative, as it gives users a near-online experience: they may view the same HTML or PDF content that intranet users see.
- *Hardcopy*. To some, nothing beats paper when it comes to reference material. One can write in the margins and easily flip back and forth between different sections. Further, if the organization does not wish to give electronic copies (online or CD-ROM) to external entities, hardcopy is probably the last option.

A ‘sweet spot’ for efficiently reaching all audiences may be to publish security policies in Adobe PDF format, make it available online, and print hardcopies for those who want or need hardcopy references. Offline (CD-ROM and hardcopy) content should carry a disclaimer stating that the policy being viewed may have been superseded.

## **AWARENESS**

As Chapter 5 will explore more fully, the organization cannot hope to secure its information assets unless its people are aware of its security policies. This goes beyond awareness. Do the organization’s employees know *how* to protect its information assets? All employees, regardless of job function, must be armed with basic information: where to find security policies, how to use them in their individual job functions, and what to do if they think that policies are being violated.

### **Where are the security policies?**

All employees need ready access to the organization’s security policies. This does not mean that every employee needs access to every topic and section. But in every organization, there are common principles that everyone needs to know.

Not all employees have access to online policies and it may not be practical to have the complete bound set in all locations. This makes sense when considering how much security policy information manufacturing assemblers or customer support centre representatives need to know.

In order to reach effectively all employees in the organization, it will be valuable to reduce the entire collection of security policies and requirements to a set of enduring principles. Whether you call this ‘information hygiene’ or ‘working smart to protect the company’, it will be a better use of company resources to deliver simple, wide-reaching messages to diverse audiences.

### **Promotions and campaigns**

Security information must be marketed just like any other topic. To get above – or at least on the level with – the cacophony of other messages, promotions and campaigns, the organization’s basic security values must be built into an effective campaign in order to reach the masses.

In order to be successful, a security campaign must be easy to recognize and understand, and security messages will need to fit into ‘sound bites’. Different tactics work in different companies: posters and e-mail messages may work in some, whereas voice-mail broadcasting or ‘town meetings’ may work well in others. All these forms of information marketing are working towards the same

goal: to get employees thinking and talking about information security, and that information security is vital to the viability and success of the organization.

At a minimum, employees need to know:

- What are the guiding security principles that everyone needs to know?
- Are they doing the right thing to protect information assets?
- Do they know who to call if they suspect a violation of security policy?
- Do they know where to find more information, whether it is online or if they want to talk to someone?

Those responsible for information security want the organization's employees to know a lot more about how to protect its information assets. They should want information security to be integrated into the company's culture. To make such fundamental changes, they need to have a completely different skill than security: they must know how to reach people and change how they think about something. This will require the talents of people in other parts of the company, or even someone outside the organization who knows how to get into its consciousness effectively.

### **Attestations and self-testing**

Some organizations want or need to go beyond publishing security policies and putting up some posters, and actually know that each person has not only read the company policies but also understands them. The security or legal department can create an 'I have read and understood company security policy' document that each employee should read and sign. This would give the organization a written record that would allow it to track who has – and who has not – acknowledged awareness and understanding of security policies.

Some organizations may want to go so far as to test employees on their knowledge of security policy. This can give the organization the assurance that individual employees not only have signed an 'I have read and understood company security policy' document but have also demonstrated actual working knowledge of the organization's security policy.

There are two reasons why an organization would wish to consider testing its employees on security policy:

- *Competence.* The organization may want its employees to take knowledge of security policies more seriously by telling them that they will be tested on it. The organization can offer interactive or paper-based exams, testing employees on their knowledge of the entire suite of security principles and policies, or just the portions that are most important for the organization. Consistently low test scores in specific topics can indicate areas where more user awareness or education is needed.

- *Accountability.* Employees who tried to feign ignorance of security policy despite having signed a ‘I have read and understood ...’ letter would have a difficult time disavowing test results that showed that they actually did understand them.

## **ENFORCEMENT AND EFFECTIVENESS**

Security policies and requirements are effective only if they are enforced. But to be enforceable, policies and requirements need to be measurable and results attributable to individuals or departments. Without accountability and enforcement, the policies and the security department might as well not exist: a security department that fails to enforce its policies will not get much respect.

### **Detection points**

Where are the enforcement and measuring points, and how can the security group effectively measure and enforce security? Several examples are shown here, but given the wide variety of business processes, organizations can use these or find other situations.

- *Project requirements.* At a project’s requirements phase, the security group can strive to get its standard suite of security requirements, as well as any project-specific requirements, inserted into formal requirements statements.
- *Proposals.* Vendors that respond to requests for proposals should describe, in detail, their compliance or non-compliance to each security requirement. The security group must be sure to review carefully the responses and press for changes or improvements prior to acceptance.
- *Assessments.* Hands-on detailed inspections of networks, servers and applications will reveal gaps in security policy compliance. These gaps can be documented, prioritized and described in terms of the risk to the organization and effort required to remediate them. Points of leverage are described later in this section.
- *Security logs.* The contents of security logs from applications, servers, intrusion detection systems, firewalls, virus management systems and network management systems will reveal attempts (successful or not) to violate security policies and requirements.
- *Security tip lines or hot lines.* The security department can establish one or more avenues for employees in the organization to report activities they suspect to be in violation of security policies. While employees can be encouraged – or even rewarded – if their tips are substantiated as genuine risks to the organization, they also need the assurance that they will not be the victims of retaliation.

## Enforcement points

In the previous section, some of the means for detecting security violations were discussed. This section describes the points of leverage that can be used to enforce situations into compliance. Some of these leverage points are described here.

- *Reprimand and termination of employees.* Employees who by their actions violate security policies may need to be reprimanded or have their employment terminated. Every security policy should state that the consequences of non-compliance may result in reprimands up to and including termination.
- *Refusal to purchase, install or support.* Depending upon where the violation was detected, there are several opportunities for the security department to bring about compliance. The security group can influence the purchasing department to delay or cancel the purchase of products or services that are noncompliant. The security group can convince engineering, installation or operations groups that the product, application or service should not be installed or supported until matters of compliance are resolved.
- *Shutdown or disconnection of networks, servers, services and applications.* The security department may be able to convince the operations department that a noncompliant server or application should be shut down or disconnected until it can be brought into compliance. This extreme measure needs to be carried out with caution, since the impact may be widely felt, and may possibly be to the detriment of the organization's revenue stream.

The security department should not wait until a crisis to begin the dialogue with other groups to influence their actions. Instead, strong relationships need to be established well in advance, and norms of interaction well established and documented. Then when the security department requests another department to become its leverage point in a security compliance matter, the whole affair and responses to it should not bring surprises, but it should proceed more like normal business. Properly constructed security processes, even when controversy may erupt, can be carried out matter of fact.

## Security exceptions

Up to this point, it may appear that the security department has at its disposal only extreme measures to bring about compliance. On the contrary, extreme measures should be used only in extreme situations – the punishment should fit the crime (or risk of the crime).

Most security compliance matters can fall into the middle ground – in between doing nothing and ‘pulling the plug’. Often a compromise can be negotiated that will permit a project to proceed but still gives the security department assurances

that systems and applications will eventually be compliant. The security exception request is the central written instrument that fulfils this role.

Take, for instance, an example of a new application that an organization is purchasing. The project team has inserted security requirements into the request for proposal and system requirements statements, but upon assessing the responses, there are some security requirements that are not met by the application. These gaps are documented and prioritized in terms of risk, as well as effort to fix. If these matters are discussed early on in the project, many of the requirements can be made to be compliant. However, some requirements may go unmet. Should the security department insist that the project does not continue until every requirement is satisfied? Perhaps, but if the security department does this too often, it will no longer be invited to participate in new projects.

Often it is acceptable to find middle ground. In order to keep the project on schedule, the development or integration group can make a written commitment to bring the project into full compliance at some date in the future. This written commitment is the security exception request. The senior manager who signs the exception request pledges to close all documented security matters in a specified timeline and that all fixes will be budgeted and resources allocated.

The group that must manage and operate the application or system should be a participant and signer in the security exception request, since that group will play a role in managing the security gap and may still be responsible for keeping the application or system secure, despite the gaps.

Finally, the security department itself should agree to the proposed fixes and timelines and be an approval party in the security exception request. Being responsible for securing the entire enterprise, the security department also needs to agree that the project can continue despite the risks and that any stated mitigation prior to compliance will sufficiently reduce risk.

As the deadline for promised compliance approaches, the security department must contact all parties and remind them of the commitments that were made when the exception request was signed. This is the place where the real effectiveness and value of the security department is tested: it must firmly resist the perpetual continuation of security exceptions, but insist that remediation be completed as originally specified. It is only at this juncture in the exception request process that the security department should bring out the big guns and suggest or threaten that additional steps be taken to bring about eventual compliance.

## **Risk agreements**

The security department may elect to insert into its security exception process a provision that will permit a project to continue, despite a known security gap that will not be mitigated in the foreseeable future. Rather than perpetually renewing

security exception requests that everyone knows will never be fulfilled, the security department should instead turn to the risk agreement as the final solution.

The risk agreement provides for the approval of a specific perpetual violation of a security policy or requirement that will not be remediated any time soon. In the risk agreement, the senior executive who wants the application, server, product or service to continue operating, despite the gap(s), assumes full responsibility for any events or circumstances that arise from the existence of the security gap. This notion of accountability of a senior executive will give him or her an opportunity to consider seriously whether continued noncompliance is really the best solution. Full accountability brings the sober realities into sharp focus and helps the organization move beyond the matter of the security gap.

A group of senior executives should approve a risk agreement – this will ensure that all senior stakeholders are willing to make the risk decision and that it is being done for the good of the organization and not just one senior executive’s job.

## **SUMMARY**

---

The basis for information security policies lies in the fundamentally important aspect of the business: protecting information assets. Information that is vital to the continuing operation of the business is worth keeping, and assets worth keeping are deserving of protection. Security policies define how assets are protected.

They also define responsibilities for various parties in the organization, as well as measures to be taken to protect information. These measures include user accounts, users’ access to information, access control functions, audit logging, confidentiality, integrity, availability and incident response.

Security policies are arranged in a hierarchy. The top-most document in the security policy hierarchy is the charter document, which declares the existence of security policies, who is responsible for writing and enforcing them, and the consequences for noncompliance. The next tier in the hierarchy is the guiding principles, which describe the fundamental values of information and information asset protection. Next are the security policies, the more detailed but still nontechnical guidance defining the principles of protection. Security requirements make up the next lower tier – these provide the details describing how the security policies are to be carried out. Finally, procedures and checklists provide the step-by-step details describing how requirements are performed.

Security policies often need to reach different audiences. Vendors and others who provide applications, software or hardware to the organization need to understand the organization’s security requirements. Business trading partners use policies and requirements in order properly to secure the organization’s information. Developers

and integrators need the guidance of policies and requirements in order to be able to build security-compliant systems.

There is rarely a need for organizations to develop security policies from scratch. There are a number of high-quality collections available, including ISO 17799, the *Information Security Policies Made Easy* collection by Charles Cresson Wood, and the *RUSecure Information Security Policy Suite*.

The organization must consider how different audiences will access security policies. If all parties do not have access to online policies, alternative means must be developed, including hardcopy or CD-ROM.

Security policies are of little value unless everyone in the organization knows that they exist and can find them when needed. In order to compete for 'mind space', the security department will need actively to promote security awareness to ingrain security into the corporate culture. Some organizations will want to take the additional steps of having employees sign statements attesting to having read and understood the organization's security policies. Organizations can also test some or all of their employees to measure actual working knowledge of security policies.

Security policies are also of little value unless the organization has ways of measuring and enforcing compliance. However, the security department has many opportunities not only to detect but also influence outcomes. For instance, it can insist that security requirements be a part of every request for proposal that is sent to hardware, software and service vendors. The responses from those vendors should be carefully examined to gauge the degree of compliance to policies and requirements. The results of security audits and assessments also provide much valuable information. Audit trails provide data on events, whether they are transactions or changes made to systems. Security tip lines and hot lines can record information regarding suspected security violations in the organization.

