BUSINESS | PEOPLE | TECHNOLOGY

Bridging the gap between business and technology

Cybersecurity Without the Jargon

A PRACTICAL GUIDE TO PROTECTING YOUR BUSINESS, YOUR DATA, AND YOUR PEOPLE

CHARLES EDWARDS

trg networking

BRIDGING THE GAP BETWEEN BUSINESS AND TECHNOLOGY

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

For Small and Mid-Sized Organizations

Dr. Charles Edwards

TRG Networking Inc.

First Edition - 2026

Copyright and Publication Information

This book is intended as an educational resource. It provides general cybersecurity and technology risk-management information and is not a substitute for legal advice, regulatory advice, formal compliance assessment, incident-response services, or security services tailored to a specific organization.

Cybersecurity requirements and technologies change over time. Organizations should evaluate current laws, contractual obligations, regulatory requirements, vendor guidance, and authoritative cybersecurity standards before making security or compliance decisions.

References to products, services, standards, organizations, or technologies are provided for educational purposes and do not necessarily constitute an endorsement.

Published in association with TRG Networking Inc.

First Edition - 2026

Author: Dr. Charles Edwards

Title: Cybersecurity Without the Jargon: A Practical Guide to Protecting Your Business, Your Data, and Your People

Dedication

To the business owners, organizational leaders, educators,
IT professionals, and employees who carry the responsibility
of protecting information and keeping their organizations running.

Cybersecurity is not only about technology.

It is about people, trust, resilience, and the ability
to continue serving those who depend on us.

And to my family, whose encouragement, patience,
and support have made every professional journey possible.

About the Author

Dr. Charles Edwards is an information security and technology professional, educator, and Chief Operating Officer/Chief Technology Officer of TRG Networking Inc. His work brings together cybersecurity, computer networking, engineering, technology management, cyber forensics, and education.

Dr. Edwards holds a Ph.D. in Information Systems, a Master's degree in Electrical Engineering, and a Master's degree in Computer Engineering. This combination of security, engineering, and computing has shaped his practical approach to technology: understand the business problem first, identify the risk, and then apply technology in a way that people can realistically manage.

As an educator at Howard Community College, Dr. Edwards teaches in areas that include cybersecurity, cyber forensics, operating systems, and computer programming. His teaching philosophy emphasizes practical understanding. Students should not simply memorize terminology; they should understand what a technology does, why it matters, how it can fail, and how to apply it responsibly.

That same philosophy is at the heart of this book. Cybersecurity is often presented in language designed for specialists. Yet many of the people who must make cybersecurity decisions are business owners, executives, managers, nonprofit leaders, government contractors, educators, and professionals whose primary responsibility is not information security.

About the Author - Continued

The goal of Dr. Edwards's work is to help close that gap. A leader does not need to become a security engineer to make better cybersecurity decisions. A leader does, however, need to understand the organization's risks, know what questions to ask, recognize when an answer is incomplete, and understand which protections deserve priority.

Through TRG Networking Inc., Dr. Edwards is involved in technology and cybersecurity services that support organizations in areas such as managed IT, cybersecurity, Microsoft 365 and cloud services, Azure cloud hosting, business continuity, secure adoption of artificial intelligence, help desk support, and cybersecurity readiness.

Cybersecurity Without the Jargon reflects both sides of his professional experience: the educator's responsibility to explain difficult concepts clearly and the technology leader's responsibility to turn those concepts into practical action.

The book is written for readers who want to understand cybersecurity without being overwhelmed by acronyms, fear-based messaging, or unnecessary complexity. Its central message is straightforward: cybersecurity risk cannot be eliminated, but it can be understood, prioritized, managed, and reduced.

A Message to the Reader

If you are responsible for a business, nonprofit organization, government-contracting company, school, professional practice, or growing enterprise, cybersecurity has probably become part of your job whether you asked for it or not.

You may have an IT provider. You may have internal technical staff. You may use Microsoft 365, cloud applications, laptops, mobile devices, remote workers, vendors, and online banking. You may also have contracts or regulations that require specific safeguards. Yet at some point, someone in leadership still has to ask: Are we adequately protected?

That question is harder than it sounds. Security products are plentiful. Acronyms are everywhere. Vendors often describe every problem as urgent. News headlines focus on catastrophic breaches. Meanwhile, smaller organizations have limited budgets, limited staff, and a business to run.

This book is for that reality.

It will not ask you to become a cybersecurity specialist. Instead, it will help you understand what matters, why it matters, what can go wrong, and what reasonable action looks like. Each chapter will translate a cybersecurity issue into business terms and finish with practical questions and next steps.

The goal is not fear. The goal is informed action.

Preface - Why I Wrote This Book

Cybersecurity has developed a language of its own. Ransomware. Multifactor authentication. Zero trust. Endpoint detection and response. Vulnerability management. Business email compromise. CMMC. NIST. Artificial intelligence governance. For technology professionals, these terms may be part of everyday conversation. For a business leader, they can make cybersecurity seem more complicated than it needs to be.

At its core, cybersecurity is about protecting a few very important things: your people, your information, your systems, your customers, your reputation, and your ability to continue operating.

I wrote this book because organizations frequently encounter two opposite problems. Some make cybersecurity so complicated that leaders disengage and leave every decision to technical staff or vendors. Others assume that because nothing serious has happened yet, their security must be adequate. Neither approach provides a sound basis for managing risk.

Small and mid-sized organizations face an additional challenge: resources. They rarely have unlimited security budgets or large internal security teams. Current NIST small-business guidance similarly emphasizes practical risk management and recognizes that implementation varies with an organization's size, resources, sector, and requirements.

The right question is therefore not, “How do we buy every security product?” The better question is, “What risks matter most to us, and what should we do first?”

Preface - A Practical Approach

This book approaches cybersecurity from the perspective of the person responsible for the organization, not from the perspective of the person configuring a firewall.

We will examine how attackers gain access, why identity has become so important, why email remains a major avenue of attack, what ransomware can do to operations, why cloud services still require security decisions, how remote work changes risk, why backups must be tested, what government contractors need to understand about protecting sensitive information, and how artificial intelligence is creating both opportunity and new forms of exposure.

Throughout the book, the emphasis will remain practical. Each chapter will include realistic scenarios, plain-English explanations, common mistakes, safeguards, questions to ask your IT team or service provider, and a short action checklist.

The book also builds toward a larger goal. By the final chapters, the reader should be able to develop a practical 30-day cybersecurity action plan and identify the areas of the organization that deserve further attention.

Cybersecurity is not about eliminating every possible risk. No organization can do that. It is about understanding risk, making intelligent decisions, implementing reasonable safeguards, detecting problems, responding effectively, and recovering when something goes wrong.

You do not need an unlimited cybersecurity budget.
You need a practical plan.

How to Use This Book

This book is designed to be read from beginning to end, but it can also be used as a practical reference. If your immediate concern is ransomware, Microsoft 365, government-contracting requirements, artificial intelligence, or incident response, you can go directly to that chapter and return to the earlier material when needed.

Each chapter follows a consistent structure so you always know what to expect.

- The Business Problem - why the topic matters to the organization.
- What You Need to Know - the essential concepts explained in plain English.
- A Realistic Scenario - a practical example showing how the issue may appear in an organization.
- What Went Wrong? - the weaknesses or decisions that allowed the problem to develop.
- What Should Have Happened? - reasonable safeguards and better decisions.
- What This Means for Your Organization - the leadership and operational implications.
- Questions to Ask Your IT Provider or Internal IT Team - questions that help leaders evaluate their current position.
- TRG Practical Takeaway - a concise set of priority actions.
- Before You Move On - a short checklist for applying the chapter.
- Key Terms and Chapter Summary - plain-English reinforcement of the major ideas.

Do not feel that every recommendation must be implemented at once. Cybersecurity is a risk-management process. Prioritize the actions that reduce the greatest risk, address contractual or regulatory obligations, and protect the systems and information your organization depends upon most.

Table of Contents

Table of Contents - Continued

CHAPTER 1

Cybersecurity Is a Business Problem

Cybersecurity is often treated as an IT problem. I understand why. When people hear the word cybersecurity, they think about computers, firewalls, passwords, networks, cloud systems, and hackers.

But that is not where the real problem ends. When something goes wrong, the impact is felt by the entire business. Employees may not be able to work. Customers may not be able to reach you. Payroll can be delayed. Important files can become unavailable. Confidential information can be exposed. Contracts can be affected. Very quickly, what started as a technology problem becomes a business problem that management has to deal with.

That is why I tell clients that cybersecurity cannot be left entirely to the IT department or to an outside technology company. Technology is certainly part of the solution, but leadership still has to understand what the organization depends on, what information matters most, what could stop the business from operating, and how quickly the company could recover.

> The question I would ask any business owner is simple: If your systems were unavailable tomorrow morning, what part of your business would stop first - and how long could you operate that way?

1.1 Why Cybersecurity Belongs in the Leadership Conversation

A business leader does not need to know how to configure a firewall or investigate malicious code. Leadership does need to understand what the organization depends on, which information is sensitive, what an interruption would cost, which obligations apply, and whether the organization can recover. These are management questions before they are technical questions.

Consider how much of an ordinary workday now depends on digital systems. Employees sign in to cloud applications. Invoices are sent electronically. Customers communicate by email. Files live in shared storage. Payroll, banking, scheduling, customer records, contracts, proposals, and vendor relationships may all depend on systems that are reachable through the Internet. When those systems fail, the business problem becomes visible immediately.

Cybersecurity therefore belongs alongside financial risk, legal risk, operational risk, personnel risk, and physical security. It is not a separate universe. It is one more way that the organization can lose the ability to perform its mission.

1.2 A Realistic Scenario: Monday Morning

Imagine a 45-person professional-services company. On Monday morning, several employees cannot open shared files. Others see unusual file extensions and ransom notes. The accounting team cannot reach a folder containing current invoices. A project manager cannot retrieve a proposal due that afternoon. The help desk discovers that an administrative account was used overnight to access multiple systems.

At first, management asks a technical question: “How quickly can IT fix this?” Within an hour, the questions change. Can employees work? Are client files exposed? Should customers be notified? Can payroll run? Are backups safe? Is the attacker still inside the network? Does the company need outside incident-response help? Does insurance require immediate notice? Does a contract require reporting? Who is authorized to speak to clients?

Nothing about those questions is merely technical. The incident has become an operational, financial, legal, contractual, and reputational event. That is why cybersecurity planning cannot begin after the ransom note appears.

1.3 What Are We Actually Protecting?

Organizations sometimes describe cybersecurity as protecting computers. Computers matter, but they are not the ultimate objective. The real objective is to protect the organization’s ability to function and the assets on which that function depends.

People

Employees, customers, vendors, partners, and leaders all interact with systems and information. Attackers routinely target people because convincing a person to approve a login, reveal a password, open a document, or change a payment instruction can be easier than defeating a well-configured security device.

Information

Customer records, employee information, financial data, intellectual property, contracts, credentials, business plans, regulated information, and internal communications can all have value. Some information is valuable because it can be sold. Some is valuable because it can be used for fraud. Some is valuable because disclosure would damage the organization or its clients.

Systems and Services

Email, identity systems, line-of-business applications, cloud platforms, websites, file storage, accounting systems, remote-access tools, and network services keep the

organization moving. Availability matters just as much as confidentiality. A perfectly confidential system that nobody can access during a crisis is still a business failure.

Trust and Reputation

Customers and partners make decisions based partly on trust. A breach does not automatically destroy that trust, but poor preparation, delayed communication, misleading statements, or repeated failures can. Reputation is difficult to quantify until it is damaged.

1.4 Cyber Risk Is Business Risk

Risk can be discussed in many ways, but a useful business-level model is straightforward: something valuable exists, a threat could affect it, a weakness makes that threat more likely or more damaging, and the organization experiences a consequence. Cybersecurity is the process of understanding that chain and deciding what to do about it.

This does not mean eliminating every weakness. That is impossible. It means identifying the risks that matter most and applying safeguards proportionate to the organization's needs, resources, obligations, and tolerance for disruption.

For a small organization, the highest priorities may be protecting email accounts, requiring multifactor authentication, keeping systems patched, maintaining recoverable backups, training employees to recognize fraud, and having a simple incident-response plan. A government contractor handling controlled information may have additional contractual and compliance requirements. A healthcare organization may face a different combination of privacy, availability, and regulatory concerns.

1.5 The Cost of "We Are Too Small to Be a Target"

One of the most persistent cybersecurity misconceptions is that attackers only care about large corporations. In reality, attackers often operate at scale. They send phishing messages to thousands of recipients, scan large portions of the Internet for exposed systems, reuse stolen passwords automatically, and exploit opportunities wherever they find them.

A smaller organization may not be selected because of its name. It may simply be selected because an account was compromised, a system was exposed, a vendor relationship was useful, or an employee responded to a convincing message. From the attacker's perspective, opportunity can matter more than prestige.

The correct question is not "Why would anyone target us?" It is "What could someone gain by compromising us, and what could we lose?" Those answers may include money, credentials, customer information, access to a larger partner, sensitive

documents, computing resources, or the ability to extort the organization by interrupting operations.

1.6 Cybersecurity Is Also About Availability

Security discussions often focus on stolen information. That is only one part of the problem. An organization can suffer a severe cyber incident even if no confidential data is publicly released. If employees cannot access email, files, applications, or customer systems, the business may still be unable to operate.

This is why business continuity and disaster recovery belong in cybersecurity planning. Leaders should know which systems are essential, how long the organization can tolerate their loss, what alternatives exist, where backups are stored, and whether restoration has actually been tested.

Leadership question

> *When someone says, "We have backups," ask the next question: "When did we last restore from them successfully?"*

1.7 The Human Side of Cybersecurity

Employees are sometimes described as the weakest link in cybersecurity. That phrase is convenient, but it can be counterproductive. People are part of the security system. They can make mistakes, but they can also recognize suspicious behavior, report unusual requests, challenge fraudulent payment changes, protect credentials, and stop an incident before it grows.

The goal should not be to make employees afraid of clicking anything. The goal should be to create a culture in which people know what suspicious activity looks like, understand how to report it, and are not punished for reporting a mistake quickly. Early reporting can dramatically reduce the damage from a compromised account or malicious attachment.

Leadership sets that culture. If managers bypass security controls because they are inconvenient, employees notice. If executives share passwords, approve exceptions without review, or treat security training as a nuisance, the organization receives the message that cybersecurity is optional.

1.8 Technology Alone Is Not a Cybersecurity Program

Organizations can purchase excellent security products and still remain poorly prepared. A tool may be installed but not monitored. Alerts may be generated but nobody may be responsible for reviewing them. Backups may run but never be tested.

Multifactor authentication may be enabled for most users but excluded from the accounts attackers would value most. Policies may exist but be unknown to employees.

A cybersecurity program connects technology to responsibility. Someone owns each critical process. Important systems are inventoried. Access is reviewed. Updates are applied. Backups are protected. Incidents are reported. Vendors are evaluated. Leadership receives meaningful information instead of a list of technical alerts.

This is also why buying more products is not always the correct first step. Before purchasing another security platform, an organization should understand the gaps it is trying to close and whether existing capabilities are being used properly.

1.9 What Good Cybersecurity Governance Looks Like

Governance sounds formal, but at a small or mid-sized organization it can be simple. It means deciding who is responsible, what matters most, what rules apply, how risk decisions are made, and how leadership knows whether important protections are working.

At a minimum, leadership should be able to answer several questions:

- Who is accountable for cybersecurity decisions?
- Which systems and information are most critical to the organization?
- Who has administrative or privileged access?
- Is multifactor authentication required where it matters most?
- How are vulnerabilities and software updates handled?
- How are backups protected and tested?
- What happens when an employee reports a suspicious message or compromised account?
- Which vendors can access systems or sensitive information?
- What contractual, regulatory, insurance, or reporting obligations apply?
- Who makes decisions during a serious cyber incident?

If these questions cannot be answered, the organization has identified useful work to do. “I do not know” is not a failure; it is a finding. The dangerous situation is not knowing and assuming everything is fine.

1.10 Measuring Cybersecurity in Business Terms

Technical teams naturally measure technical things: blocked attacks, vulnerabilities, patch levels, malware detections, firewall events, and endpoint alerts. Those measures can be useful, but executives also need measures connected to business outcomes.

For example, leadership may want to know the percentage of accounts protected by multifactor authentication, the number of privileged accounts, the age of unresolved critical vulnerabilities, the success rate of backup restoration tests, the time required to disable a departed employee's access, the number of employees who completed security training, and the time required to respond to a reported incident.

Good metrics should lead to decisions. A dashboard full of numbers that nobody acts upon is decoration, not governance.

1.11 Common Leadership Mistakes

Mistake 1: Treating cybersecurity as an IT-only responsibility

IT may operate many safeguards, but leadership owns business risk. Technical staff cannot decide alone how much downtime the organization can tolerate, which contractual obligations are most important, or how much risk leadership is willing to accept.

Mistake 2: Assuming compliance equals security

Compliance requirements can provide valuable structure, but passing an assessment does not guarantee that an organization is secure. Security should support the business every day, not only during an audit.

Mistake 3: Waiting for an incident to define responsibilities

A crisis is the wrong time to decide who calls the insurer, who contacts legal counsel, who communicates with customers, who preserves evidence, and who has authority to shut down a system.

Mistake 4: Focusing only on prevention

Prevention matters, but no defense is perfect. Organizations also need detection, response, recovery, and lessons learned. Resilience assumes that something will eventually go wrong and prepares the organization to continue.

Mistake 5: Believing more spending automatically means more security

Security spending should follow risk and priorities. Basic controls that are consistently implemented can be more valuable than expensive products that are poorly configured or ignored.

1.12 What Should Have Happened in the Monday-Morning Scenario?

Return to the professional-services company from the beginning of the chapter. A better-prepared organization would not necessarily prevent every ransomware attempt,

but it would make the incident harder to start, easier to detect, and easier to recover from.

Administrative access would be limited and protected with strong authentication. Employees would know how to report suspicious activity. Important systems would be patched. Endpoint protections would be monitored. Backups would be separated from ordinary user access and restoration would be tested. The organization would know which systems must be restored first. An incident-response plan would identify decision-makers, technical contacts, legal and insurance contacts, communication responsibilities, and escalation procedures.

Most importantly, leadership would not be inventing the plan while the business was already disrupted.

1.13 What This Means for Your Organization

The first step in cybersecurity is not purchasing software. It is understanding dependence. What technology does the organization depend on? What information would cause harm if exposed? Which accounts could create the most damage if compromised? Which vendors have access? How long could the organization operate without email, accounting, shared files, or customer systems?

Once those questions are answered, technical decisions become easier to prioritize. Cybersecurity stops being an endless list of threats and becomes a manageable set of business risks.

1.14 Questions to Ask Your IT Provider or Internal IT Team

- What are our five most critical systems or technology services?
- Which accounts have administrator-level privileges, and how are they protected?
- Is multifactor authentication required for email, cloud services, remote access, and privileged accounts?
- How quickly could we restore our most critical systems after ransomware or a major outage?
- When was our last successful backup restoration test?
- How do employees report suspected phishing, fraud, or compromised accounts?
- Who monitors security alerts, and what happens after a serious alert is generated?
- How quickly are critical security updates applied?
- Which third parties have access to our systems or sensitive data?
- Do we have a written incident-response plan, and when was it last exercised?

1.15 TRG Practical Takeaway

Cybersecurity should be discussed in the same language as other business risks: impact, likelihood, responsibility, priority, cost, and recovery. Leaders do not need to perform every technical task, but they do need enough understanding to ask meaningful questions and make informed decisions.

Three actions to take now

1. Identify the systems and information your organization cannot operate without. 2. Confirm who is accountable for cybersecurity decisions and incident response. 3. Ask when your organization last tested recovery from a serious technology disruption.

1.16 Before You Move On

- ☐ We can identify our most important systems, information, and business processes.
- ☐ Leadership understands that cybersecurity is an organizational risk, not only an IT task.
- ☐ We know who is responsible for cybersecurity decisions and incident escalation.
- ☐ We know whether our critical accounts use multifactor authentication.
- ☐ We know where critical backups are stored and whether restoration has been tested.
- ☐ Employees know how to report suspicious activity or mistakes quickly.
- ☐ We know which vendors have access to important systems or sensitive information.
- ☐ We have considered contractual, regulatory, insurance, and reporting obligations.
- ☐ We have a basic incident-response plan with named decision-makers.
- ☐ We can explain our highest cybersecurity priorities in business terms.

1.17 Key Terms in Plain English

Asset: Anything the organization values and needs to protect, including information, systems, services, devices, identities, reputation, and business processes.

Threat: Something capable of causing harm, such as a criminal group, malicious insider, compromised account, ransomware operation, or destructive event.

Vulnerability: A weakness that can be exploited or can increase the effect of an incident.

Risk: The possibility that a threat will affect something the organization values and create a harmful consequence.

Control or Safeguard: A measure used to reduce risk. It may be technical, administrative, physical, or procedural.

Resilience: The ability to prepare for disruption, continue essential functions, respond effectively, and recover.

Incident Response: The organized process for detecting, containing, investigating, communicating about, and recovering from a cybersecurity incident.

Multifactor Authentication (MFA): A sign-in process requiring more than one type of proof, reducing the value of a stolen password by itself.

1.18 Chapter Summary

Cybersecurity is not successful when an organization owns the most security products. It is successful when the organization understands what it depends on, protects what matters, reduces avoidable risk, detects problems, responds with discipline, and can recover.

The rest of this book builds on that foundation. Before we discuss phishing, ransomware, Microsoft 365, artificial intelligence, government-contracting requirements, or incident response, we first need to understand the assets that make the organization worth protecting. That is the focus of Chapter 2: Know What You Are Protecting.