The security department has plenty of enforcement areas. Security policies should state that reprimand and termination are possible consequences to violation of security policies, and the security department should not be afraid to insist on these measures as appropriate. Security can also request that the purchasing department does not acquire noncompliant products, as well as request that operations does not install or operate noncompliant applications and systems. It can request that operations shuts down or disconnects networks, servers, services and applications as necessary.

There is plenty of open ground between the extremes of inaction and the measures described in the previous paragraph. The security department can facilitate negotiated settlements – called security exceptions – whereby other groups commit to fixing noncompliant systems and applications within a specified timeframe. Systems that cannot be remedied can be the subject of a different kind of settlement, called the risk agreement, where the group that wants the operations department to support a noncompliant application must accept responsibility for the risks associated with it.

## **NOTE**

- 1 SANS is an acronym for Systems Administration, Networking and Security. SANS is a source for security research and other information for technologists and security professionals.

# Security is about people's behaviour

- Introduction 93
- Technology is not the solution 93
- The 'people threat' 94
- Mitigating the threat 99
- Trust 100
- Summary 102



## INTRODUCTION

Despite all the technical controls and mechanisms that an organization can put in place to protect its information assets, the security and integrity of its information ultimately relies upon people. Even in the finest security environment where there are best-in-class authentication, authorization and access control mechanisms protecting information, people with access to protected information can all too easily damage, alter, copy or remove the information, sometimes without leaving a trail.

Organizations that rely solely on technology for their information security are in for enormous disappointments. Because *people* create, develop, manage and share the organization's most valuable information assets, the same processes and technologies that are used to perform these tasks can also harm the organization. Computers and software are not smart enough to distinguish between proper and improper handling of information.

Ultimately, trust must rest with the people in the organization. With or without security technology, the security of an organization's information assets depends upon its people. This chapter will explore these concepts and explain why and how the organization's best line of defence is with its employees and contractors.

## TECHNOLOGY IS NOT THE SOLUTION

Many a CIO, CSO or CEO, believing that technology can solve security problems, has purchased firewalls, encryption, virtual private networks, intrusion detection systems, public key infrastructures, token authentication and biometrics. Sooner or later, these managers will discover that – for the most part – technology is not the solution because technology is not the problem.

Technology is *related to* the security problem, and here is why.

### **Technology is an amplifier**

Technology is a tool that *amplifies* the effort of the individual so that he or she can do the work of many people at once, or in other cases do in a few minutes what would otherwise take weeks or months. Technology amplifies a person's work in order to make the organization more accurate, more efficient or more responsive to the needs of a customer or market.

The benefit of technology is that it amplifies the effort of the individual. But what if the individual wishes to cause harm to others? A malevolent person using technology as a tool can cause great harm to vast numbers of individuals or large organizations. Consider the Code Red or NIMDA worms in 2001: each was probably written by a lone individual, and yet the effort of a single person was able to adversely and seriously affect hundreds of thousands of computer systems around the world.

Peter Gregory's law of technology is:

*The capacity for technology to amplify productive work can be equalled or exceeded by the capacity for technology to amplify destructive work.*

For the most part, security problems in organizations are a result of people's deeds and misdeeds, and not so much about technology.

## **THE 'PEOPLE THREAT'**

Few, if any, organizations can refute the statement, 'The organization is the people who comprise it.' In other words, take away the people – and with them their knowledge, experience, contacts and work patterns – and the organization will cease to function almost immediately.

On the other hand, people, not technology, are at the root cause of nearly all security incidents. It is a well-known fact that the majority of security incidents are 'inside jobs'. In other words, most security events (over 70 per cent by many estimates) are not a result of attacks from outside the organization but rather are security breaches perpetrated from within.

The remainder of this section will discuss the ways in which people's behaviour can result in security incidents.

### **The knowledge gap**

Most people in an organization want to do the right thing, including protecting the organization's information assets. The problem is, many people do not know what the 'right thing' is.

The role of security policies in an organization is to define formally 'the right thing'. Lack of awareness with regards to security policies falls into five levels:

- *Existence of security policies.* Employees may not even know that the organization has security policies.
- *Accessibility of security policies.* Employees must know where to find the organization's security policies.
- *Comprehension of security policies.* Employees may know that the organization has security policies, but they might not know what the policies say. The policies may be difficult to understand, hard to read or hard to find.
- *Applicability of security policies.* Employees may not realize that there are security policies that specify *what* they need to be doing, or *how* they need to be doing what they are doing.

- *Relevance of security policies.* Employees might not understand that specific policies that they know about apply to their function or activities.

These five points – existence, accessibility, comprehension, applicability and relevance – are links in a chain. Every employee handling company information must have full awareness of each of these five aspects of security policy so that they can carry out their jobs correctly and thus protect the organization's information assets. This is explored more fully in Chapter 4.

Some real-life examples of the kinds of things that employees can do – when they are unaware of the facts of security policies – include the following:

- *Homework exposed.* An employee might transfer some secret company documents to his or her home computer to work on them over the weekend. Because a home system is frequently not equipped with up-to-date anti-virus protection, a successful Trojan horse attack causes the secret documents to be e-mailed to everyone on the employee's personal e-mail address list – including former colleagues who now work for competitors, suppliers and customers. *[Security policies generally state that company information is not to be transferred to personally owned computers.]*
- *Group account abuse.* A new system administrator, eager to show quick results, creates a group account on an extranet server to give several individuals access to proprietary information. One of the individuals misuses his access privilege and commits corporate espionage by selling the documents to a competitor. *[Security policies generally state that there should be no group accounts, since this creates an accountability gap where it is difficult or impossible to attribute actions associated with the account to any single individual.]*
- *Sharing account information.* An employee shares her userid and password with a colleague, since that colleague's account was accidentally locked; the explanation given was that her colleague's Caps Lock key was on and after attempting to enter the password three times, the account was locked. Later, if the colleague makes a mistake, the employee herself would likely be blamed for it, since the actions would be attributed to the user account, which she shared with someone else. *[Security policies generally state that users should not share their login and password information with others.]*

## The privilege gap

User account privileges specify what information a person is allowed to view, alter, move or remove, as well as what functions a person is permitted to perform. Here is a sample of the kinds of events that can and do occur in organizations:

- *Developer access to production systems.* An application developer is given one-time access to production systems during an emergency where he needs to repair a program. The privilege is never removed and the developer begins to alter routinely production software, circumventing configuration management and change management processes. *[Security policies generally permit such emergency arrangements, but they must be promptly rescinded when the emergency is over.]*
- *Job change.* An employee in the purchasing department generates accounts payable (A/P) requests as a routine part of her job. Later, she transfers to the accounts payable department, where she now approves A/P requests. Her ability to generate A/P requests was never removed and she is in a situation where she can not only *generate* A/P requests but also *approve* them. She is then able to extort money from the organization and does so for many months until an audit uncovers the scheme. *[Security policies require ‘separation of power’ so that no single individual has too many privileges.]*
- *Abuse of inadequate privilege controls.* An employee discovers that all users of a web-based application have complete privileges. The application displays only the intended functions to each employee but does not block an employee from executing any function if they know the URL. If employees e-mail the privileged function URLs to each other and bookmark them, everyone can perform all functions. *[Security policies generally require stronger privilege controls in order to enforce effectively separation of powers; also, employees should restrain themselves from abusing the lack of controls in this situation.]*

Abuse of privilege is one of the biggest problems in information technology. This is because the task of privilege management is very time consuming, and often an organization does not provide adequate resources to this ‘overhead’ function. As a result, security controls and the integrity of company processes and information suffer.

## Audible conversations

People are by nature gregarious beings – we love to congregate and talk about everything. We especially enjoy talking about the failings of other employees, the latest investigations, and top-secret stuff of almost every kind. People love to share ‘hard-to-get’ information because of the attention it brings. This hard-to-get information may frequently be a confidential matter that should not be shared with others.

Not all spoken conversation is bad, however. Face-to-face and telephone conversations are one of the primary modes of communication, and, gossip aside, spoken conversations are downright essential for the business. To assure the confidentiality of business information, people need to be aware of their surroundings.

### ***Who is listening?***

People need to think before they speak – this needs to be as natural as the conversation itself. Before they begin talking, they need to do a quick ‘look around’ to see if any other people are within range to overhear what is about to be said. Before the conversation begins, people may need to:

- shut the door
- find an empty office or conference room
- walk down the hall away from others
- go outside
- lower their voice to a whisper
- decide to discuss it later in more secure surroundings.

People should avoid restrooms (someone in a cubicle could overhear) and stairwells (sound carries great distances – someone on another floor could overhear). From the awareness perspective, asking people to think ‘who is listening?’ will also give them pause before sharing information that should not be shared.

Prior to a sensitive telephone conversation, people should think for a moment about the phone they are using. If they are using a cellular phone, they should make sure that the connection is digital (which is very difficult to eavesdrop) and not analogue (which is easy to eavesdrop). Cordless phones may be risky, too, for the same reason – a wired landline may be preferable.

### **Shortcuts**

People are under pressure to get more work done in less time. Basic human nature equips us with the ingenuity to seek and use techniques that will make our jobs easier to do by changing procedures or skipping steps. System administrators, software developers, clerks and managers all want to do only what is necessary and move on to the next task. The ragged edges on an incomplete job may provide opportunities for abuse by others who either take the time to find the gaps or uncover them accidentally.

Here is a sample of events that can happen anywhere:

- *Inadequate user account management (UAM) processes.* An employee transfers from one department to another and needs the special privileges from his old department removed. However, the transferred employee is still training his replacement and sometimes needs to fill in for her. The department manager convinces the user account management (UAM) department to let the transferred employee keep his old privileges. But the lack of follow-up by the UAM department means that the transferred employee's excess privileges will probably never be trimmed.

- *More inadequate UAM processes.* Employees in the UAM department are overworked: they are under constant pressure to create accounts for new employees. Some UAM technicians no longer remove the accounts of terminated employees, since nobody notices anyway. Some disgruntled former IT employees discover that their accounts are still active – they access their former employer’s systems and post sensitive data on a website.
- *Failure to protect information.* An organization requires its employees to encrypt sensitive files on their notebook computers. A vice-president is travelling with his notebook computer, which contains long-range product and marketing plans. He is busy and non-technical, and unfamiliar with the encryption programs on his system; thus, he rarely encrypts sensitive files. His system is stolen while he is travelling and the information that should have been encrypted is posted on a website.
- *Excessive access privileges.* A UNIX system administrator, anxious to get a newly installed application running, suspects permission problems and loosens access permissions for the application to the point that there are virtually no file access permissions on the system. Once the system is running, the UNIX administrator is reluctant to put permissions back to the way they were, but instead leaves the system in a highly vulnerable state and moves on to the next project. Months or years later, an intruder finds this same server an easy target.

## Theft of information

As more of society’s functions are automated and its information placed in computers, the value of the information and functionality increases. The onset of e-commerce created thousands of computers – set up by relatively inexperienced people and using software that is repeatedly exploited – that contain high-value information such as credit card numbers, taxpayer identity numbers, customer lists and so forth. A person motivated by greed and lacking ethics and appropriate right/wrong judgement can break into systems containing high-value information and use it for personal gain. Recent history is flush with the stories of hackers who have broken into e-commerce sites, stolen thousands of credit card numbers and posted them on other websites. Information found in customer lists – names, addresses, dates of birth, etc. – can be sold to and used by identity theft crime rings.

The spate of security holes in Microsoft-based operating systems and applications has made it easy for intruders to break into and steal information. Hackers do not, however, have the exclusive ability to steal. Company insiders with easy access to information such as credit card numbers find it too tempting to take this information for personal gain, whether they use it themselves or sell it to others. Those who are old enough to remember credit card vouchers with ‘carbons’ may recall a time where unscrupulous restaurant employees would fish the carbons out

of the rubbish and use or sell them. Carbonless vouchers reduced this risk. But today, credit card numbers are listed by the thousands in online systems. Poor application architectures and the failure to install security patches produce a lucrative opportunity for an intruder or insider.