Chapter 1 References and Further Reading

This chapter draws on current authoritative guidance, including the NIST Cybersecurity Framework 2.0, NIST cybersecurity resources for small businesses, and relevant CISA guidance. Because cybersecurity guidance changes over time, readers should verify current authoritative sources when making time-sensitive security or compliance decisions.

1.19 A Simple Cyber Risk Decision Framework

When leaders are presented with a cybersecurity concern, the conversation can quickly become technical. A simple decision framework helps bring the discussion back to business priorities. Start by identifying the business asset or process at risk. Then ask what could happen, how disruptive the consequence would be, what protections already exist, what important gap remains, and what practical action would reduce the risk.

For example, suppose the organization learns that several employees can approve financial transactions through email and online banking. The business asset is money and the payment process. The threat may be account takeover or impersonation.

Existing safeguards might include passwords and a bank approval process. The gap may be that email accounts do not use multifactor authentication and payment changes can be accepted through email alone. A practical response could combine MFA, independent verification of payment changes, and dual approval for higher-value transactions.

This framework prevents a common mistake: beginning with a product instead of beginning with the risk. The question is not, “Should we buy Product X?” The question is, “What business problem are we trying to reduce, and is this the most appropriate way to reduce it?”

Leaders should also distinguish between accepting a risk and ignoring it. Risk acceptance is a conscious decision made after understanding the possible consequence and alternatives. Ignoring a risk means no informed decision was made. Good governance documents important risk decisions so that the organization understands why a particular course was chosen.

1.20 Leadership Exercise: Map One Critical Business Process

Before leaving this chapter, choose one business process that would create a serious problem if it stopped tomorrow. It might be payroll, customer billing, order processing, email, patient scheduling, contract delivery, remote access, or access to shared project files. Do not begin with technology. Begin with the business process.

Write down the people who perform the process, the systems they use, the information they need, the vendors involved, and the accounts or credentials required. Then ask what happens if each dependency is unavailable for one hour, one day, three days, or one week. This exercise often reveals dependencies that were previously invisible.

Next, identify the person who can answer how each dependency is protected and recovered. If nobody can answer, record that as an action item. If the answer is based on an assumption rather than a test, record that too. The purpose is not to assign blame. The purpose is to replace assumptions with knowledge.

Finally, identify one improvement that could be completed within the next 30 days. A useful first improvement is specific, owned by a named person or role, and easy to verify. “Improve security” is not an action. “Require MFA for all Microsoft 365 administrator accounts by the end of the month” is an action.

- ☐ Critical business process selected
- ☐ Key people and systems identified
- ☐ Sensitive information identified
- ☐ External vendors/dependencies identified

- ☐ Maximum tolerable interruption discussed
- ☐ Recovery method identified
- ☐ Recovery method tested or test scheduled
- ☐ One 30-day improvement assigned

1.21 A Note for Small and Mid-Sized Organizations

Smaller organizations sometimes compare themselves with large enterprises and conclude that meaningful cybersecurity is unaffordable. That comparison is rarely useful. A 30-person organization does not need to reproduce the security department of a global corporation. It needs safeguards appropriate to its own risks, obligations, systems, and resources.

The strongest starting point is disciplined use of fundamentals: know what you have, protect important identities, keep systems current, reduce unnecessary privileges, train people, protect and test backups, plan for incidents, and review important vendors. More advanced tools can be added when the risk and business case justify them.

This is the approach we will follow throughout the book. Each chapter will separate essential actions from enhancements so that readers can make progress without assuming that cybersecurity requires an unlimited budget.

Chapter 2
Know What You Are Protecting

Cybersecurity Without the Jargon

Dr. Charles Edwards | TRG Networking Inc. | First Edition - 2026

Before you can protect an organization, you have to know what you are protecting. That sounds obvious, but this is one of the places where many cybersecurity programs get off track.

A company may buy a firewall, renew antivirus, turn on multifactor authentication, and schedule backups. Those are all useful steps. But I still want leadership to be able to answer a more basic question: Which systems, information, identities, and business processes would hurt us most if they were stolen, changed, exposed, or unavailable?

By the end of this chapter, you should be able to identify the major assets your organization depends on, understand why they matter to the business, and begin deciding which ones deserve the greatest protection and recovery priority.

2.1 Start With the Business, Not the Technology

Imagine a 40-person professional services company. It has laptops, Microsoft 365, a line-of-business application, cloud file storage, accounting software, a website, mobile phones, and several outside vendors. A traditional technology inventory might list hardware, software, IP addresses, and licenses. That is useful, but it does not tell leadership what matters most.

A business-centered inventory asks different questions. Which system creates invoices? Where are customer contracts stored? Who can approve wire transfers? Which mailbox receives password-reset messages? What data must remain confidential? Which application would stop operations if it were unavailable for two days? Who knows how to restore the accounting system? Which vendor has administrative access?

Those questions reveal dependencies. Cybersecurity is ultimately about protecting the organization’s ability to achieve its mission. Technology is part of that mission, but it is not the whole mission.

2.2 The Five Things Every Organization Is Really Protecting

- People and identities — employees, executives, contractors, administrators, customers, and the accounts that represent them.
- Information and data — financial records, customer information, employee records, intellectual property, contracts, credentials, regulated information, and business communications.
- Systems and technology — endpoints, servers, networks, cloud services, applications, websites, databases, security tools, and communications platforms.
- Business operations — the processes that allow the organization to sell, deliver, bill, communicate, manufacture, support customers, and meet contractual obligations.
- Trust and reputation — the confidence customers, employees, partners, regulators, and the public place in the organization.

These categories overlap. A compromised executive account can expose confidential email, authorize fraudulent payments, interrupt operations, and damage customer trust in the same incident. That is why asset protection must be viewed as a business discipline rather than merely a technical inventory exercise.

2.3 People and Digital Identities

In modern organizations, a person's digital identity can be as important as the physical device he or she uses. An attacker who steals an employee's Microsoft 365 credentials may not need to compromise the employee's laptop at all. The account itself can provide access to email, files, calendars, collaboration tools, password-reset messages, and business relationships.

Privileged identities deserve special attention. Global administrators, domain administrators, cloud administrators, finance approvers, payroll personnel, and executives may have access or authority that makes their accounts unusually valuable. Service accounts and application identities can be equally important even though no human logs into them every morning.

Organizations should know which identities exist, who owns them, what privileges they possess, how they are authenticated, and what happens when an employee or contractor leaves. Dormant accounts, shared administrator credentials, excessive privileges, and forgotten vendor access are common examples of identity risk.

Practical Question: If one account in your organization were compromised today, which account could cause the greatest damage? Leadership should know the answer.

2.4 Information: Not All Data Has the Same Value

A common mistake is to treat all organizational data as though it requires the same protection. A publicly available brochure and a spreadsheet containing employee Social Security numbers are both files, but their confidentiality requirements are very different. Likewise, a blank form may have little value, while the completed version may contain sensitive personal or financial information.

A practical data inventory should identify what information the organization collects, creates, receives, stores, transmits, and destroys. It should also consider why the organization has the information in the first place. Keeping unnecessary sensitive data creates risk without creating business value.

- Public information — material intentionally available to customers or the public.
- Internal information — routine operational information not intended for public distribution.
- Confidential information — contracts, pricing, strategy, employee records, customer information, and other restricted business data.
- Highly sensitive or regulated information — authentication secrets, financial data, protected personal information, controlled government information, or other data subject to legal, contractual, or regulatory safeguards.
- Intellectual property — designs, source code, methods, research, proposals, trade secrets, and proprietary knowledge.

Leadership Question: If someone copied every file your organization stores tonight, which five types of information would concern you most tomorrow morning? Those are strong candidates for your highest data-protection priorities.

2.5 Confidentiality, Integrity, and Availability

Cybersecurity professionals often describe information protection using three simple goals: confidentiality, integrity, and availability. These concepts are useful because they show that protecting data means more than preventing someone from reading it.

Confidentiality means information is available only to people or systems authorized to see it. Integrity means information remains accurate and is not changed improperly. Availability means information and systems are accessible when the organization needs them.

Consider payroll. Confidentiality protects employee compensation and banking information. Integrity prevents an attacker from changing direct-deposit details.

Availability ensures payroll can actually be processed on time. The same asset can therefore require all three forms of protection.

A public website may have low confidentiality requirements but high integrity and availability requirements. A confidential acquisition document may have extremely high confidentiality requirements even if temporary unavailability would be tolerable. Understanding the difference helps organizations spend money where it matters.

2.6 Systems, Applications, and Cloud Services

Organizations frequently know about the equipment they purchased but have a less complete picture of cloud services and software adopted over time. A department may subscribe to an online application using a corporate credit card. An employee may connect a productivity tool to Microsoft 365. A vendor may host a critical database. These services become part of the organization's technology environment whether or not the IT department originally selected them.

A useful inventory includes endpoints, servers, network equipment, websites, databases, cloud platforms, software-as-a-service applications, remote-access systems, backup systems, security platforms, and business applications. For each important system, identify an owner, purpose, location, administrator, dependency, and recovery requirement.

The goal is not to create an enormous spreadsheet nobody maintains. The goal is to know enough about the environment to make decisions when something changes or fails. NIST Cybersecurity Framework 2.0 similarly emphasizes inventories of hardware, software, services, systems, supplier services, and designated data, and prioritizing assets according to classification, criticality, resources, and mission impact.

2.7 Business Processes Are Assets Too

Some of the most important assets cannot be plugged into a wall. Consider the process for approving a new vendor, issuing a payment, onboarding an employee, responding to a customer, restoring a server, or submitting a government deliverable. These processes combine people, technology, information, and authority.

Attackers often exploit business processes rather than technical vulnerabilities. A fraudulent invoice may succeed because the payment process allows one person to change banking information without independent verification. A business email compromise may succeed because employees trust an urgent message that appears to come from an executive.

Documenting critical processes helps an organization identify where cybersecurity controls should exist. Sometimes the best security improvement is not another software product. It is a better approval process, a second verification step, or a clear separation of duties.

2.8 Third Parties and Vendors

Your organization may protect its own environment carefully and still depend on companies that can affect your security. Managed service providers, cloud vendors, payroll companies, accountants, software providers, contractors, consultants, and other partners may store your data or have access to your systems.

Vendor risk is therefore part of asset protection. The organization should know which third parties possess sensitive data, which have administrative or remote access, which provide mission-critical services, and what happens if one becomes unavailable.

Contracts should address appropriate security responsibilities where necessary, but paperwork alone is not enough. Access should be limited to what is needed, reviewed periodically, and removed when the relationship ends. A supplier inventory is also specifically recognized in NIST CSF 2.0 asset management.

2.9 Physical Assets Still Matter

Cloud computing has not eliminated physical security. Laptops are stolen. Mobile phones are lost. Network closets are left unlocked. Printed documents are discarded improperly. Backup media may be stored in locations that are not adequately protected.

An asset inventory should therefore include relevant physical technology and facilities. Smaller organizations do not need elaborate controls everywhere, but they should understand where critical equipment resides, who can access it, and what would happen if the location became unavailable because of fire, water damage, power loss, theft, or another event.