Unlike physical or monetary assets, it can be difficult even to know when information has been stolen – because it can still be present and there may be few, if any, clues indicating that the information was taken.

## **MITIGATING THE THREAT**

If it is so easy for an insider to take information out of an organization, what is stopping people from doing this more often? This section will explore some of these ideas in detail.

### **Ethics**

Fundamentally, most people know it is wrong to take information out of the organization for personal gain or inappropriate disclosure. Because no organization wants to take on the job of teaching morality to adults, this makes screening for behavioural characteristics an important part of the hiring process. Behavioural interviewing, background checks and reference checks are essential for uncovering these traits.

### **Conscience**

Most people have had the childhood experience of taking someone else's belongings, and vividly recall the trauma of having lived with the secret that gnawed at their conscience, as well as the humiliation of being caught or confessing to the deed. Those without this personal experience may have seen another child in a similar circumstance who was publicly exposed. These childhood lessons create powerful memories that influence one's long-term behaviour.

### **Security policies and awareness**

Every organization needs to have written security policies that define the concepts of the corporate ownership of proprietary information. Every employee and contractor must be made aware of these policies and the consequences of failure to conform to them.

Organizations may consider having each employee and contractor sign a statement attesting to their knowledge of security policies and the consequences of nonconformity. This written record gives the organization a paper trail, and the

act of signing such a document sends the message that the organization takes its security policies seriously. This is covered in Chapter 4.

### Notices of confidentiality

Markings or labels on all company documents that state that they are confidential and proprietary (and potentially the level of confidentiality) are a constant reminder that such information is the property of the organization and that it should not be disclosed. While such markings cannot prevent information from being mishandled, it may trigger the hesitation needed to keep proprietary information safe.

### Fear

For most, the betrayal of trust is a venture into the unknown, with many real and imaginary consequences awaiting the would-be conspirator.

- *Being caught.* Many employees believe that an organization's security controls and event-logging capabilities are better than they actually are and that individuals inside the organization will discover their illicit actions.
- *Criminal prosecution.* The prospect of being caught and prosecuted can be a powerful deterrent. For this to be effective, an organization must not only have a policy stating that it will prosecute those who illegally disclose or mishandle information, but it must consider doing so visibly as an example for any others considering taking similar actions.
- *Losing employability.* The prospect of being caught can result in difficulty in being re-employed in the future. Even if an employee is not tried and prosecuted, news of malevolence has a way of getting around. With people moving from organization to organization, hiring managers generally have quite an extensive network of contacts in other companies and frequently call upon each other when an individual applies for a job. Stories of thievery and dishonesty tend to follow an individual as long as they remain in the same line of work or within the same geographical area.

## TRUST

The organization's employees operate the business: they develop information, transform information, communicate information and use information to get their jobs done. Frequently – if not exclusively – the information they are using is proprietary and confidential.

Employers implicitly trust their employees to use the right information, to develop good information, to improve information and to communicate information.

Employers must trust that their employees will handle information properly: that information will not be shared with or sent to people who do not have a need-to-know regarding the information.

### **Technology is not a barrier**

Short of taking away their technology (which causes a few problems of its own), there is very little that an employer can do to prevent employees from spiriting information out of the organization. The same tools and technologies that facilitate the development and collaboration of information within the organization are used illicitly to take the information out. And, of course, people can always walk out with information, with no help from technology at all.

Beyond a few best-practice remedies such as limiting the size or frequency of outbound e-mail messages, or even random or comprehensive scanning of outbound e-mail, it is unlikely that the organization can prevent – or even detect – an employee who would inappropriately send proprietary or confidential information out of the enterprise. In fact, the implementation of measures or controls may have the opposite effect: such controls, if discovered, could send a ‘we do not trust you’ message from employer to employees. This distrust will result in resentment, betrayal of trust and a higher risk of losing the best employees.

Whether an organization's most valuable information is paper-based or stored electronically, there is practically nothing that the organization can do to prevent employees from taking that information out, whether they do it for personal gain or ‘just for the hell of it’.

### **Building trust**

Employers must manage the sometimes delicate trust balance. Knowing that employees could walk out at any time with information that could cause great damage if disclosed, employers must build a good working relationship with all its employees, based on trust and value. Such a theme could go like this:

- we value you – you are the lifeblood of the business
- we trust you – we have carefully chosen you to be on our team
- we will develop you – we will train you to give you even better judgement on handling proprietary and confidential information.

The trust relationship is vital to employees' performance and satisfaction. Trusted and valued employees, with the proper training, will want to do the right thing and will be more likely to remain loyal to the organization.

## SUMMARY

The protection of information assets and infrastructure depends primarily upon the actions of employees and contractors. Ultimately, there are few – if any – technological barriers that will effectively stop an insider from altering, misusing or disclosing proprietary information.

Many conditions and circumstances will lead people to misuse the organization's information assets. First, they may not know how to identify or handle proprietary information. Also, they may have more privileges for access to and handling of information than they need. This can be a result of improper use of or application of privilege management, or job changes that result in employees having more privileges than they should. Next, confidential information may be inappropriately disclosed through improper procedures – proprietary documents may be left exposed where others can view them or conversations of a sensitive nature may be overheard.

Further, people will cut corners to make their jobs easier and so that tasks will take less time. Security-related tasks and procedures, often thought of as unnecessary bolt-ons, are sometimes the first to be eliminated.

Fear plays a significant role in deterring individuals from betraying the organization. Fear of termination, the inability to be re-employed, and criminal prosecution can all have a significant impact on one's lifestyle and freedom. Only the most hardened and determined – or foolhardy – individuals will proceed with a plot to steal, destroy or disclose proprietary information given these consequences.

There are, however, statements and actions on the part of the organization that will influence people's behaviour and, in many cases, give them pause before they carry out their plans. These include the establishment of a security policy that defines confidential and proprietary information and the consequences of misuse or disclosure. The security policy must be well publicized and it must be publicly and vigorously enforced.

The other matter influencing employee behaviour is trust. Ironically, the more an employer is apt to trust its employees with the freedoms and access to its information assets, the more likely it is that employees will be trustworthy and do the right thing. Since, ultimately, the organization's employees and contractors have nearly complete control over its information assets and infrastructure, employees and contractors who are valued and trusted will have a tendency to reflect that trust and be loyal to the organization.

# Protecting corporate information beyond the corporate boundaries

- Introduction 105
- The new world 105
- Regaining control 107
- Summary 112



## INTRODUCTION

Extranets have the potential to take the organization to risk–reward extremes. The power to leverage additional resources or create greater efficiencies is enough to make the experienced executive weak-kneed.

While the world is apparently rushing to the extranet bonanza, it is not risk-free by anyone's imagination. On the contrary, when taking a serious look at the risks involved, one will have to work hard to develop the business case to be able to proceed. However, the risks are mostly surmountable, at least to the point of being able to consider moving forward with an extranet. This chapter will explore many avenues for risk mitigation and will prepare the reader for the many possible outcomes.

## THE NEW WORLD

Imagine a world where your organization routinely transfers and stores its assets at the premises of your suppliers and customers. Visualize these assets as being out of your direct control, entrusted to your trading partners' care. Now, imagine that these are *information* assets that we are talking about. You know, the kind that are difficult to control. The kind that, when stolen, are still there.

This is not the stuff of fantasy or futurespeak but a description of what is being done today in almost every large technology-enabled organization in the world. We have given our suppliers and customers the ability to peer into our information systems and our data, either by sending them copies of (some of) our information assets or by giving them electronic access to our information systems.

Organizations are no longer isolated. Their information is no longer locked in the castle, but distributed throughout the kingdom, and even beyond the kingdom. Welcome to the extranet era.

The threats and vulnerabilities to extranet environments include the following:

- *Transitive trust.* When an organization connects its enterprise network to its trading partners' networks, the organization's network is then only as safe as that of the trading partners. If the trading partners' networks have any inferiorities in defence, whether anti-virus, perimeter, access control, etc., those inferiorities will make one's own enterprise equally vulnerable. Consider that the trading partners are likely have extranet partners of their own and so forth. So while for a moment you think that you can protect your enterprise by exercising complete control (!) over them, can you control *their* trading partners' security and business practices? The bottom line: one's own enterprise network is only as secure as its weakest connected neighbour.

- *Lack of physical control.* One important ‘defence in depth’ characteristic of the enterprise network is not only the logical control but also the physical control one has over it. The organization locks its buildings, especially its data closets, phone rooms and data centres, and it records who enters which rooms at what times. Even if the customer’s network to which the enterprise network is connected has the same standard of physical access controls as your own, their physical access is under *their* control, not *yours*.
- *Lack of logical control.* Does the trading partner’s organization to which your network is connected have the same standard of logical protection: intrusion detection, firewalls, remote access, anti-virus, access control, etc.? Whether it has the same standards or not, logical control of the trading partner’s enterprise is under *their* control, not *yours*.
- *Lack of organizational control.* When an employee is the subject of misconduct, the traditional organizational command and control structure deals with the matter through disciplinary action: the employee can be reprimanded, reassigned or terminated. But how does one control someone in the trading partner’s organization? How will you even know when an employee in the trading partner’s organization is acting inappropriately?
- *Lack of segregation within applications.* A common method for organizations to give partners, suppliers or customers access to relevant information is to issue userids to individuals in these outside organizations. The intention is to give individuals in these partner organizations access to *their* data. The problem with this approach is that most legacy enterprise applications were not designed with this sort of functional segregation in mind: many (if not most) enterprise applications were designed to give a user access to all information. Artificially partitioning data so that a user or group of users has access to some subset of information is a function that many enterprise applications lack today. However, building in this functionality may be cost prohibitive. Scarce resources and other issues make customization on this level infeasible.
- *Business operations integrity.* The supplier to which your enterprise is connected may have similar connections to other large customers, some of which could be your competitors. Such suppliers could find themselves in the position of having insider information into two or more competitors’ operations. Will these suppliers practise restraint and act appropriately? Can they make sound business decisions with this insider knowledge? Can they and will they act with integrity?

While painting a bleak picture is not this section’s intention, one must take a sobering and realistic view of the implications of connecting one’s enterprise network to that of a trading partner.

## Security is about people, not technology

Chapter 5 explored the concept of information security ultimately depending upon people's behaviour: in most contexts, technology can do little – if anything – to stop an insider who is intent on violating security policy. This theme of security-depends-upon-people applies to trading partners as well. And arguably, the employees of a trading partner who are intent on committing an unscrupulous act, whether it be the theft, inappropriate disclosure or destruction of information, have fewer constraints holding them back:

- *Absence of loyalty.* Employees of a trading partner, if loyal to anyone, are most likely going to be loyal to *their* employer and not to *your* organization. They may very likely not know much about or care about your business.
- *Out of sight, out of mind.* Away from the prying eyes of your organization, they may succumb to the temptation of misusing your information. The data they can access may not be vital to their lifestyle – it belongs to someone else and may be of little import to them.
- *No fear.* They may know enough to know that they can do many things to your data without leaving tracks. Frequently – especially in older applications – many compromises need to take place in order to pound a 'square peg' intranet-centric application into an extranet-centric 'round hole' environment. The lack of adequate controls or audit trails may mean that their misdeeds may not be associated with them or even their employer.

While trading partner relationships have been a mainstay of business for centuries, the advent of extranet communications to link one's suppliers, partners and customers to one's intranet applications and data is an unprecedented step into untested waters. Technologies are just beginning to catch up, and people's understanding and behaviour have even more catching up to do.

## REGAINING CONTROL

The preceding section probably has some readers wondering why and how extranets ever got their start – or if all of the early adopters did so out of ignorance or foolishness. While certainly some early adopters embraced extranets because they were 'the next big thing' or merely because they were cool, some realized there were efficiencies to be gained by tightening up one's supply or value chain through increased automation: Electronic Data Interchange (EDI) just wasn't cutting it any more. The extranet visionaries realized that the risks were surmountable if properly calculated and managed – that the gains did outweigh the risks.

There are no magic cures or works-for-everyone formulas that can eliminate the risks associated with trading partner extranets. Some of the reasons for this are explained here.