2.10 A Realistic Scenario: The File Server Nobody Thought Was Critical

A regional construction company believed its accounting platform was its most important system. Leadership invested in backups and support for that application. During a ransomware incident, however, a shared project drive became unavailable. That drive contained active project drawings, subcontractor documents, change orders, photographs, and correspondence used by project managers every day.

The accounting system remained operational, but field and office teams could not access the information needed to coordinate active projects. Work slowed immediately. Employees began searching old email attachments and personal downloads for copies. Different versions of documents appeared. Nobody was certain which drawing was current.

The organization had not lost only a file server. It had lost access to an operational process. The technical label had hidden the business value.

What Went Wrong: The organization identified important technology but did not connect that technology to the work employees performed every day. Criticality was determined by perception rather than by business impact.

What Should Have Happened: Leadership should have mapped the major business processes and the systems and information each process depended upon. Project delivery would then have revealed the shared drive as a high-priority recovery asset.

2.11 Determine What Is Critical

Not every asset deserves the same investment. Prioritization begins by asking what would happen if an asset were unavailable, exposed, altered, or destroyed. Consider financial impact, operational disruption, contractual obligations, legal or regulatory consequences, safety, customer impact, and reputational harm.

A simple classification can be enough for many smaller organizations. The labels matter less than the discussion used to assign them.

- Critical — loss or compromise could stop essential operations, cause major financial harm, create serious legal or contractual exposure, or materially affect customers.
- High — significant disruption or exposure would occur, but temporary workarounds may exist.
- Moderate — the asset supports operations but its loss would be manageable for a limited period.
- Low — the asset is useful but can be recreated, replaced, or temporarily lost with limited business impact.

Classification should not be performed by IT alone. Finance knows which systems affect cash flow. Operations knows what stops delivery. Human resources knows which records are sensitive. Sales knows what customer information matters. Executives understand contractual and reputational consequences.

The objective is not to declare everything critical. If everything is critical, nothing has truly been prioritized.

2.12 Impact Questions That Clarify Priority

- If this asset were unavailable for four hours, what would happen?
- What if it were unavailable for two days?
- What if the information became public?
- What if an attacker changed the information without our knowledge?
- What if the data were permanently destroyed?
- Would customers be affected?
- Would we violate a contract, law, regulation, or reporting obligation?
- Could we continue operating manually?
- How much would recovery cost?
- Who has the authority and knowledge to make decisions about this asset?

These questions convert abstract security discussions into operational consequences. Leaders may disagree about technical severity, but they usually understand the meaning of missed payroll, halted production, lost customer records, delayed contracts, or an inability to invoice.

2.13 Find Where the Data Actually Lives

An organization may believe customer documents are stored in one approved location when copies also exist in email, local Downloads folders, personal cloud drives, collaboration platforms, mobile devices, backup systems, and vendor portals. The more copies that exist, the harder it becomes to control access and deletion.

Data mapping does not need to begin as an expensive consulting exercise. Start with important information types and ask where each is created, stored, transmitted, backed up, and shared. Then identify who can access it and whether that access is still necessary.

This exercise often reveals shadow data — information stored outside the places leadership expects. It may also reveal that an organization is retaining sensitive information long after the business need has ended.

Example: A company may have a formal SharePoint location for contracts, but employees may also keep copies in email, desktop folders, Teams conversations, personal OneDrive folders, and downloaded PDF files. Protecting only the formal repository does not protect every copy.

2.14 Data Retention: More Is Not Always Better

Organizations often keep information because storage is inexpensive and deleting it feels risky. But retaining sensitive information indefinitely can increase the consequences of a breach. If the organization no longer has a legitimate business, legal, or contractual reason to keep certain data, continued retention may create unnecessary exposure.

Retention decisions should be coordinated with legal, regulatory, contractual, operational, and records-management requirements. Cybersecurity should not independently decide what must be deleted. Its role is to make the risk visible and ensure retained information receives appropriate protection.

2.15 Ownership: Every Critical Asset Needs a Responsible Person

A technical administrator is not necessarily the business owner of an asset. IT may operate an accounting server, but the finance leader understands the business requirements for the accounting process. The business owner should help define acceptable downtime, access requirements, retention needs, and recovery priorities.

Assigning ownership prevents a familiar problem: everyone assumes someone else is responsible. An asset owner does not need to perform technical maintenance. The owner needs to understand why the asset matters and participate in decisions about its protection.

Practical Rule: If nobody can answer who owns an important system, dataset, or business process, the organization probably has a governance gap.

2.16 The Asset Life Cycle

Assets change over time. Laptops are purchased and retired. Employees join and leave. Cloud services are adopted and abandoned. Vendors change. Applications are replaced. Data is created, copied, archived, and eventually destroyed.

Security therefore has to follow the asset throughout its life cycle. A secure laptop deployment can still create risk if the device is later discarded with recoverable data. A properly created user account can become dangerous if it remains active after employment ends. A vendor integration that was justified three years ago may become forgotten access today.

An inventory is useful only when it is maintained. Major organizational changes, acquisitions, new applications, new vendors, and changes in business processes should trigger updates.

2.17 Common Mistakes Organizations Make

- Buying security tools before identifying the organization's most important assets.
- Maintaining a hardware inventory while ignoring cloud services, data, identities, vendors, and business processes.
- Assuming every asset has the same security requirements.
- Allowing sensitive information to accumulate in email, local drives, and unapproved cloud applications.
- Failing to remove access for former employees, contractors, or vendors.
- Treating backups as proof of recoverability without knowing what is actually backed up.
- Allowing one person to be the only source of knowledge for a critical system or process.
- Failing to identify third parties that hold sensitive information or administrative access.
- Classifying everything as critical, which makes prioritization meaningless.
- Leaving asset inventories untouched until an audit or incident forces an update.

2.18 What This Means for Your Organization

Asset management is not busywork for the IT department. It is the foundation for nearly every security decision that follows. You cannot design meaningful access controls if you do not know which accounts and systems matter. You cannot create a recovery plan if you do not know which processes must return first. You cannot protect sensitive data if you do not know where it is stored.

For smaller organizations, perfection is not required. Start with the assets that would create the greatest harm if they failed or were compromised. Build enough visibility to make informed decisions. Then improve the inventory over time.

A one-page list of the ten most important business assets, with owners and recovery priorities, can be more valuable than a hundred-page inventory nobody trusts.

2.19 Questions to Ask Your IT Provider or Internal IT Team

- Do we have a current inventory of company-owned computers, servers, network devices, and mobile devices?
- Do we maintain an inventory of software, cloud services, and important SaaS applications?
- Which systems and accounts have administrative privileges?
- Where is our most sensitive information stored?

- Which third parties store our data or have remote or administrative access?
- Which five systems would we restore first after a major outage?
- Do our backup priorities match our business priorities?
- How quickly is access removed when an employee or contractor leaves?
- Do we know which applications employees are using outside the approved technology environment?
- Who owns each of our critical systems and datasets?
- When was the asset inventory last reviewed?
- What process ensures new technology and new vendors are added to the inventory?

2.20 TRG Practical Takeaway

You cannot protect what you do not know you have, and you cannot prioritize what you have not connected to the business. Begin with visibility. Identify the people, identities, information, systems, vendors, and processes your organization depends upon. Then decide which ones matter most.

- Identify your critical business processes first.
- Map the systems, information, people, and vendors those processes depend upon.
- Create practical inventories of hardware, software, services, identities, suppliers, and important data.
- Assign a business owner to every critical asset.
- Prioritize assets by impact, not by price or technical complexity.
- Know where sensitive data is stored and how many copies exist.
- Review access when employees, contractors, vendors, and business needs change.
- Connect backup and recovery priorities to business criticality.
- Update inventories as the organization changes.

2.21 Before You Move On

- ☐ Can we name our five most important business processes?
- ☐ Can we name the systems and information each process depends upon?
- ☐ Do we know where our most sensitive information is stored?
- ☐ Do we know which employees and vendors have privileged access?
- ☐ Have we assigned owners to critical assets?
- ☐ Have we classified assets by business impact?
- ☐ Do our recovery priorities reflect those classifications?
- ☐ Do we know which cloud services and SaaS applications the organization uses?
- ☐ Do we have a process for removing obsolete accounts, systems, and vendor access?

- ☐ Is our inventory reviewed when the business changes?

2.22 Key Terms in Plain English

Asset: Anything of value to the organization that supports its mission or business operations, including people, data, identities, systems, services, facilities, and processes.

Asset inventory: A maintained record of important assets the organization owns, manages, uses, or depends upon.

Business owner: The person responsible for defining the business requirements and importance of an asset or process.

Confidentiality: Keeping information from people or systems that are not authorized to see it.

Integrity: Protecting information and systems from improper or unauthorized change.

Availability: Ensuring information and systems are accessible when authorized users need them.

Criticality: The degree to which an asset is essential to the organization's mission or operations.

Data classification: Grouping information according to sensitivity, business value, or protection requirements.

Shadow IT: Technology or services used without appropriate organizational visibility or approval.

Shadow data: Copies of information stored outside expected or approved locations.

Privileged account: An account with elevated authority to administer systems, access sensitive information, or perform high-impact actions.

Supplier risk: Cybersecurity risk created through vendors, contractors, service providers, and other external dependencies.

2.23 Leadership Exercise: Build Your First Critical-Asset List

Before reading the next chapter, gather a small group representing leadership, operations, finance, IT, and any department that handles particularly sensitive information. Give the group 30 minutes. Do not begin by asking for a list of servers.

Instead, ask each person to write down the five business activities the organization could least afford to lose. Compare the lists. Then identify the people, information, systems, applications, vendors, and facilities required for each activity.

For every resulting asset, record four items: a business owner, the primary business purpose, the impact if unavailable or compromised, and the desired recovery priority. Do not worry about creating a perfect enterprise inventory during the first meeting.

The exercise will often reveal disagreements. Those disagreements are useful. If finance believes the accounting platform must be restored first while operations believes the project-management system is more important, leadership has discovered a recovery-priority question before an emergency forces the decision.

Repeat the exercise periodically and whenever the organization changes significantly. Asset knowledge is not a one-time project; it is part of managing the business.

2.24 Connection to NIST Cybersecurity Framework 2.0

The approach in this chapter aligns closely with the Asset Management category of the NIST Cybersecurity Framework 2.0 Identify function. NIST describes assets broadly, including data, hardware, software, systems, facilities, services, and people, and emphasizes managing them according to their importance to organizational objectives and risk strategy.

The framework calls for inventories of hardware; software, services, and systems; supplier-provided services; and designated data. It also calls for assets to be prioritized based on classification, criticality, resources, and mission impact, and for assets to be managed throughout their life cycles.

The value of this guidance is not that every small organization must build a complex compliance program. The value is the underlying discipline: know what exists, know why it matters, know who depends on it, and manage it throughout its useful life.

2.25 Chapter Summary

Cybersecurity begins with knowing what the organization is trying to protect. Assets include far more than computers. They include people and identities, information, systems, cloud services, business processes, vendors, facilities, intellectual property, and trust.