## Reducing risk

An organization can take several measures to mitigate the risks outlined in the first half of this chapter. None of these will eliminate risk, however, and some may not be relevant to every situation. This may be thought of as a starting point, with other steps that may need to be followed as business conditions dictate.

### *Legal contracts*

Every trading partner extranet relationship must be defined and enforced through a legal contract. From the information security perspective, a contract should describe the following:

- the precise nature of any electronic connection between the organizations – this should include a high-level description of the business activities that are supported by any connection(s)
- a description of the functions that the trading partner may perform with your data
- a statement that says that there are no other permitted uses for your data, including both individual records and aggregate or derived data
- a statement that specifies where, in what form and for how long the trading partner may store your data
- a description of the kinds of audit events that the trading partner must create and maintain, and clear definitions stating which kinds of events *are* audit events
- the contract should enumerate or refer to security policies that the trading partner must adhere to on any systems or networks that process your data; a reference to an external (outside of the legal contract, that is) security policy might be preferable, as best practices and changing conditions may require changes to security policies during the lifetime of the legal contract (this will be described in more detail under ‘Audits’ below)
- the contract must include a ‘right to audit’ clause that permits you to perform or order an audit on the trading partner’s information systems and practices, should undesirable practices be suspected
- the contract must contain financial penalties that will be imposed upon the trading partner if it is found to be violating any part of the contract.

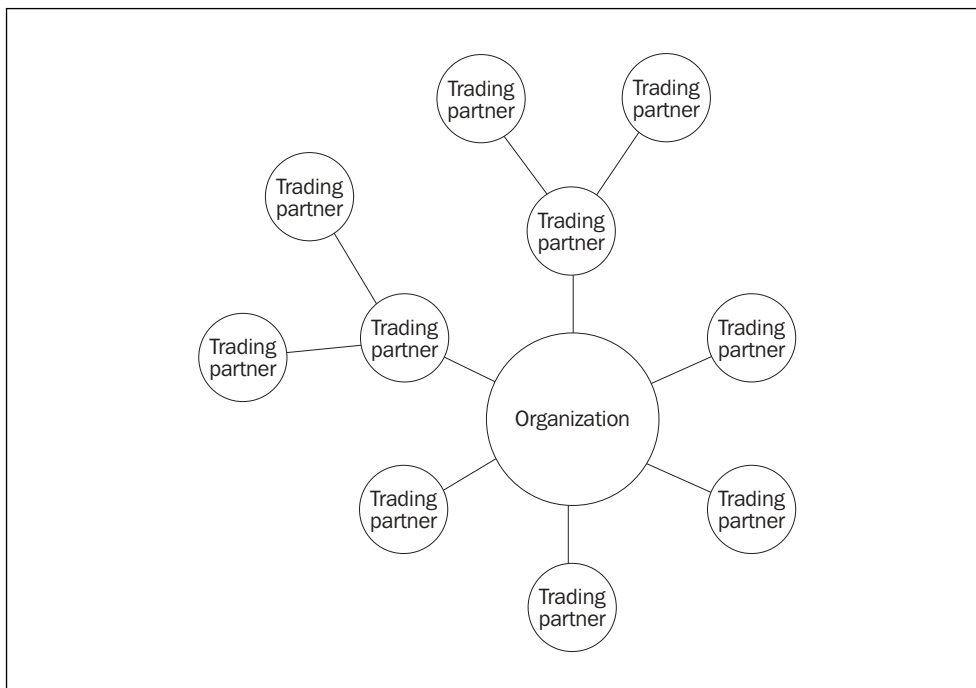
### **Policies and requirements**

An organization considering a trading partner extranet connection must impose upon that trading partner a list of minimum security policies and requirements that will ensure the highest possible protection for the organization's information assets and infrastructure. This is the one category where poor information hygiene will become painfully evident sooner or later.

Generally speaking, the organization should assert on to its trading partner its own set of high-level security policies and requirements. However, it would probably be unreasonable for the organization to require that the trading partner's *entire* enterprise conform to these policies and requirements. A compromise might be for the trading partner to provide an acceptable degree of physical or logical segregation between its internal systems and the systems that manage your data. For instance, if a particular trading partner has little in the way of security practices, it might be reasonable to require that the system(s) that stores and processes your data resides in a locked room with no network connection to the rest of the trading partner's enterprise.

Likewise, should the trading partner participate in other extranet environments, you are going to have to see how the trading partner protects itself from trading partners. This means finding out what policies and requirements each trading partner asserts upon *its* trading partners. Figure 6.1 illustrates this concept.

**Fig. 6.1** The trading partner web



This point cannot be overemphasized, because the actions of your trading partners' trading partners is *completely* out of your control. While one's own security mechanisms constitute the first line of defence, those in the trading partners' networks form an additional ring of protection (however, the same trading partners' networks also present an additional layer of risks).

### **Access to information**

For the purposes of this book it is presumed that a trading partner requires access to some portion of the organization's information in order to fulfil whatever task the trading partner is supposed to perform. The trading partner should have access to *only* the information it requires to perform its assigned function. While this may be dreadfully obvious, it might not be all that easy to implement. Customization of reports, query screens and other user interfaces can be expensive, and the controls beneath it all – at the database level – might be more expensive still.

In an ideal situation, the trading partner will have verifiable access to only certain databases, tables, fields and records – and all can be neatly tied back to required functions that are described in legal contracts. But situations are seldom ideal and this leads to the need for additional measurements and controls.

### **Measurements and controls**

The business functions that the trading partner is performing need to be measured. This is the most difficult need to describe because of the wide variety of possible functions that a trading partner might be performing. Some examples will be described here.

- *Order processing and payment acceptance.* A third-party company processes orders and accepts payments for the organization for products sold online. This third party performs similar services for other organizations, including competitors. This kind of arrangement requires a full settlement relationship where the organization and the trading partner must perform daily transaction reconciliation and balance all orders, payments, refunds, credits and other transactions. The organization must ensure that the trading partner is not accepting payments from (or even originating) uncollectable sources such as credit card fraud, and that it is not producing shipping orders for which no payments have been received. These and several other fraudulent activities must be measured and controlled.
- *Customer support centre.* A third party performs telephone support functions for a software product vendor. The support services, which the customer pays for through annual contracts, consist of telephone support and occasionally sending out new software media or bug fixes. In this environment the organization wants to ensure that the third-party support centre is not giving out free support to

customers who have not paid for it. So again, the trading partner and the organization must utilize a system of recording the identities of callers and have some way of auditing these calls to ensure that the trading partner is not synthesizing call records in order to profit somehow from the resources entrusted to it.

- *Warranty exchange fulfilment.* A third party performs product exchange services by taking telephone support calls, working with the customer to understand the nature of the product difficulty and, when certain criteria are met, by shipping a replacement product to the customer. In this example, the trading partner has limited access to the organization's customer database so that it can verify account and product-type information. Like the other examples here, this trading partner has been entrusted to manage valuable organization resources. Details of all transactions must be tracked in order to detect inappropriate distribution of company products.

Not only must the organization protect itself against erroneous and fraudulent actions on the part of its trading partner, but the trading partner itself also may require means for protecting it and the organization against erroneous and fraudulent actions on the part of customers and other parties.

### **Audits**

The organization must obtain a guaranteed right to audit all relevant trading partner transactions. Without this right, a trading partner cannot be held accountable for its actions and the organization will not be able to confirm whether the trading partner is performing its functions correctly.

Everything that the trading partner does on behalf of the organization may be an audit event. Consider the example where the trading partner has access to the organization's customer records. Arguably, even the act of viewing records should be an audit event. Consider the upward trend in identity theft and the fact that basic customer information is a readily saleable commodity. Trading partners must be accountable for even viewing the organization's information – not just customer information but anything they need to perform their job.

Audit records must reflect the actions of identifiable individuals. People must be held accountable for their actions. This is a powerful deterrent against unscrupulous acts, but also a useful means for assisting investigations. A strong and irrefutable association of audit events to individuals requires strong user access policies and procedures: the association of a person's name on an audit record is only as strong as the entire process of providing and maintaining user accounts. But if users fail to conform to security policies and share userids and passwords, then of course the audit data is no good and the entire operation is in real trouble.

## SUMMARY

In this writer's view, almost everything in the world is a 'glass half full', so why does this chapter paint such an ominous picture of extranets? The reason behind this mood is all security related. Bruce Schneier, one of the world's leading cryptologists and author of *Secrets and Lies* and *Applied Cryptography*, recently went on record to say that we are building information systems and infrastructures faster than we can work out how to secure them. So arguably, if we are only beginning to work out how to secure our information assets and infrastructures, why is it that we now want to throw open our doors and let in our trading partners? Of course, we trust them and need them to facilitate efficient and responsive business, but if we can't keep the ruffians off our own payrolls (where we ought to have some control), what risks are we taking on by bringing in whole communities of users that we don't even know or have any control over?

So why build extranets at all? This writer's optimistic side says that the answer is 'because we can', but also because extranets hold the promise of delivering a more efficient and responsive supply/value chain, driving down costs and driving up profits and market share. But extranets are not a light undertaking: any but the most trivial and minimal extranet project will shake the business to the core, and few stakeholders will be unshaken and unmoved.

Extranets are fraught with risks at every level – including network connectivity, access to corporate data and business processes – and there is no way to build sufficient technical controls to protect everything. At every step of the way, there are opportunities for things to go wrong – some obvious, others less so. The core issues may be an enterprise network with a strong perimeter defence (which is tunnelled through in the case of an extranet) but weak internal controls, or perhaps a legacy application that cannot be taught to hide some records, fields or functions from some users. It could also be a lack of good controls to stave off fraud that could drain any hope of the organization's viability, much less profitability.

Audit records are only going to be as good as the processes and technologies supporting user accounts – otherwise the name associated with any audit event might as well not be there. And finally, the overall health of the enterprise network is going to be no better than that of the trading partner, or even the trading partner's trading partner(s).

The lawyers will earn their money with extranet contracts. Without the leverage to audit the trading partner and hold them accountable for their conduct, the extranet can be a fine way to send the organization into a tailspin.

# Privacy

- Introduction 115
- What is personal information? 115
- It's all about trust 117
- Privacy policy 118
- How security supports privacy 120
- Privacy certifications 121
- Summary 122



## **INTRODUCTION**

People by nature do not wish to be intruded upon. The convenience of the Internet brings with it at least one downside – in order to perform online transactions, one has to give up some privacy by sharing information with the organization with which one is doing business. The issue with privacy is that so many organizations that collect personal or private information seem to share it with other organizations against our will and without our consent.

Privacy – how an organization collects, handles and shares private information – is becoming an increasingly visible issue in the market. Some companies are making their privacy policy a competitive differentiator, and consumers are beginning to take notice and tailor their shopping preferences accordingly.

This chapter will approach privacy by first describing what personal information is. This is followed by discussions on identity theft, trust, privacy policy and P3P, and then the dependency of privacy on security policy. The chapter concludes with a discussion of privacy and security certifications.

## **WHAT IS PERSONAL INFORMATION?**

Personal information is information that is specifically associated with an individual. Some examples of this personal information are:

- full name
- date of birth
- tax identification number
- retirement or social insurance identification number
- place of employment
- driver's licence number
- residence address
- e-mail address
- telephone number
- bank account number(s)
- credit card number(s)
- mother's maiden name
- names of other family members, especially children.

This information is collectively known as personally identifiable information, or PII.

While customs, cultures and individual preferences introduce some variability on this list, most individuals are particular about sharing their PII. According to

the Center for Democracy and Technology in the USA ([www.cdt.org](http://www.cdt.org)), ‘consumers are still shocked to learn that information about their activities ranging from online browsing to grocery shopping is used for a variety of purposes and made available to other companies without their permission’. Organizations that collect personal information from customers need to consider carefully what information they are gathering and why they are gathering it.

## **Identity theft**

People are becoming painfully aware of the exponential rise in the abuse of personal information. The number of complaints in the USA to credit reporting agencies and various branches of the federal government including the Federal Trade Commission, the Department of Justice and the Social Security Administration are in the hundreds of thousands each year, ranging from complaints about the abuse of individuals’ PII to identity theft.