Effective protection requires more than an inventory. Assets must be connected to business purpose and prioritized according to the consequences of loss, exposure, alteration, or unavailability. Sensitive data must be located, critical processes must be understood, owners must be assigned, and third-party dependencies must be visible.

The objective is not perfect documentation. The objective is better decisions. When leaders know what matters most, security spending becomes easier to prioritize, recovery plans become more realistic, and technical controls can be aligned with actual business risk.

In the next chapter, we will build on this foundation by examining how cybercriminals actually get in — through people, passwords, email, vulnerabilities, exposed services, vendors, and other paths that connect directly to the assets identified here.

Chapter 2 Reference Note

Primary standards reference: National Institute of Standards and Technology (NIST), The NIST Cybersecurity Framework (CSF) 2.0, especially the Identify (ID) function and Asset Management (ID.AM) category. The framework is used here as authoritative guidance for asset identification, inventories, prioritization, supplier-service awareness, data inventories, and asset life-cycle management.

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

Chapter 3

How Cybercriminals Actually Get In

Authored by Dr. Charles Edwards

Most cyberattacks do not begin with some extraordinary technical maneuver that only a highly trained hacker could understand. In many cases, the attacker finds one ordinary weakness and uses it as the doorway into the organization.

It may be a convincing email. A password stolen somewhere else. An old server that has not been patched. A remote-access system exposed to the Internet. Or a trusted vendor account that has been compromised.

Attackers usually do not need to defeat every security control. They need one workable path inside.

Once they have that first foothold, the situation can change quickly. One compromised mailbox can lead to password resets, fraudulent invoices, stolen information, or deeper access into the environment.

In a ransomware incident, the attacker may spend days or weeks learning the organization before anyone sees the visible damage. That is why understanding the first point of entry matters so much.

3.1 The Initial Access Problem

Cybersecurity professionals often use the term initial access to describe the first successful entry an attacker gains into a computer system, network, application, or online account.

Think of a business building with twenty doors. Nineteen doors may have excellent locks. But if one loading-dock door is accidentally left unlocked, the burglar does not need to defeat the other nineteen locks.

Cybersecurity works in much the same way. An organization may have firewalls, antivirus software, cloud security, encrypted laptops, strong backup systems, security policies, and an information technology staff. But an attacker might still enter because one employee's password was stolen through a phishing message.

The lesson is not that security controls are useless. They are extremely important. The lesson is that cybersecurity must be approached in layers.

If one security measure fails, another should make it difficult for the attacker to continue. This concept is often called defense in depth.

3.2 Phishing: Still One of the Most Effective Doors

Phishing is one of the most common ways cybercriminals attempt to gain access to organizations.

A phishing message is designed to convince someone to take an action that benefits the attacker.

- clicking a malicious link
- opening an infected attachment
- entering a username and password into a fake website
- approving a fraudulent payment
- providing confidential information
- installing software
- approving an unexpected multifactor authentication request

3.3 Social Engineering: Attacking People Instead of Computers

Phishing is part of a larger category of attacks known as social engineering. Social engineering means manipulating people into doing something they normally would not do.

Instead of attacking a firewall, the criminal attacks trust. Instead of breaking encryption, the criminal creates urgency. Instead of defeating a password system, the criminal convinces someone to reveal the password.

Cybercriminals frequently use urgency, fear, authority, curiosity, opportunity, and helpfulness. Employees should therefore be taught more than simply not to click suspicious links. They should learn to recognize the behavioral techniques attackers use.

3.4 Business Email Compromise

One particularly damaging form of social engineering is called Business Email Compromise, often abbreviated BEC.

In a BEC attack, the criminal impersonates or takes control of a trusted business email account. The attacker might pretend to be the CEO, a company executive, the accounting department, a supplier, an attorney, a customer, or another trusted employee.

The message may request a wire transfer, a change in bank account information, payment of an invoice, employee tax information, payroll changes, gift card purchases, or confidential documents.

A good cybersecurity process should make it difficult for one email to authorize a significant financial transaction. For example, the company might require verbal confirmation using a previously known telephone number before banking information or payment instructions can be changed.

The important principle is that security should not depend entirely on whether one employee recognizes a sophisticated scam. The business process itself should provide protection.

3.5 Stolen Passwords

Another major entry point is the stolen password.

People often imagine that criminals obtain passwords by sitting at a computer and trying millions of combinations. That does happen, but criminals frequently obtain passwords in easier ways.

- phishing
- malware
- data breaches
- fake login pages
- password-stealing browser extensions
- compromised computers
- insecure password storage
- reused credentials from another service

3.6 Why Password Reuse Is So Dangerous

Password reuse creates a chain reaction. If one password is used for several accounts and one of those services is breached, the other accounts may also become vulnerable.

This is why every important account should have a unique password. A reputable password manager can make this practical by generating and storing long, unique passwords.

3.7 Multifactor Authentication Changes the Equation

Multifactor authentication, usually called MFA, adds another layer of protection. Instead of requiring only something you know, such as a password, the system requires an additional form of verification.

MFA may use an authentication application, a security key, a device-based approval, a biometric factor, or another verification method.

MFA is extremely valuable because a stolen password by itself may no longer be enough. It should be enabled wherever practical, especially for email, administrative accounts, remote access, financial systems, cloud services, and other sensitive applications.

3.8 MFA Fatigue

One technique criminals use is sometimes called MFA fatigue or MFA bombing. The attacker already has the victim's password and repeatedly attempts to log in. Each attempt causes the employee's phone to display an approval request.

Eventually the employee may approve one simply to make the notifications stop.

Employees need to understand a simple rule: If you did not initiate the login, do not approve the authentication request. Unexpected MFA requests should also be reported because they may indicate that someone already knows the employee's password.

3.9 Unpatched Software

Not every attack depends on tricking a person. Sometimes criminals simply find vulnerable technology.

Software contains defects, and some defects are security vulnerabilities. When manufacturers discover these vulnerabilities, they often release updates or patches to correct them.

The problem occurs when organizations do not install those updates. Once information about a vulnerability becomes public, attackers may begin searching the Internet for organizations that have not patched it.

That is why patch management is not merely an IT housekeeping task. It is a cybersecurity control.

3.10 Internet-Exposed Services

Businesses increasingly need employees, vendors, and administrators to access systems remotely. Remote access is useful, but it can also create another doorway.

Attackers routinely search the Internet looking for systems that expose remote desktop, VPN gateways, administrative portals, file-transfer systems, web applications, databases, and other management interfaces.

An exposed service is not automatically insecure. The danger occurs when it is combined with weaknesses such as default passwords, weak passwords, outdated software, missing MFA, poor configuration, unnecessary Internet exposure, or excessive user privileges.

A fundamental question every organization should periodically ask is: What systems can someone reach directly from the Internet?

3.11 Malware and Malicious Downloads

Another entry method is malicious software, commonly called malware.

Malware can arrive through email attachments, fake software updates, compromised websites, pirated software, malicious advertisements, infected documents, USB devices, or software downloaded from untrusted sources.

Malware may steal passwords, record keystrokes, install remote-access software, encrypt files, capture screenshots, search for financial information, disable security tools, or download additional malware.

Modern attacks frequently involve several stages. The first malicious program may simply establish access so additional tools can be installed later.

3.12 Third-Party and Vendor Access

Organizations do not operate in isolation. They depend on managed service providers, software companies, accountants, payment processors, cloud providers, consultants, contractors, suppliers, and other partners.

Every connection between organizations creates potential risk. If a provider's account is compromised, an attacker may be able to reach the company through that trusted relationship.

This is sometimes called third-party risk or supply-chain risk.

The lesson is not to stop working with vendors. The lesson is to understand who has access to your systems, what level of access they have, and how that access is protected.

3.13 Human Error Does Not Mean Human Failure

Employees are sometimes described as the weakest link in cybersecurity. That description can be misleading.

People make mistakes, but poorly designed systems often turn small mistakes into major incidents. An employee clicking one bad link should not automatically allow an attacker to take control of the entire organization.

A stronger cybersecurity environment assumes that mistakes will occasionally happen. It therefore uses layers such as MFA, email filtering, endpoint protection, network segmentation, limited user privileges, backups, monitoring, patching, and security awareness training.

The goal is not to create perfect employees. The goal is to build an organization in which one mistake does not become a catastrophe.

3.14 From One Account to the Entire Network

Initial access is often only the beginning. Once attackers gain entry, they may try to understand the environment, steal additional credentials, obtain higher privileges, locate valuable systems, move to other computers, disable defenses, steal information, and prepare the final attack.

This movement from one system to another is commonly called lateral movement.

A typical progression may look like this: phishing email -> employee account compromised -> attacker accesses workstation -> additional credentials discovered -> administrator account compromised -> file server reached -> sensitive information stolen -> ransomware deployed.

Notice how the original incident may have begun with something that appeared relatively minor: one employee entered a password on the wrong website. The final result can be an organization-wide emergency.

3.15 Privilege Escalation: Getting More Power

When attackers first enter a system, they may not immediately have complete control. They might begin with the account of an ordinary employee.

That account may allow them to read email, access shared files, use business applications, connect to cloud services, and communicate with other employees. But the attacker may want much more.

The process of gaining greater levels of access is called privilege escalation.

Think of an office building. A visitor badge may let someone enter the lobby and several meeting rooms. An employee badge may provide access to office floors. A master access card may open almost every door in the building. Cybercriminals want the digital equivalent of that master access card.

An administrator account may allow an attacker to create new accounts, reset passwords, install software, disable security controls, access confidential information, modify system settings, and control multiple computers.

This is why administrator accounts require additional protection. Employees should not routinely use administrative privileges for ordinary work unless those privileges are actually necessary.

The security principle behind this idea is called least privilege. Least privilege means giving people only the access they need to perform their responsibilities - and no more. This limits the damage that can occur if an account is compromised.

3.16 Lateral Movement: Moving Through the Organization

After gaining access to one computer or account, attackers often begin exploring. They want to understand the environment.

They may ask what other computers are available, which servers the user can reach, where important files are stored, which accounts have administrative privileges, whether the account can access financial systems, whether backup systems are available, and whether credentials are stored somewhere on the computer.

The process of moving from one system to another is called lateral movement.

This is one reason network segmentation is important. If every system can communicate freely with every other system, compromising one computer may provide a path to many more.

Network segmentation does not guarantee that an attacker will be stopped. It creates barriers. Every additional barrier gives defenders another opportunity to detect and stop the attack.

3.17 Persistence: Making Sure the Attacker Can Come Back

Cybercriminals know that their access may eventually be discovered. An employee might change a password, a computer might restart, security software might remove malware, or an administrator might disable a suspicious account.

For that reason, attackers often try to establish persistence. Persistence means creating a way to regain access even if the original entry point disappears.

Attackers may attempt to create additional user accounts, install remote-access software, add malicious scheduled tasks, modify startup settings, create hidden administrator accounts, steal additional credentials, create new authentication tokens, or compromise multiple computers.