Identity theft, in case the few who don’t yet know about it are readers of this book, occurs when a perpetrator is able to gather enough of an individual’s PII that they can begin to apply for credit in the individual’s name and commit other abuses for their personal gain. The victim of identity theft usually finds his or her credit ruined, as the perpetrator has ‘run up’ one or more credit cards with no intention of paying the bills.

So it should be no wonder that individuals may be a bit shy when asked to share a piece of PII with a corporation. They have good reason to be genuinely afraid that their information could be misused and result in their identity being stolen.

## **PII a commodity**

For all the expense, grief and hassle that identity theft causes, the street value for the PII on an individual is as low as US\$50.

It is not difficult to obtain this kind of information. Customer care representatives in large companies have easy access to customer information. Care reps work in high-stress environments: they work in relatively noisy, crowded rooms, they are verbally abused by frustrated customers, and they work long hours for low pay. Their working conditions cause many to feel resentment towards their employer, making it relatively easy for small circles of care reps to collaborate in conspiracies to collect and sell customer information for profit. Many care reps, frustrated by their working conditions, will feel a sense of power and revenge for having sold some of their employer’s confidential information to an outsider. Everyone has a price, including those who are the stewards of our private information.

## **IT'S ALL ABOUT TRUST**

The information revolution, the 'new' or 'digital' economy, or whatever you would like to call it, has two noteworthy characteristics:

- individuals are purchasing goods and services online, and through this and other activities are sharing personal information in order to complete these transactions
- databases on genealogy and birth records are proliferating, giving researchers an unprecedented view into many individuals' personal history.

Personal information is flooding the Internet. Credit card numbers, birth dates, names, addresses and phone numbers are showing up on more and more computer systems. Not all of them are secure. There have been several spectacular hacking incidents whereby hackers have broken into e-commerce, banking and healthcare databases and carted off sometimes hundreds of thousands of credit card numbers, tax identification numbers and other PII. In many cases this information is sold to identity theft rings or even posted to websites.

### **Misuse of personal information**

According to a February 2002 Harris Interactive survey, 75 per cent of US consumers have a major concern with providing their personal information to a company. They are afraid that the company will supply their private information to another company without their permission.

The proliferation of spam (unwanted junk e-mail, usually a business solicitation) is another demonstration of how quickly and easily a person's e-mail address will propagate from one organization to the next. Many consumers believe that their e-mail address, given in implied confidence to a website, was sold or distributed to other Internet operators who send spam. This feeling of betrayal is making individuals think twice about sharing even more valuable and sensitive information such as a credit card number on the Internet, especially with a 'no name' online business.

### **Theft of personal information**

Consumers are not just afraid that unscrupulous companies will deliberately sell their personal information. They also fear that their information will be stolen from sites that have done a less-than-thorough job of securing their information. The same Harris Interactive survey cited earlier, states that 69 per cent of consumers fear

that hackers will steal the personal data that they submit online. This perception can have a chilling effect on e-commerce if not mitigated in some way.

## **PRIVACY POLICY**

Most corporations that do business with consumers – whether they conduct their business online or not – have developed and published a privacy policy. Businesses' primary motivations for developing a privacy policy include:

- *Accountability* – organizations are coming under increasing pressure to be publicly accountable for the information they collect from individuals.
- *Statutory compliance* – in many jurisdictions throughout the world, laws are being enacted that require businesses and governments to collect PII only under certain conditions, protect that information and share it only under certain circumstances.

The publication of a privacy policy is quickly moving from legal disclaimer status to a statement of competitive advantage. The manner in which an organization collects and manages private information has become a competitive differentiator. Over time these concerns will have greater impact on organizations as they begin to change their business processes and practices – and eventually even their business models – in order to retain competitive advantage.

The basics of a privacy policy describe the following:

- *Why private information is collected.* The organization needs to state the purpose behind the collection of private information. Is it collected in order to provide services to the customer? Is it used for market research? Does the organization sell or trade the information?
- *What private information is collected.* Which items specifically are collected?
- *How private information is collected.* Is information obtained from the customer directly and, if so, how? Is information obtained from other sources (such as credit bureaus, the government, banks, etc.)?
- *How private information is used internally.* What does the organization do with the data once it has obtained it?
- *Who in the organization can see private information.* Which individuals or job functions have the ability to see this information? Is viewing private information an audit event?
- *How private information is protected.* What measures are taken to ensure that the customers' private information is not compromised or disclosed? How does the organization ensure that the information is not disclosed to the wrong customer? This would include access control, encryption and perimeter defences.

- *How the private information is disclosed to others.* Under what conditions and for what purposes is private information shared with other organizations?
- *Where the customer can view the private information to make sure it is accurate and up to date.*
- *Who can answer the customer's questions.* If the customer has comments or questions about the private information that the organization may have on file, whom should they contact and how should they contact them?
- *Where the customer can view the privacy policy.*
- *Whether the privacy policy will change.* Under what conditions will the privacy policy change? Will the customer be informed and, if so, how?

The organization's privacy policy is a commitment to its customers. At the very least it is a moral contract between organization and customer, and in some countries it is considered a legal contract.

## Privacy policy and P3P

P3P, or Platform for Privacy Preferences Project, is a standard that was developed by the World Wide Web Consortium (W3C) to permit users to gain more control over the use of personal information on websites they visit.

### **How it works**

P3P requires two components to function: a browser that supports P3P and a website that supports P3P. In a web browser that supports P3P (e.g. Microsoft IE version 6 and above), users set their privacy preferences, either by selecting a general setting (e.g. low, medium, high), or by changing preferences at a more detailed level. A website that supports P3P performs a similar act: at the direction of the person(s) responsible for privacy policy, the webmaster configures the website's P3P settings to match the organization's privacy policies.

When a user with a P3P-enabled browser visits a P3P-enabled website, the browser queries the site for its P3P policies, downloads them and compares them to the user's P3P privacy preferences. If the website's attributes are equal or better than the user's preferences, nothing happens. However, if any of the website's P3P settings do not meet the user's preferences, the user will see a dialogue box pointing out the fact that the site they are visiting violates the user's preferences. The user can then choose to ignore this and proceed to browse the site or they can go elsewhere.

### **Implications**

Many website operators are upset over P3P, for a number of reasons. P3P lacks sufficient granularity to be of value in some situations and it exposes practices that

some website operators do not want their customers to know about. A good example of the granularity issue has to do with the organization that shares personal information with other organizations. There are many reasons why an organization would share personal information with another:

- *Credit check.* In order to sign up a customer with a service, the site might do a credit check to assess the customer's creditworthiness.
- *Catalogue and payments.* A site may have outsourced the catalogue and shopping cart portion of its site to an organization that specializes in this activity.
- *Other outsourcing.* An organization may have outsourced any of several other functions, such as customer support, collections and payments.
- *Market research.* An organization may share its customer information with a company that performs market research on the organization's behalf.
- *Selling personal information.* The site may be overtly or covertly selling personal information to other organizations.

The trouble with P3P is that all these information-sharing activities fall under one umbrella: does the organization share its information with outside companies or not? This issue has a lot of organizations concerned. Larger companies that legitimately outsource some of their functions are placed in the same category as sites that sell information to anyone and everyone for cash. The legitimate uses of information sharing – most of which are done in strict confidence – may be maligned by the general public who may think that all information sharing is bad.

Finally, P3P is still in its infancy and it is not yet known whether the majority of consumers will pay attention to the P3P settings in their browsers or whether they will change to the lowest settings and ignore any warning messages. For this reason, many organizations are still sitting on the sidelines and waiting to see whether consumers pay attention to P3P or not.

## **HOW SECURITY SUPPORTS PRIVACY**

The organization that collects and stores PII on its employees, contractors or customers needs to determine how to protect it. Ideally the organization will make these information protection decisions top-down, beginning with security guiding principles and security policies. The policies that an organization needs to develop fall into these general categories:

- *Classification of PII.* There are two or three levels of sensitivity of information collected on an individual. Items such as tax identification number, credit card numbers, date of birth and probably a few others would fall into the category of the most sensitive data. Other items such as name, address and telephone number

might be categorized as less sensitive. This classification is extremely important because information-handling policies for these categories will determine how the organization uses the information. Local laws may also influence this classification process.

- *Storage of PII.* The organization must determine what safeguards will protect stored PII. Will the most sensitive data be encrypted or stored only on systems that have an extra degree of hardening? Will employees be permitted to store any PII on their personal computers? Will a data warehouse or decision support system store credit card numbers or tax identification numbers in their entirety, with just the last few digits or nothing at all?
- *Transmission of PII.* The organization must also determine how it will transmit PII. Will any or all PII be encrypted when sent over the Internet? Will it be encrypted when transmitted within the organization's local private network?
- *Viewing of PII.* The organization must decide who can view PII. Which employees will be able to view what fields? Will viewing PII constitute an audit event that must be logged? Will any employee be able to view an entire credit card number or tax identification number? If not, the organization must decide who can view an entire credit card number, for example, and who can see only the last few digits.
- *Sharing PII with external organizations.* Under what circumstances will the organization share any PII with an outside company? How will the security of this PII, once it leaves the confines of the organization, be protected from inappropriate distribution or disclosure?
- *Aggregate information.* If the organization creates aggregate information on its employees or customers (for example, average income or age), how will this information be protected, stored, viewed and shared?
- *Alignment with corporate privacy policy.* Do the security policies governing the use of PII align with and support the stated privacy policy?

Policies on all of these topics will drive enterprise business processes, system and application architectures, and data flow models. They will translate into security requirements that will influence the implementation details of internally developed, purchased, outsourced and managed enterprise services.

## **PRIVACY CERTIFICATIONS**

A study conducted by Harris Interactive (February 2002) states that 91 per cent of US consumers say they would be more likely to do business with a company that verified its privacy practices with a third party. The study found that 62 per cent said

third-party security verification would allow them to be satisfied with the company and 84 per cent thought that third-party verification should be a requirement.

P3P does not constitute third-party verification. Since the website operator performs its own P3P configuration to match its privacy policy, there is no objectivity or verification. An unscrupulous website operator could create a written privacy policy that says one thing while setting its P3P configuration to indicate something different. The website operator may not even be *legally* obliged to configure its P3P server accurately, but that does not make it immune to civil litigation or simply to the loss of its customers.

There are several options available to the organization that wishes to obtain a third-party endorsement of its security and privacy. A few of these are described here.

- *TRUSTe*. Pronounced like ‘trustee’, this is a non-profit organization that has a well-known certification programme. In order to earn its ‘privacy seal’, an organization must subject itself to an audit by TRUSTe to verify its information-handling policies regarding customers’ private information.
- *Webtrust*. This is a certification that attests to the organization’s compliance with stringent principles and criteria developed for different types of businesses. Webtrust was developed by the American Institute of Certified Public Accountants (AICPA).
- *Systrust*. This is a professional service that provides an attestation to the availability, security, integrity and maintainability of an information system. This service was developed by AICPA and the Canadian Institute of Chartered Accountants (CICA).
- *SAS70*. Also known as the Statement on Auditing Standards No. 70, this is an internationally recognized auditing standard developed by AICPA. It allows organizations to disclose their control activities and processes to their customers and their customers’ auditors in a uniform reporting format.

The presence of one or more of these attestations can help the organization in two distinct ways. First, it gives potential customers and business partners an assurance that the organization has subjected itself to and successfully obtained an objective statement of security and integrity. Second, it can help stave off independent audits required by each potential large customer or business partner.

## **SUMMARY**

Privacy is freedom from unauthorized intrusion. In the digital world, this translates into the growing concern that people’s personal and private information is being propagated to other companies without their approval and that it is not adequately protected, which is resulting in hacking attacks that violate their privacy.

PII, or personally identifiable information, is the industry term that describes the pieces of information that are used to identify an individual. Some of these include date of birth, tax and social insurance identification numbers, credit card numbers, driver's licence number, etc. The epidemic of identity theft arises directly from the fact that many organizations have some or all of these items of information stored on systems that lack adequate – or even basic – protection.