Removing the obvious symptom does not necessarily remove the attacker. Organizations must investigate how far the attacker went, what systems were accessed, what credentials may have been stolen, and what persistence mechanisms may have been created.

3.18 Attackers Usually Want Time

Many people imagine that cyberattacks happen immediately. In reality, sophisticated attacks may unfold much more slowly.

An attacker may spend days or weeks inside an environment before launching the final attack. During that time, the criminal may study the network, identify important employees, learn business processes, collect passwords, locate sensitive information, identify backup systems, determine when administrators are less likely to notice unusual activity, and decide how to cause maximum damage.

This period is dangerous because the organization may appear completely normal. This is one reason cybersecurity monitoring is important.

3.19 Data Theft Often Happens Before Ransomware

Modern ransomware operations are often more complex than simply encrypting files. Attackers may first steal large quantities of information. This process is called data exfiltration.

The attacker may copy customer records, employee information, contracts, financial documents, intellectual property, medical information, credentials, email archives, business plans, legal documents, and other sensitive data.

Only after stealing the information does the attacker encrypt systems. The criminal may then threaten the organization in two ways: pay us or your files remain encrypted; pay us or we will publish the information we stole.

This technique is sometimes called double extortion.

A good backup protects availability, but it does not undo stolen data. Backups are extremely important, but they are not a complete ransomware strategy.

3.20 Why Backups Are Still Essential

Although backups do not solve every cybersecurity problem, they remain one of the most important protections against destructive attacks.

A reliable backup can help an organization recover from ransomware, hardware failure, accidental deletion, corrupted files, system failure, and certain natural disasters.

But simply saying, “We have backups,” is not enough. An organization needs to know what is being backed up, how often, where backups are stored, whether they are encrypted, whether attackers can access them, how long they are retained, whether restoration has been tested, and how long recovery would take.

A backup that has never been tested is only an assumption.

3.21 Attackers Target Backups Too

Experienced cybercriminals understand that backups can allow organizations to recover without paying a ransom. Therefore, attackers may deliberately search for backup systems.

If they gain administrative access, they may attempt to delete backups, encrypt backup files, disable backup software, erase recovery snapshots, destroy cloud backups, or modify retention policies.

This is why backup systems should not be treated like ordinary file storage. A stronger backup strategy includes separate administrative credentials, restricted access, immutable backups when appropriate, offline or isolated copies, monitoring, and multifactor authentication.

The objective is simple: an attacker who compromises the production network should not automatically gain the ability to destroy every recovery copy.

3.22 A Complete Attack Scenario

Consider a fictional company called North Valley Manufacturing. North Valley employs approximately 85 people. The company has Microsoft 365, several Windows servers, accounting software, a file server, cloud backups, remote-access capabilities, and an outside IT provider.

Monday - 8:43 AM: Jennifer, an accounts-payable employee, receives an email that appears to come from Microsoft. The subject says, “Immediate Action Required: Mailbox Storage Limit Reached.” She clicks the link and enters her username and password into a Microsoft-looking page.

Monday - 9:10 AM: The attacker attempts to access Jennifer's Microsoft 365 account. The company uses MFA. After several repeated approval prompts, Jennifer eventually taps Approve without carefully reading it. The attacker is now inside the account.

Monday - 11:30 AM: The attacker searches Jennifer's mailbox and learns the name of the company president, the bank the business uses, the accounting manager, how invoices are normally approved, and which vendors receive large payments.

Tuesday: The attacker creates a mailbox rule that moves messages containing words such as wire, fraud, suspicious, bank, and payment into a hidden folder.

Wednesday: Using Jennifer's account, the attacker obtains a legitimate vendor invoice and changes the bank account information.

Thursday: The attacker discovers an old IT troubleshooting document containing a service account password that should have been removed years earlier.

Friday: The criminal uses the newly discovered credentials to access a server. The account has more privileges than necessary. The attacker begins moving laterally through the environment.

The following week: The attacker identifies the primary file server, several administrative accounts, the backup system, and a remote management tool used by the IT provider. Security logs contain unusual activity, but nobody is actively reviewing them.

Wednesday - 2:00 AM: The attacker begins copying sensitive files outside the organization. Several hundred gigabytes of information are transferred. No alert is generated.

Thursday - 1:15 AM: The attacker attempts to disable security software. Backup snapshots are deleted. Several servers are encrypted. Employee workstations begin displaying a ransom message.

Thursday - 7:05 AM: Employees cannot open files. By 7:30 AM, the organization realizes it has a serious incident. Operations stop, accounting cannot process payments, production documents are unavailable, shared files cannot be accessed, and customers begin calling.

What originally began with one phishing email has now become an organization-wide crisis.

3.23 The Attack Was Not Caused by One Failure

It would be easy to blame Jennifer. She clicked a phishing link, entered her password, and approved an MFA request. But doing so would miss the larger cybersecurity lesson.

Several weaknesses contributed to the incident:

1. The phishing email reached the employee.
2. The employee entered credentials into a fraudulent website.
3. The MFA system allowed simple push approval without additional verification.
4. Security awareness training did not adequately address MFA fatigue.
5. An old document contained a password.
6. The service account had excessive privileges.
7. The password had not been rotated.
8. Network segmentation was insufficient.
9. Security logs were not actively monitored.
10. Data exfiltration was not detected.
11. The attacker was able to reach backup systems.
12. Backup deletion was possible.
13. The organization did not detect the attacker during the earlier stages of the intrusion.

A mature cybersecurity program assumes that some controls will occasionally fail. The goal is to prevent those failures from becoming catastrophic.

3.24 Cybersecurity Is a System, Not a Product

Organizations sometimes approach cybersecurity as a shopping problem. They ask: What software should we buy?

Security technology matters, but cybersecurity is not created by purchasing one product.

An organization may buy an expensive firewall, endpoint protection, email security, cloud monitoring, vulnerability scanning, and backup software, yet still remain vulnerable if nobody monitors alerts, administrator accounts use weak passwords, employees share credentials, software is not patched, backups are never tested, terminated employees retain access, or vendors have excessive privileges.

Technology is one part of cybersecurity. People and processes matter just as much.

A strong security program combines People + Process + Technology. All three must work together.

3.25 Questions Every Business Owner Should Ask

About Email

- Does our email system use multifactor authentication?
- Are phishing protections enabled?
- Are suspicious attachments scanned?
- Are employees trained to recognize impersonation attempts?
- Do we have a procedure for reporting suspicious messages?

About Passwords

- Are employees required to use unique passwords?
- Do we use a password manager?
- Are administrator passwords protected differently?
- Are old or unused accounts removed?
- Are default passwords changed?

About Remote Access

- What systems can be accessed remotely?
- Is MFA required?

- Are remote-access systems patched?
- Who is authorized to use them?
- Are failed login attempts monitored?

About Software Updates

- Who is responsible for patching?
- How quickly are critical vulnerabilities addressed?
- Are servers, workstations, network devices, and applications all included?
- Do we know which systems are no longer supported by the manufacturer?

About Backups

- What systems are backed up?
- How frequently?
- Can ransomware reach the backups?
- Are there isolated or immutable backup copies?
- When was the last successful restoration test?

About Vendors

- Which vendors can remotely access our systems?
- What level of access do they have?
- Is MFA required?
- Are vendor accounts disabled when not needed?
- Do we periodically review vendor access?

About Monitoring

- Who receives security alerts?
- Who reviews them?
- What happens outside business hours?
- Can we identify unusual login activity?
- Would we know if large amounts of data were leaving the organization?

3.26 Questions to Ask Your IT Provider

1. Is MFA enabled for every administrator account?
2. How are privileged credentials protected?
3. Do your technicians share accounts or have individual accounts?
4. Can your remote management tools reach all of our systems?

5. How do you secure your own remote-access platform?
6. How quickly are critical security patches applied?
7. How do you monitor suspicious login activity?
8. Who receives security alerts after business hours?
9. Can our backups be deleted from an ordinary administrator account?
10. When did you last test restoration of our systems?
11. What happens if one of your technician accounts is compromised?
12. How quickly would you notify us of a security incident?
13. Do you maintain an incident-response plan?
14. Are our systems segmented to limit attacker movement?
15. Can you provide us with a current inventory of systems and accounts you manage?

3.27 Practical Safeguards

Protect Accounts: Use multifactor authentication wherever possible, especially for email, administrative accounts, financial systems, remote access, cloud platforms, and backup systems.

Use Unique Passwords: Do not reuse business passwords across multiple services. Use a reputable password manager.

Limit Administrative Access: Employees should not have administrator privileges unless required. Separate ordinary user accounts from administrative accounts where practical.

Patch Systems: Develop a documented process for installing security updates. Prioritize critical vulnerabilities, especially on Internet-facing systems.

Secure Remote Access: Require MFA, remove unnecessary exposure, and monitor suspicious login activity.

Segment the Network: Do not allow one compromised computer to communicate freely with every critical system.

Protect Backups: Maintain multiple backup copies, restrict administrative access, use isolated or immutable storage where appropriate, and test restoration regularly.

Monitor the Environment: Collect and review security logs. Investigate unusual behavior and create alerts for high-risk events.

Train Employees: Training should include realistic examples of phishing, business email compromise, MFA fatigue, fraudulent payment requests, password theft, and impersonation.

Verify Financial Changes: Bank account changes and significant payment requests should be verified through a second communication method. Never rely solely on email.

Review Vendor Access: Know which outside organizations have access, remove unnecessary accounts, and require appropriate security controls.

3.28 Chapter 3 Cybersecurity Checklist

☐ MFA is enabled for email.
☐ MFA is enabled for administrators.
☐ MFA is enabled for remote access.
☐ Employees use unique passwords.
☐ A password manager is available.
☐ Default passwords have been removed.
☐ Old accounts are regularly disabled.
☐ Employees receive phishing-awareness training.
☐ Employees know how to report suspicious email.
☐ Financial requests require verification.
☐ Bank-account changes require secondary confirmation.
☐ MFA fatigue is included in security training.
☐ Operating systems receive security updates.
☐ Business applications are patched.
☐ Firewalls and network appliances are patched.
☐ Unsupported software is identified.
☐ Internet-facing systems are reviewed.
☐ Users receive only the access they need.
☐ Administrator accounts are limited.
☐ Privileged accounts are periodically reviewed.
☐ Service accounts are documented.
☐ Old passwords are not stored in documents.
☐ Remote access requires MFA.

☐ Unused remote-access services are disabled.
☐ Vendor access is controlled.
☐ Remote login activity is monitored.
☐ Critical systems are backed up.
☐ Backup jobs are monitored.
☐ Backups are protected from ordinary administrator accounts.
☐ At least one backup copy is isolated or protected from alteration.
☐ Restoration tests are performed.
☐ Security alerts have a responsible owner.
☐ Critical alerts are reviewed promptly.
☐ Failed login attempts are monitored.
☐ Administrative changes are logged.
☐ Suspicious outbound data transfers can be detected.
☐ Vendor accounts are documented.
☐ Vendor access is limited.
☐ Vendor MFA is required where appropriate.
☐ Unused vendor accounts are disabled.