Consumers are concerned about whether a company with which they share personal or private information will pass it on to other organizations without their authorization. The proliferation of spam (unwanted junk e-mail) makes many people wonder whether a website they recently registered with has sold their e-mail address to spammers. While spam is a mere annoyance, people are genuinely concerned about identity theft, and many balk at giving credit card numbers and other information to websites. This fear is justified, given that in the USA alone there are hundreds of thousands of cases of identity theft each year.

The starting place for protecting personal and private information is security policy and privacy policy. These two policies define the basic rules for the collection, care and protection of private information. Security policies drive the technical architectures and business processes that actually handle and protect private information.

P3P, or Platform for Privacy Preferences Project, is an emerging standard that is used to allow consumers to gauge electronically a company's privacy policy and steer away from it if its policies do not meet the individual's privacy needs.

Security policy supports privacy by defining classes of customer data and the handling requirements for this data. This will determine how an organization stores, processes, transmits and views customer personal information.

There are a variety of security and privacy 'seals' that organizations obtain in order to give consumers and trading partners more confidence that the enterprise is properly collecting and processing private information. These certifications can also help the organization to stave off individual and time-consuming audits that may otherwise be required by trading partners.



## Action items

- Most important and urgent action items (Quadrant I) 128
- Most important but less urgent action items (Quadrant II) 130
- Important and urgent action items (Quadrant III) 132
- Important and less urgent action items (Quadrant IV) 133
- Epilogue 135



Presuming that you have read most or all of the previous chapters in this book, you are probably wondering why I am including a chapter called ‘Action items’ when it appears to be painfully obvious what you need to do next. But it’s *not* all that obvious to most readers which things should be done first, which ones next, and which ones not at all. I can’t tell you the answers because they’re different in every organization, and if I did, you wouldn’t believe me anyway.

To benefit the greatest number of readers of this book, it will be assumed that your organization is in terrible shape: there is no perception that security consists of any more than a firewall and anti-virus software, and security plays no strategic role (and barely a tactical role).

The work to be done is organized into a Steven Covey-like (from *The Seven Habits of Highly Successful People*, Econo-Clad Books, 1990) Important-Urgent quadrant, with these labels on the quadrants:

- Quadrant I: Most Important and Urgent
- Quadrant II: Most Important but Less Urgent
- Quadrant III: Important and Urgent
- Quadrant IV: Important and Less Urgent.

Figure 8.1 shows these quadrants with several items in each quadrant. The items listed are somewhat arbitrary and may not fit your organization’s priorities and needs.

**Fig. 8.1** Enterprise security action item quadrants

		<i>Urgent</i>	<i>Less Urgent</i>
<i>Most Important</i>	<i>Quadrant I – Most Important and Urgent</i>	Roles and responsibilities, short-term Firewalls Anti-virus Security patches Initial assessments Basic security awareness Emergency contact lists Change default passwords	<i>Quadrant II – Most Important, Less Urgent</i>
	<i>Quadrant III – Important and Urgent</i>	Automated anti-virus updates Security logs	<i>Quadrant IV – Important and Less Urgent</i>
<i>Important</i>			Roles and responsibilities, long-term Security architecture and strategy Configuration management Security requirements Formal security awareness

The remainder of this chapter assumes that your organization's sense of importance and urgency for each of the issues here exactly matches how they are placed in the quadrants. Of course each reader, knowing his or her enterprise's state of affairs, will have a different idea of what is urgent and not, and what is important and not.

## **MOST IMPORTANT AND URGENT ACTION ITEMS** **(QUADRANT I)**

The most urgent and important things to do will ensure that the enterprise can be made as secure as possible from outside threats *right now*.

### ***Roles and responsibilities, short-term***

Right now, you need to find out who in the organization is responsible for protecting the enterprise information assets and infrastructure. Is this person or group responsible for security strategy? Security operations? Security administration?

Once you know *who* is responsible for protecting information assets and infrastructure, ask them *how* they are protecting them – not just with what technologies but with what processes as well. Is there a change management board that approves security changes? Is there a configuration management process that is capturing the details of each change?

Details are not needed for these questions at this time – general answers will do.

### ***Firewalls***

You need to determine whether firewalls are protecting the enterprise from external (trading partners, Internet) networks. Do *all* ingress points have firewalls? Are they configured correctly? Is a single person or group responsible for configuring all firewalls or are different groups responsible for administration on various firewalls? If the latter, it is unlikely that any single person or group understands the 'big picture' of protecting the enterprise from the kinds of external threats that firewalls are designed to handle.

Is there an approval process governing firewall rule changes? Are the details of changes captured in a configuration management process? (The configuration management process itself is covered in Quadrant IV.)

### ***Anti-virus***

Are all desktops and servers protected with up-to-date anti-virus software and definition files? Do not worry too much about the details on how desktops and servers are updated: whether they are installed and up to date is what matters today.

Unless the enterprise has an automated virus management system that keeps track of which systems have anti-virus installed (and which versions of definition files), this is a labour-intensive task. But chances are this work has already been done, unless your organization is either very lucky or not connected to the Internet. This is because NIMDA and Code Red pretty well exposed nearly every organization's unprotected systems.

### ***Security patches***

Are all of the required security patches installed on all servers? Are any security patches installed on clients' systems? Is there any recordkeeping regarding which patches are installed on which systems? Is there a testing or change management process that governs which patches are installed when?

### ***Initial assessments***

It may be a good idea to get a competent IT or EDP (electronic data processing) auditor to give your systems a thorough going-over, especially on servers that host Internet services and servers used by large numbers of users (or very critical users). It is important to determine who will manage and pay the auditors so that there is no conflict of interest or collusion. Every auditor's results must be objective – this is especially important now.

### ***Basic security awareness***

How aware of information security basics are most people in the organization? Are they aware of basic security practices such as using good passwords and not sharing each other's accounts?

### ***Emergency contact lists***

This item assumes total *un*preparedness in the event of a security emergency, natural or man-made disaster. At the very least you need to develop a complete catalogue of the organization's critical personnel, vendors, suppliers, utilities, etc. and distribute hardcopies to the critical personnel and others as needed. Copies should be stored in one or more off-site locations as well.

This is a pitiful excuse for a business contingency plan, but it's just slightly better than nothing.

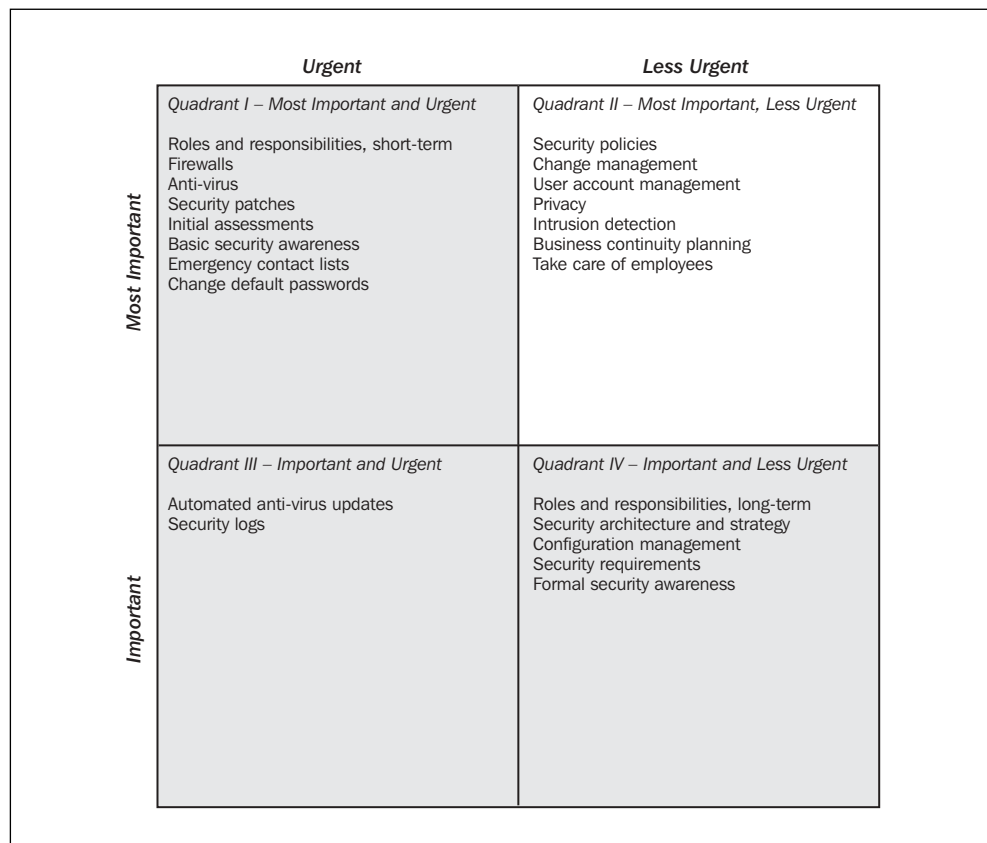
### ***Change default passwords***

All accounts that are preconfigured with default passwords must have those passwords changed. Otherwise an intruder with knowledge of those default passwords can easily break into the system.

## MOST IMPORTANT BUT LESS URGENT ACTION ITEMS (QUADRANT II)

You are fortunate to be allocating resources in this quadrant (Figure 8.2) for this means that the urgent items are under control. The items in this quadrant are very important and pay long-term dividends.

**Fig. 8.2** Quadrant II



### ***Security policies***

The organization's security policy must be examined. Is it up to date, relevant, accessible, comprehensive and comprehensible? Are employees required to read and understand it? Is it enforced?

### ***Change management***

A significant portion of information systems quality problems are rooted in key processes such as change management. Are changes made to your production environment subject to a review and approval process? Does the process (if it exists) work properly and meet the organization's quality and security needs?

### ***User account management***

How are user accounts on your information systems managed? Is there a central UAM function or do system and network administrators build userids on demand? Is there a process for deactivating terminated employees' accounts? Do user accounts expire after an interval of non-use?

Userids and passwords are the first line of defence protecting information assets. If userids and passwords are unmanaged, the opportunities for abuse will skyrocket.

### ***Privacy***

Does your organization collect any information from customers, such as birth dates, tax identification numbers or credit card numbers? If so, does your organization have a privacy policy?

The proliferation of organizations building customer databases containing private information and selling or sharing these databases has brought near hysteria to many consumers, so much so that slow, conservative governments are taking notice and passing privacy laws and regulations.

If your organization collects private information and does not have a privacy policy, this is a Quadrant I (Most Important and Urgent) action item.

### ***Intrusion detection***

Since you are in Quadrant II, this means you have firewalls placed at all of your network ingresses. However, without intrusion detection systems (IDSs), you may never know (at least until it's much too late) whether someone or something has penetrated the firewall.

Remember that insiders are capable of malicious behaviour, too, and increasingly organizations are placing IDSs well inside the perimeter in order to detect anomalous behaviour that may be originating within the enterprise network.

### ***Business continuity planning***

How well prepared is the organization for the kinds of unexpected events that will threaten the very existence of the organization? Business continuity planning (BCP) defines the steps to be taken to ensure that the most critical business functions continue to operate despite any event that may have interrupted them.