3.29 Key Terms

Business Email Compromise (BEC): A form of fraud in which criminals impersonate or compromise a trusted business email account to manipulate victims into transferring money, changing payment information, or providing sensitive information.

Credential Stuffing: The practice of using stolen usernames and passwords from one service to attempt access to other services.

Data Exfiltration: The unauthorized transfer of information from an organization's systems to an external location.

Defense in Depth: A cybersecurity strategy that uses multiple protective layers so that the failure of one control does not automatically result in complete compromise.

Initial Access: The first successful entry an attacker gains into an organization's account, computer, application, or network.

Lateral Movement: The process by which an attacker moves from one compromised system or account to others within an organization.

Least Privilege: The principle that users and systems should receive only the access necessary to perform their responsibilities.

Malware: Malicious software designed to damage systems, steal information, provide unauthorized access, or perform other harmful activities.

Multifactor Authentication (MFA): An authentication method requiring more than one form of verification.

MFA Fatigue: A social-engineering technique in which attackers repeatedly generate authentication requests in hopes that the victim will eventually approve one.

Network Segmentation: Dividing a network into controlled sections to limit communication and reduce the ability of an attacker to move freely.

Persistence: Techniques attackers use to maintain or regain unauthorized access.

Phishing: Fraudulent communication designed to trick a victim into revealing information, clicking a malicious link, opening an attachment, or taking another unsafe action.

Privilege Escalation: The process of gaining access rights beyond those originally obtained.

Ransomware: Malware or an attack operation designed to deny access to systems or information and demand payment.

Social Engineering: The manipulation of people into revealing information or performing actions that benefit an attacker.

Third-Party Risk: Cybersecurity risk introduced through vendors, contractors, suppliers, service providers, and other outside organizations.

3.30 Key Takeaways

Cybercriminals rarely need to defeat every security control in an organization. They need one path inside.

That path may be a phishing email, a stolen password, an MFA approval, an unpatched application, an exposed remote-access service, malicious software, or a compromised vendor.

Once inside, attackers may attempt to increase their privileges, move throughout the environment, establish persistence, steal information, destroy backups, and eventually launch ransomware or another damaging attack.

This is why cybersecurity must be layered. Strong organizations assume that individual controls can fail and prepare for that possibility.

The objective is not perfection. The objective is resilience.

3.31 Chapter Summary

Cyberattacks often begin in ordinary ways: a legitimate-looking email, a reused password, an outdated server, a remote-access service, an employee trying to be helpful, a vendor account, or a simple configuration mistake.

What makes these attacks dangerous is what happens after that first step.

A typical progression is: Initial Access -> Credential Theft -> Privilege Escalation -> Lateral Movement -> Persistence -> Data Discovery -> Data Exfiltration -> Disruption or Ransomware.

Understanding this sequence changes the way organizations think about cybersecurity. The objective should not simply be preventing anyone from ever getting in.

Organizations must also ask: If someone does get in, how quickly can we detect it? How far can that person go? What systems can they reach? Can they become an administrator? Can they reach our backups? Can they steal our information? Can we recover?

A secure organization is not one that assumes attacks will never happen. A secure organization is one that makes attacks difficult, limits their impact, detects them quickly, and has a realistic plan for recovery.

And that brings us to the next major question: If attackers are constantly looking for weaknesses, how do we determine where our own weaknesses are before they do?

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

Chapter 4

Finding Your Weaknesses Before the Attackers Do

Authored by Dr. Charles Edwards

Most organizations know they have cybersecurity weaknesses.

The difficult part is identifying them before someone else does.

A business may believe its systems are secure because nothing bad has happened yet. Employees are working. Email is functioning. Customers can access the website. Backups are running. The firewall is online.

Everything appears normal.

But looking normal is not the same as being secure.

An attacker sees the organization differently.

The attacker is looking for one outdated system, one forgotten account, one exposed portal, one reused password, one bad configuration, or one vendor connection that provides a way in.

Cybersecurity therefore requires more than simply putting defenses in place.

Organizations need to examine themselves periodically from the perspective of someone trying to find a weakness.

That is the purpose of vulnerability management, assessments, configuration reviews, penetration testing, and other forms of security evaluation.

The goal is not to prove that the organization is perfect. No organization is perfect.

The goal is to find the important weaknesses while there is still time to correct them.

4.1 Every Organization Has Vulnerabilities

A vulnerability is a weakness that could potentially be used to compromise a system, application, account, process, or organization.

Some vulnerabilities are technical.

- outdated software
- missing security patches
- insecure network services
- weak encryption
- default passwords
- incorrect firewall rules
- insecure cloud settings
- poorly configured applications

4.2 Why Weaknesses Accumulate

Many vulnerabilities do not appear because someone deliberately ignored security. They accumulate gradually.

A new application is installed. A temporary firewall rule is created. An employee changes departments. A consultant receives administrative access. A server is placed online for a six-month project. A cloud account is created for testing. A vendor needs temporary remote access.

Then time passes. The project ends. The employee leaves. The consultant's work is completed. But some of the access remains.

Temporary configurations have a way of becoming permanent. This creates what might be called security drift.

Security drift occurs when systems slowly move away from their original secure configuration.

Nothing necessarily went wrong on any single day. The risk accumulated over time. That is why periodic review is necessary.

4.3 The Difference Between a Vulnerability and a Risk

The words vulnerability and risk are often used as though they mean the same thing. They are related, but they are not identical.

A vulnerability is a weakness. Risk considers what might happen if that weakness is exploited.

Not every vulnerability deserves the same priority. Organizations need a way to identify what matters most.

4.4 Think in Terms of Likelihood and Impact

A practical way to think about cybersecurity risk is to consider two questions: How likely is this weakness to be exploited? What would happen if it were exploited?

Organizations should avoid managing vulnerabilities entirely by counting them. A report showing 2,000 vulnerabilities sounds frightening, but leadership needs to know which vulnerabilities create immediate danger, which systems are involved, which ones are exposed to attackers, which ones affect critical business functions, which ones have known exploits, and which ones could cause serious operational damage.

Numbers without context are not enough.

4.5 Asset Inventory Comes First

You cannot protect what you do not know you have.

Before an organization can effectively manage vulnerabilities, it needs an accurate inventory of its technology assets.

- desktop computers
- laptops
- servers
- network switches
- firewalls
- wireless access points
- printers
- mobile devices
- cloud systems
- websites
- business applications
- databases
- virtual machines
- Internet-facing services
- backup systems

4.6 What Is the Attack Surface?

The term attack surface describes the collection of possible entry points an attacker might attempt to use.

For a small organization, the attack surface might include the public website, employee email accounts, remote-access systems, cloud applications, laptops, mobile devices, wireless networks, vendor connections, and Internet-facing servers.

Reducing attack surface is therefore a valuable security strategy.

If a service is not needed, disable it. If an account is no longer required, remove it. If a server should not be accessible from the Internet, do not expose it. If an application is obsolete, retire it.

Every unnecessary system creates something additional that must be protected.

4.7 External Exposure: What Can the Internet See?

One of the first perspectives an organization should examine is its external exposure.

Ask: What can someone on the Internet discover about us?

An attacker may search for public websites, login portals, VPN systems, remote desktop services, email servers, cloud storage, exposed databases, old web applications, test systems, forgotten subdomains, and administrative interfaces.

Organizations are sometimes surprised to discover systems they did not realize were still publicly accessible.

This is why periodic external exposure reviews are important. Businesses should know what the Internet can see.

4.8 Vulnerability Scanning

One of the most common tools used to identify technical weaknesses is a vulnerability scanner.

A vulnerability scanner examines systems for known security issues.

Depending on the tool and configuration, it may identify missing security patches, outdated software, known vulnerabilities, open network ports, weak cryptographic protocols, default configurations, expired certificates, insecure services, and certain configuration problems.

The scanner then produces a report. But vulnerability scanning is only the beginning.

Scanning creates information. Vulnerability management turns that information into action.

4.9 A Scanner Is Not a Security Program

Organizations sometimes purchase a vulnerability-scanning product and assume the problem has been solved.

But a scanner that produces reports nobody acts on provides little protection.

The real process should look like this: Discover systems -> Scan for vulnerabilities -> Review findings -> Prioritize risk -> Assign responsibility -> Fix or mitigate the weakness -> Verify the correction -> Repeat.

Cybersecurity is not a one-time project. Systems change continuously, new vulnerabilities are discovered continuously, and the process must repeat.

4.10 What Does a Vulnerability Severity Score Mean?

Vulnerability reports often assign severity levels such as Critical, High, Medium, Low, or Informational. Many tools also use numerical scores.

These ratings are useful, but they should not be treated as automatic business decisions.

Organizations should combine technical severity with business context. Useful questions include whether the system is Internet-facing, whether exploitation is known to be occurring, whether the system contains sensitive information, whether the vulnerability can lead to administrator access, whether the affected system is critical to operations, and whether existing controls reduce the danger.

Cybersecurity prioritization requires judgment.

4.11 Known Exploited Vulnerabilities

Some vulnerabilities deserve urgent attention because attackers are already using them in real attacks.

If attackers are actively targeting a vulnerability in a system your organization exposes to the Internet, remediation may need to become an immediate priority.

This is especially important for VPN systems, firewalls, remote-access appliances, web servers, email systems, identity platforms, and other Internet-facing infrastructure.

Attackers frequently automate this work. They do not need to know your company personally. Their software may simply scan thousands or millions of Internet addresses looking for vulnerable systems.

4.12 Security Configuration Reviews

Not all weaknesses are caused by missing patches. Many serious problems occur because systems are configured incorrectly.

A configuration review examines how a system has been set up.

Examples include checking password settings, MFA requirements, administrator privileges, firewall rules, file permissions, logging, encryption settings, cloud access controls, remote-access settings, default accounts, network services, and backup permissions.

This is why organizations need both patch management and configuration management.

4.13 Default Settings Are Designed for Convenience

Manufacturers want products to be easy to install. That often means default configurations are designed for compatibility and convenience.

Security may require additional changes such as disabling unnecessary services, changing default administrator accounts, requiring MFA, increasing logging, limiting network access, enforcing stronger password rules, and restricting file-sharing permissions.

Organizations should not automatically assume that a product is securely configured simply because it works correctly.

4.14 Penetration Testing

A vulnerability scan searches for known weaknesses. A penetration test goes further.

In a penetration test, authorized security professionals attempt to determine whether weaknesses can actually be exploited.

The goal is to simulate aspects of a real attack in a controlled and authorized manner.

A penetration tester might attempt to exploit vulnerable systems, gain access to accounts, bypass security controls, move from one system to another, obtain higher privileges, access sensitive information, or demonstrate how one weakness can be combined with another.

A scanner may report that a server has a vulnerability. A penetration tester may demonstrate that the vulnerability allowed access to the server, exposure of administrator credentials, and access to confidential records.

4.15 Penetration Testing Is Not About Embarrassing the IT Department

A good penetration test is not a contest between the tester and the IT staff.