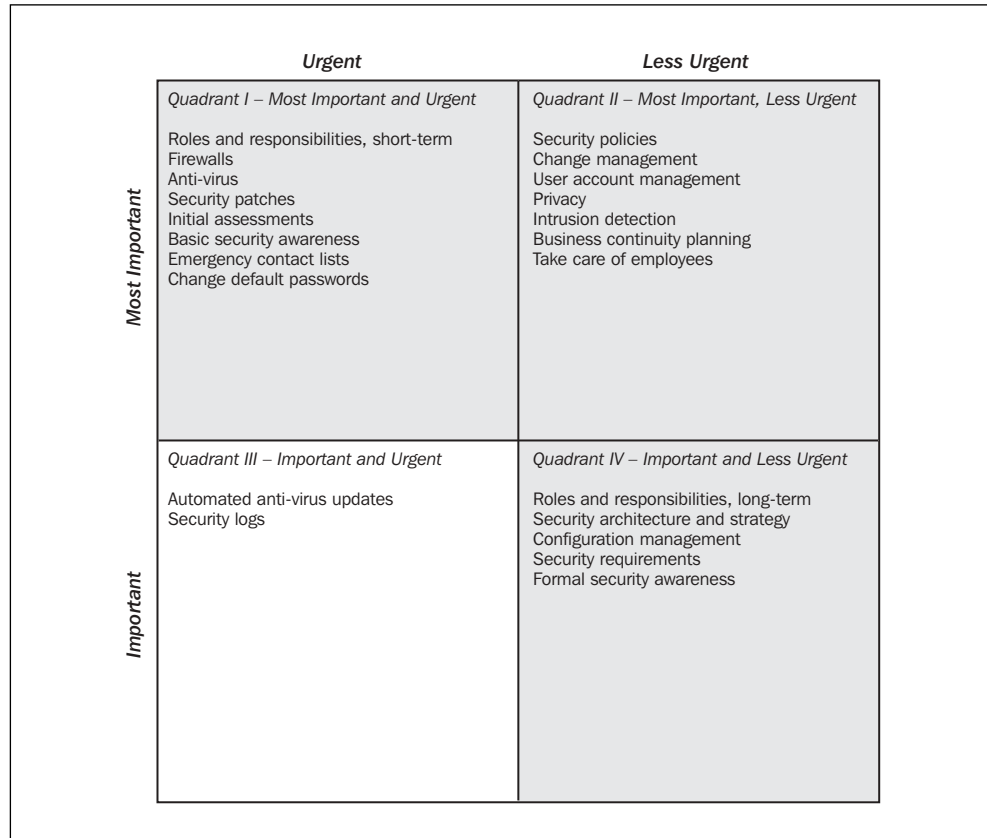
### ***Take care of employees***

It is truly amazing that, even in this era, some employers don't show their employees how much they are valued. Remember two things: the majority of security incidents are inside jobs – you know who the perpetrators are. And employees, being insiders, have the *means* and plenty of *opportunities*. Take care not to hand them a *motive*.

## IMPORTANT AND URGENT ACTION ITEMS (QUADRANT III)

If, after handing out assignments to cover everything in Quadrant I, you still have available resources (ha!), cover these items next (Quadrant III).

**Fig. 8.3** Quadrant III



### ***Automated anti-virus updates***

The latest versions of the most popular anti-virus products (e.g. Norton, McAfee) have automated definition file update mechanisms that can relieve users of the task of updating definition files manually. Further, they may be able to configure your anti-virus products to ‘push’ new definition files out to the enterprise in real-time – very handy in an emergency.

In newer Windows-based environments (if you’re running Macintosh, UNIX, Linux or green screens, you probably aren’t too worried about viruses), you’ll be able to force desktop systems to reinstall anti-virus software if it is old or missing, and further you will be able to prevent users from removing or disabling anti-virus software. Getting these capabilities set up will enable those personnel to get out of the urgent quadrants (or help others still stuck there).

## Security logs

Where are the security logs? In a large enterprise, chances are they are in hundreds of places with no correlation or aggregation functions. System, audit and security log correlation and aggregation capabilities are still in their infancy, and given typical IT resource constraints, chances are few know where the logs are and seldom (if ever) are they examined.

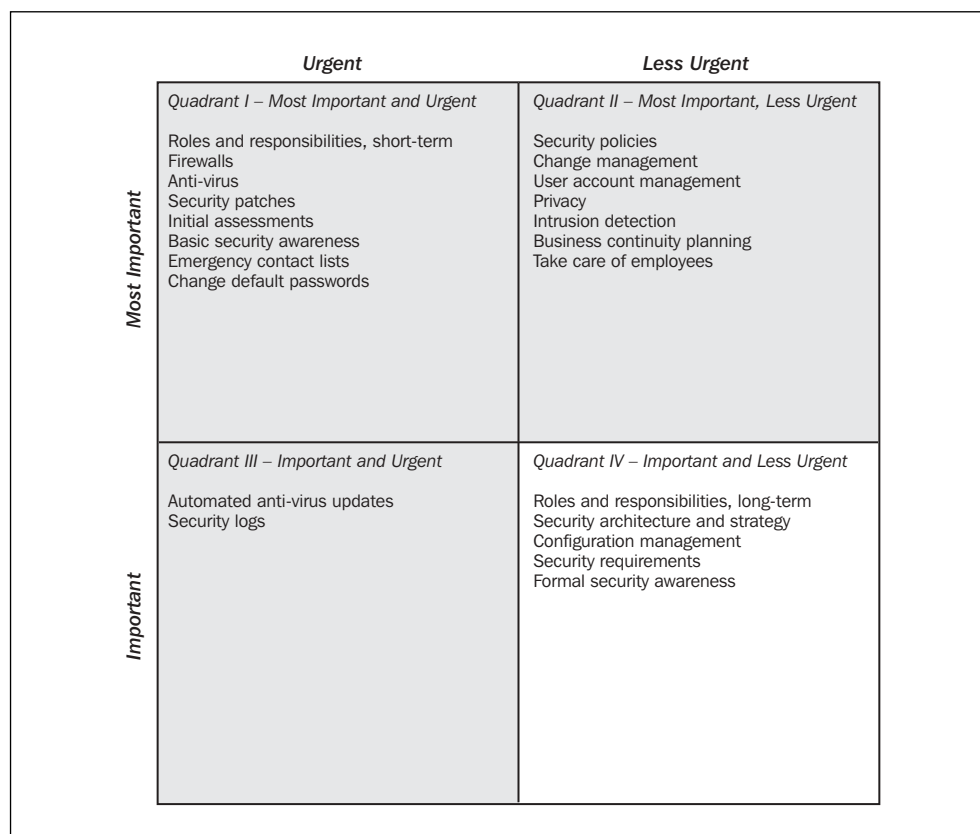
This can become a ‘boil the oceans’ (or world hunger) problem if not properly sized and managed. For now, it is probably enough for the most important security logs to be identified, and if there is a resourceful systems engineer available, an out-paging capability can be built to send the most urgent events to pagers or message-enabled cell phones.

## IMPORTANT AND LESS URGENT ACTION ITEMS

### (QUADRANT IV)

In this quadrant you will be able to work on activities that will bring long-term value to the organization. The strategic tasks build repeatable processes and direction that will make security integral to the organization.

**Fig. 8.4** Quadrant IV



### ***Roles and responsibilities, long-term***

In some organizations, security is concentrated in an information security group, while in others information security is a cross-functional responsibility with practitioners from several groups performing various tactical duties. If information security is cross-functional, are the various security people/groups organized or are their efforts co-ordinated? If information security takes the form of an information security group, does this group mete out its policies and practices from its ivory tower or does it regularly involve partners in architecture, development, operations, marketing, etc.?

### ***Security architecture and strategy***

Organizations with more than several hundred employees need to have someone (or *part* of a someone in smaller firms) thinking about enterprise security architecture. This is the discipline of modelling the organization's future security infrastructure that will supply centralized authentication, authorization, audit and public key services. The enterprise security architect (along with other IT architects) needs to become aware of the future trend of centralized, managed security infrastructures that provide security services to the enterprise applications and information infrastructure.

The organization's security strategy may have more to do with the organization's products and services, or it may not (it depends upon the organization's business). Regardless, someone in the organization (product and/or service architects and strategists or IT technology strategists) needs to be thinking of security-enabled products, services and business processes. A security-enabled enterprise can function only if someone has taken into account the impact of security on business processes as well as its products.

### ***Configuration management***

The integrity of the enterprise information infrastructure depends upon the quality processes and tools used to build and manage it. On the technology side, this means configuration management. Simply put, this is the technology that records everything about not only the configuration of the information infrastructure but also the company's products and services.

One of the enemies of security is mediocrity. An organization cannot build any but the most trivially simple infrastructure, products or services without a configuration management function to manage versions and configurations. Depending upon the size and scale of the enterprise, configuration management can be Microsoft Excel files, UNIX-based 'SCCS' (source code control system, which is also handy for most types of configuration files) or a full-blown enterprise configuration management application.

### **Security requirements**

Any organization that builds or uses widgets, software or services in all likelihood develops requirements in order to define formally the detailed behaviour of the goods and services that it buys or builds. Security requirements need to be a part of nearly every requirements definition document that the organization develops.

Like other Quadrant IV activities, I do not expect your organization to build its security requirements in a day, but over time an amalgamation included in various projects (or the security requirements that *should have* been included) can be carefully organized, scrutinized and improved – and eventually the organization will have a standard set of security requirements that it can tack on to almost any buy or build project henceforth.

### **Formal security awareness**

Most employees want to do the right thing, especially when it comes to protecting organization assets – if they only knew what to do. It's time to get security 'baked in' to the organization's products, services – and culture. But this cannot happen overnight. People need to become aware of security principles, which with discipline will change their thinking, habits, decisions and strategies.

Security awareness happens with training, activities and events. Like other corporate initiatives, security awareness will compete for space in employees' consciousness. In larger organizations, employees are bombarded with information on benefits, communities, the state of the business, compensation and incentive programmes, new products and services – the list goes on and on. Security awareness will have to *compete* for time and mindspace, and yet a security awareness programme must be fit somehow into the cacophony of messages employees hear every day.

## **EPILOGUE**

Readers like you, with large responsibilities in a fast-changing world, felt the need to learn more about the mysterious subject of information security. Hopefully you found that it was not so mysterious after all, but a methodical discipline based upon sound, time-tested principles such as 'least privilege' and 'defence in depth'.

The mystery of information security principles and practices probably remained with most of you because the security of physical things like buildings and people is far different from the security of information. Information can be stolen, and yet it remains. It can be changed, sometimes without a trace. And for the most part, the operating systems encircling corporate information assets were designed in an era when security was not the do-or-die imperative for ordinary private industry organizations but something only for the CIA and the KGB.

My goal was to get you to experience at least one illuminating discovery while reading this book, through having made the connection between something explained in this book and something about what your organization does or how it does it. Hopefully this helped to strip away the mystery about information security and will give you an appreciation of just how slippery a profession information security really is.

Because you understand more about information security, you can ask more constructive questions and, make more valuable decisions, knowing that you are now exercising the due diligence that is implicitly or explicitly required of you. You are leading your organization towards a brighter future.

---

# References/sources for additional information

This section contains an extensive list of online and in-print security information that you or others in your organization can use. References are grouped in the following categories:

- recommended reading
- portals
- government
- privacy
- corporate certifications
- personal certifications
- intelligence and research
- policy information and sources
- trade groups and associations
- trade shows, conferences, and seminars
- periodicals
- other security information.

Many of the references defy classification – some are part-periodical, part-research, part-products and services, and part-portal – in other words, they are multi-functional. The classifications here are quite arbitrary and unscientific. No attempts are made to rank, rate or assign value to any of these sites. One person's treasure is another's trash.

Many of the portals, trade groups and intelligence/research sites have e-mail-based 'news clipping services', which can be useful for busy individuals who want to stay up to date with events, technologies and trends, but who do not have time to trawl the Internet for information. The better ones provide granular subject matter and frequency choices that let you choose what you want to see and how often. Some websites may require registration.

The websites, periodicals and books mentioned here are not your destinations – they are your starting points. Each of you has different talents and needs and your organizations have different business models, balance sheets and missions.

## **RECOMMENDED READING**

---

Allen, Julia H. (1991) *The CERT® Guide to System and Network Security Practices*, Addison-Wesley.

- Anderson, Ross (2001) *Security Engineering*, John Wiley & Sons.
- Barman, Scott (2001) *Writing Information Security Policies*, New Riders.
- Carey, Peter (2000) *Data Protection in the UK*, Blackstone Press.
- Garfinkel, Simpson (2001) *Database Nation*, O'Reilly UK.
- Kaufman, Charlie (2002) *Network Security: Private Communication in a Public World*, Prentice Hall PTR.
- King, Christopher (2001) *Security Architecture*, Osborne McGraw-Hill.
- Kovacich, Gerald L. and Blyth, Andrew (2001) *Information Assurance: Surviving in the Information Environment*, Springer-Verlag UK.
- Prosise, Chris and Mandia, Kevin (2001) *Incident Response: Investigating Computer Crime*, Osborne McGraw-Hill.
- Schneier, Bruce (1996) *Applied Cryptography*, John Wiley & Sons.
- Schneier, Bruce (2000) *Secrets and Lies: Digital Security in a Networked World*, John Wiley & Sons.

## **PORTALS**

- |  |  |
|--|--|
| <a href="http://www.commoncriteria.org">www.commoncriteria.org</a>                                     | Common Criteria – criteria for evaluation of IT security.  |
| <a href="http://www.compseconline.com/compsec">www.compseconline.com/compsec</a>                       | CompSec Online – portal that focuses on computer and telecommunications security, biometrics and smart card information. Subscription and fee-based journals and research. |
| <a href="http://www.eurocert.net/legislature.html">www.eurocert.net/legislature.html</a>               | Computer Law and Legislation in European Countries. References to legislation related to the use of computers and networks in various European countries.                  |
| <a href="http://www.firewall.com">www.firewall.com</a>   | Firewall.com, a security portal.   |
| <a href="http://www.gartner.com/security">www.gartner.com/security</a>                                 | Gartner Group's research and information security portal.  |
| <a href="http://www.hi-media.co.uk/uk_security/index.htm">www.hi-media.co.uk/uk_security/index.htm</a> | UK-based Internet security portal.   |
| <a href="http://www.infosyssec.org">www.infosyssec.org</a>   | InfoSysSec, a <i>huge</i> information security portal.   |

<a href="http://www.osso.info/IT.htm">www.osso.info/IT.htm</a>	One Stop Security Online, a general-purpose IT security portal, based in the UK.
<a href="http://www.securityadvisor.info">www.securityadvisor.info</a>	<i>Security Advisor</i> .
<a href="http://www.securityfocus.com">www.securityfocus.com</a>	SecurityFocus, a security portal with news, advisories, products and services, and a vulnerability database.
<a href="http://www.shake.net/security.cfm">www.shake.net/security.cfm</a>	Shake Communications' security portal and guide to products and services.

## **GOVERNMENT**

<a href="http://www.bcs.org.uk/dataprot/dpc.htm">www.bcs.org.uk/dataprot/dpc.htm</a>	British Computer Society's Data Protection Committee home page.
<a href="http://www.europa.eu.int/comm/internal_market/en/media/dataprot">www.europa.eu.int/comm/internal_market/en/media/dataprot</a>	European Union Directive on Data Protection home page.
<a href="http://www.dataprotection.gov.uk">www.dataprotection.gov.uk</a>	UK Data Privacy Commissioner.

## **PRIVACY**

<a href="http://www.epic.org">www.epic.org</a>	Electronic Privacy Information Center, a public interest research centre in Washington, DC.
<a href="http://www.privacyalliance.org">www.privacyalliance.org</a>	Online Privacy Alliance, a group of global corporations and associations that introduce and promote business-wide actions that promote trust and foster the protection of online privacy.
<a href="http://www.w3.org/P3P">www.w3.org/P3P</a>	Platform for Privacy Preferences Project (P3P), developed by the World Wide Web Consortium which is emerging as an industry standard providing a way for users to gain more control over the use of personal information on websites they visit.
<a href="http://www.privacyexchange.org">www.privacyexchange.org</a>	A portal that focuses on consumer privacy, e-commerce and data protection.

[www.privacy.org/pi](http://www.privacy.org/pi)

Privacy International, a human rights group formed in 1990 as a watchdog on surveillance by governments and corporations. Based in London.

[www.ftc.gov/privacy](http://www.ftc.gov/privacy)

General information on privacy from the US Federal Trade Commission.

## **CORPORATE CERTIFICATIONS**

---

[www.sas70.com](http://www.sas70.com)

SAS70 (Statement on Auditing Standards No. 70), an internationally recognized auditing standard developed by the American Institute of Certified Public Accountants (AICPA).

[www.aicpa.org/assurance/systrust/index.htm](http://www.aicpa.org/assurance/systrust/index.htm)

Systrust is an information system reliability assessment methodology developed by the AICPA and Canadian Institute of Chartered Accountants (CICA).

[www.truste.org](http://www.truste.org)

TRUSTe is a non-profit privacy organization that issues a branded online seal indicating that a website meets TRUSTe's privacy guidelines.

[www.aicpa.org/assurance/webtrust/index.htm](http://www.aicpa.org/assurance/webtrust/index.htm)

Webtrust is an information system assessment methodology that is focused on privacy, security, availability, confidentiality, consumer redress for complaints and business practices. Webtrust was developed by AICPA and CICA.

## **PERSONAL CERTIFICATIONS**

---

[www.isc2.org](http://www.isc2.org)

The global non-profit organization that issues the CISSP (Certified Information Systems Security Professional) and SSCP (System Security Certified Practitioner) professional certifications.

<a href="http://www.isaca.org/cert1.htm">www.isaca.org/cert1.htm</a>	The American Institute of Certified Public Accountants (AICPA), which issues the CISA (Certified Information Systems Auditor) professional certification.
<a href="http://www.bcs.org/iseb">www.bcs.org/iseb</a>	The Information Systems Examining Board (ISEB), a part of the British Computer Society, which issues the Certificate in Data Protection (British Computer Society) professional certification.
<a href="http://www.giac.org">www.giac.org</a>	The Global Information Assurance Certification (GIAC), a series of hands-on security professional certifications issued by the SANS (Systems Administration, Networks and Security) Institute.

## INTELLIGENCE AND RESEARCH

<a href="http://rr.sans.org">rr.sans.org</a>	SANS Institute's Information Security Reading Room, which contains white papers on numerous information security topics.
<a href="http://www.gartner.com/security">www.gartner.com/security</a>	The information security main page for the Gartner Group. Contains free and subscription-based research and analysis materials.
<a href="http://www.gigaweb.com">www.gigaweb.com</a>	Giga Information Group. Contains free and subscription-based research and analysis materials.
<a href="http://www.cio.com/research/security">www.cio.com/research/security</a>	Information security news and analysis from <i>CIO</i> magazine.
<a href="http://www.idc.com">www.idc.com</a>	Information security news and analysis from IDC.
<a href="http://www.metagroup.com">www.metagroup.com</a>	Information security news and analysis from the Meta Group.
<a href="http://www.cerias.purdue.edu">www.cerias.purdue.edu</a>	Center for Education and Research in Information Assurance and Security (CERIAS), a research and development

[www.cert.org](http://www.cert.org) department at Purdue University's Computer Science department in West Lafayette, Indiana. The CERT Coordination Center (CERT/CC), a centre of Internet security expertise, at the Software Engineering Institute, a federally funded research and development centre operated by Carnegie Mellon University in Pittsburg, Pennsylvania.

## **POLICY INFORMATION AND SOURCES**

[rr.sans.org/policy](http://rr.sans.org/policy) SANS (Systems Administration, Networks and Security) Institute Reading Room containing articles on security policy issues.

[www.iso17799.net](http://www.iso17799.net) ISO17799 is an internationally recognized information security standard from the British Standards Institution (BSI) and British Standards Publishing Limited (BSPL).

[www.baselinesoft.com/ispme.html](http://www.baselinesoft.com/ispme.html) *Information Security Policies Made Easy*, a vast and impressive collection of information security policies in electronic form, from Pentasafe.

[www.information-security-policies.com](http://www.information-security-policies.com) *RUSecure Information Security Policy Suite*, another impressive collection of canned information security policies in electronic form.

[web.mit.edu/security/www/gassp1.html](http://web.mit.edu/security/www/gassp1.html) Generally Accepted System Security Principles, a good collection of high-level information security principles.

## **TRADE GROUPS AND ASSOCIATIONS**

[www.gocsi.com](http://www.gocsi.com) Computer Security Institute (CSI), an individual membership organization

	offering training, conferences and research materials.
<a href="http://www.securityforum.org">www.securityforum.org</a>	Information Security Forum (previously known as the European Security Forum), a corporate membership organization that promotes the protection of business information.
<a href="http://www.isalliance.org">www.isalliance.org</a>	Internet Security Alliance, a corporate membership organization that promotes sound information security practices, policies, and technologies.
<a href="http://www.eema.org">www.eema.org</a>	European Forum for Electronic Business, a corporate membership association that promotes the electronic-based marketplace.
<a href="http://www.teletrust.de">www.teletrust.de</a>	TeleTrusT, a corporate membership association that promotes the security of information and communication technology.
<a href="http://www.first.org">www.first.org</a>	Forum of Incident Response and Security Teams (FIRST), a consortium of over 100 incident response and security teams that co-ordinates worldwide response to Internet security incidents.
<a href="http://www.isaca.org">www.isaca.org</a>	Information Systems Audit and Control Association (ISACA), an individual membership association that promotes IT governance, control and assurance.
<a href="http://www.isoc.org">www.isoc.org</a>	The Internet Society, an individual membership association providing leadership in addressing issues that confront the future of the Internet.
<a href="http://www.issa-intl.org">www.issa-intl.org</a>	Information Systems Security Association (ISSA), an individual membership association promoting the knowledge, skill and professional growth of its members.

<a href="http://www.htcia.org">www.htcia.org</a>	High Technology Crime Investigation Association, an individual membership association promoting the knowledge of computer crime investigation techniques.
<a href="http://www.issea.org">www.issea.org</a>	International Systems Security Engineering Association (ISSEA), an individual membership association promoting the adoption of systems security engineering as a defined and measurable discipline.

## **TRADE SHOWS, CONFERENCES AND SEMINARS**

<a href="http://www.infosec.co.uk">www.infosec.co.uk</a>	InfoSecurity Europe, a trade show with information security product and service exhibits and training classes.
<a href="http://www.isse.org">www.isse.org</a>	Information Security Solutions Europe, a trade show with information security product and service exhibits and training classes.
<a href="http://www.gartner.com/events">www.gartner.com/events</a>	Gartner Group: various conferences and seminars on all IT topics.
<a href="http://www.interop.com">www.interop.com</a>	Interop: conferences and seminars on networking.
<a href="http://www.sans.org">www.sans.org</a>	SANS (Systems Administration, Networks and Security) Institute: various conferences and seminars on systems administration and security.
<a href="http://www.misti.com/europe.asp">www.misti.com/europe.asp</a>	MISTI (MIS Training Institute) Europe: various seminars and conferences on information technology.
<a href="http://www.gigaworld europe.com">www.gigaworld europe.com</a>	GigaWorld IT Forum Europe, a conference featuring keynote speakers, seminars and time with Giga analysts.
<a href="http://www.gocsi.com">www.gocsi.com</a>	Computer Security Institute (CSI): various conferences and seminars on information security.

## PERIODICALS

<a href="http://www.infosecuritymag.com">www.infosecuritymag.com</a>	<i>Information Systems Security.</i>
<a href="http://www.scmagazine.com">www.scmagazine.com</a>	<i>Secure Computing Magazine.</i>
<a href="http://www.securitymagazine.com">www.securitymagazine.com</a>	<i>Security Magazine.</i>
<a href="http://www.securitymanagement.com">www.securitymanagement.com</a>	<i>Security Management.</i>
<a href="http://www.isr.net">www.isr.net</a>	<i>Internet Security Review.</i>
<a href="http://www.isec-worldwide.com">www.isec-worldwide.com</a>	<i>Information Security World.</i>
<a href="http://www.infosecnews.com">www.infosecnews.com</a>	<i>InfoSecurity News.</i>
<a href="http://www.2600.com">www.2600.com</a>	<i>2600 Magazine</i> , a hacker's quarterly periodical.
<a href="http://www.ieee-security.org/cipher.html">www.ieee-security.org/cipher.html</a>	<i>Cipher</i> , IEEE's (Institute for Electrical and Electronic Engineers) security magazine.
<a href="http://www.compseconline.com/compsec">www.compseconline.com/compsec</a>	Several computer/network security journals.

## OTHER SECURITY INFORMATION

<a href="http://www.sse-cmm.org">www.sse-cmm.org</a>	Systems Security Engineering Capability Maturity Model.
--	---