The purpose is not to prove that someone failed. The purpose is to identify attack paths before a criminal finds them.

If a penetration tester successfully compromises a system, that is valuable information. The organization now has an opportunity to fix the weakness without experiencing an actual breach.

The healthiest attitude is: We want the authorized tester to find the problem first.

4.16 Vulnerability Scanning Versus Penetration Testing

Vulnerability scanning is usually automated, broad, repeated frequently, focused on known weaknesses, and useful for ongoing vulnerability management.

Penetration testing is usually more hands-on, performed by skilled testers, focused on whether weaknesses can be exploited, more limited in frequency, and designed to demonstrate real attack paths.

Organizations often need both.

4.17 Security Assessments

A security assessment is broader than a vulnerability scan.

It may examine technology, policies, employee practices, physical security, cloud systems, identity management, backup processes, incident response, vendor risk, and regulatory requirements.

A strong assessment should not simply produce a long list of technical issues. Leadership needs to understand what the biggest risks are, why they matter, what should be fixed first, what it will cost to fix them, what can wait, and what happens if nothing is done.

A useful assessment turns technical findings into business decisions.

4.18 The Problem With a 200-Page Security Report

Imagine a business owner receives a security assessment containing 200 pages of technical findings.

Technically, the report may be excellent. But if the owner asks what needs to be fixed first and nobody can answer clearly, the assessment has failed to communicate risk.

Executives do not need every technical detail removed. They need the technical detail translated into decision-making information.

A useful finding explains the weakness, the business meaning, the likely impact, the priority, and the recommended action.

4.19 Internal Versus External Assessments

Organizations should consider both outside and inside perspectives.

An external assessment examines the organization from the Internet and asks what an attacker could see before gaining access.

An internal assessment assumes the attacker has already obtained some level of access and asks what someone can do once inside.

Both perspectives matter.

4.20 Cloud Security Reviews

Businesses increasingly store critical information in cloud services.

Moving information to the cloud does not eliminate the organization's responsibility for security.

Cloud weaknesses may include missing MFA, excessive administrator privileges, public storage, overly broad sharing permissions, inactive accounts, risky third-party applications, weak conditional-access policies, and insufficient logging.

Cloud security reviews should therefore become part of normal cybersecurity management.

4.21 Identity Is Now Part of the Perimeter

For many years, businesses thought of the firewall as the boundary between trusted and untrusted systems. That model has changed.

Employees may now work from home, hotels, customer locations, airports, mobile devices, and personal networks. Applications may reside entirely in the cloud.

As a result, a user's identity has become one of the most important security boundaries.

Identity security is cybersecurity.

4.22 Dormant Accounts: The Forgotten Door

One of the simplest weaknesses to understand is the old account nobody removed.

Organizations need a clear offboarding process. When employees, contractors, or vendors leave, access should be reviewed and removed promptly.

An unused account should not remain a permanent doorway.

4.23 Excessive Privileges

Another common weakness occurs when people accumulate access over time.

This is sometimes called privilege creep.

Periodic access reviews help identify this problem. Managers should ask whether each person still needs the access they have.

If an account is compromised, the attacker inherits whatever access the user has.

4.24 The Importance of Configuration Baselines

A security baseline defines how a system should normally be configured.

For example, an organization might establish that all company laptops must have disk encryption enabled, endpoint security installed, automatic screen locking, current security updates, local firewall enabled, restricted administrator privileges, and approved software only.

Baselines create consistency.

4.25 Fixing Vulnerabilities Without Breaking the Business

Security teams cannot always install every update immediately.

A patch may affect an important application, a server may need to remain available during business hours, or a specialized system may require vendor approval before changes.

This does not mean the vulnerability should be ignored. It means risk must be managed intelligently.

Possible temporary safeguards might include restricting network access, disabling an affected service, increasing monitoring, blocking external access, adding firewall controls, limiting user privileges, or isolating the system.

These measures are sometimes called compensating controls.

4.26 Vulnerability Management Needs Ownership

One of the most common reasons vulnerabilities remain unresolved is surprisingly simple: nobody owns the problem.

Every significant finding should have an owner, a priority, a target date, a remediation plan, and a verification step.

Without ownership, reports become archives. With ownership, they become action plans.

4.27 What Should Be Fixed First?

When deciding priorities, consider whether the vulnerability is exposed to the Internet, whether attackers are actively exploiting it, whether it affects a critical business system, whether exploitation could provide administrator access, whether the system contains sensitive information, whether compromise could disrupt operations, whether strong existing controls reduce the risk, whether a fix is readily available, and what happens if remediation is delayed.

This produces a much better priority list than simply fixing vulnerabilities in numerical order.

4.28 Do Not Forget Printers, Cameras, and Other Devices

Cybersecurity inventories often focus on computers and servers. But modern networks contain many other devices.

Examples include printers, security cameras, door-access systems, conference room equipment, voice-over-IP phones, environmental controls, smart televisions, industrial devices, scanners, and Internet-connected appliances.

Every network-connected device should be considered part of the security environment.

4.29 Shadow IT

Employees sometimes begin using technology without formal approval.

This unofficial use of technology is often called shadow IT.

The answer is not always to prohibit everything. Employees usually adopt these tools because they are trying to get work done.

Organizations should understand the business need and provide secure alternatives.

4.30 A Practical Vulnerability Management Cycle

A strong vulnerability management process can be summarized in seven steps.

Step 1 - Know What You Have

Maintain an inventory of systems, applications, cloud services, and important devices.

Step 2 - Identify Weaknesses

Use vulnerability scans, configuration reviews, security assessments, vendor alerts, and other sources.

Step 3 - Understand the Risk

Determine how the weakness affects actual business operations.

Step 4 - Prioritize

Address the most dangerous weaknesses first.

Step 5 - Remediate or Mitigate

Patch, reconfigure, disable, isolate, replace, or otherwise reduce the risk.

Step 6 - Verify

Confirm that the weakness was actually corrected.

Step 7 - Repeat

Repeat the process because systems and threats continuously change.

4.31 A Business Scenario: The Forgotten Remote Server

Consider a fictional accounting firm called Harbor Financial Services.

The company employs 42 people. Several years ago, Harbor used an outside software consultant to develop a client-reporting application.

During development, the consultant created a remote server so programmers could access the application from home. The project was completed and the final application was moved elsewhere. Everyone assumed the development server had been retired.

It had not.

The server remained connected to the Internet. Its operating system had not been patched in more than two years. The consultant's old administrator account still existed.

No one at Harbor remembered the server. But attackers did not need to remember it.

An automated scan identified the outdated service. The attacker exploited the vulnerability and obtained access.

From that server, the attacker discovered credentials that allowed access to another internal system. Within several days, the attacker had reached sensitive client files.

Harbor's primary firewall was functioning correctly. Its antivirus software was working. Employees had completed cybersecurity awareness training. The breach still occurred.

Why? Because the organization did not know the forgotten server still existed.

The first failure was not patching. The first failure was visibility.

You cannot secure an asset you have forgotten.

4.32 What Harbor Could Have Done Differently

Several basic controls could have prevented or reduced the incident.

Harbor could have maintained an accurate asset inventory, periodically scanned its external Internet exposure, removed unused systems, disabled old consultant accounts, required MFA for remote administration, patched Internet-facing systems, segmented development systems from sensitive resources, and monitored unusual access.

Cybersecurity maturity often comes from consistently performing basic activities well.

4.33 Questions Every Business Owner Should Ask

Leadership should periodically ask practical questions about systems, vulnerabilities, cloud services, vendors, and testing.

About Systems

- Do we have an accurate list of our computers and servers?
- Do we know what systems are exposed to the Internet?
- Are any old systems still operating?
- Are unsupported systems documented?

About Vulnerabilities

- Do we perform vulnerability scans?
- How often?
- Who reviews the results?
- Who is responsible for remediation?
- How quickly are critical findings addressed?

About Cloud Systems

- Which cloud services contain company data?
- Is MFA required?

- Who has administrative access?
- Are old accounts disabled?
- Are sharing permissions reviewed?

About Vendors

- Which vendors can access our systems?
- Are old vendor accounts removed?
- Is vendor access monitored?
- Do vendors use MFA?

About Testing

- Have we had an independent security assessment?
- Have we performed a penetration test?
- When was the last one?
- Were the findings actually corrected?

4.34 Questions to Ask Your IT Provider

Ask your IT provider the following questions.

1. Do you maintain a current inventory of our managed systems?
2. Which of our systems are accessible from the Internet?
3. How often do you scan for vulnerabilities?
4. How are critical vulnerabilities prioritized?
5. Who reviews the findings?
6. How quickly are critical patches applied?
7. How do you handle systems that cannot be patched immediately?
8. Are unsupported systems identified and reported to us?
9. Do you review cloud security configurations?
10. Do you review user and administrator accounts?
11. **How are accounts for former employees and vendors removed?**
12. Have you identified any unnecessary Internet-facing services?
13. Do you conduct configuration reviews?
14. Do you recommend independent penetration testing?
15. How do you verify that vulnerabilities were actually fixed?

4.35 Chapter 4 Cybersecurity Checklist

Use this checklist as a practical review of the organization's asset inventory, vulnerability management, configuration security, accounts, testing, and cloud systems.

Asset Inventory

☐ We maintain an inventory of computers.
☐ We maintain an inventory of servers.
☐ Network devices are documented.
☐ Cloud services are documented.
☐ Internet-facing systems are identified.
☐ Unsupported systems are identified.
☐ Old or unused systems are removed.

Vulnerability Management

☐ Vulnerability scanning is performed regularly.
☐ Results are reviewed by a responsible person.
☐ Critical findings are prioritized.
☐ Each important finding has an owner.
☐ Remediation deadlines are established.
☐ Corrective actions are verified.

Configuration Security

☐ Default passwords are changed.
☐ Unnecessary services are disabled.
☐ Administrator privileges are restricted.
☐ Security logging is enabled.
☐ Firewall rules are periodically reviewed.
☐ Cloud permissions are periodically reviewed.

Accounts

☐ Former employee accounts are disabled promptly.
☐ Dormant accounts are reviewed.
☐ Vendor accounts are documented.
☐ Administrative accounts use MFA.
☐ Privileges are periodically reviewed.

Testing

☐ External exposure is periodically assessed.

☐ Internal security assessments are performed.

☐ Penetration testing is considered based on risk.

☐ Previous assessment findings are tracked to closure.

Cloud Systems

☐ MFA is enabled.

☐ Administrator accounts are limited.

☐ Sharing permissions are reviewed.

☐ Third-party applications are reviewed.

☐ Logging is enabled where appropriate.

4.36 Key Terms

Asset Inventory: A documented record of the hardware, software, cloud services, applications, and devices an organization uses.

Attack Surface: The collection of systems, accounts, applications, devices, and services that could potentially be targeted by an attacker.

Compensating Control: A temporary or alternative security measure used to reduce risk when the preferred corrective action cannot be implemented immediately.

Configuration Review: An examination of system settings to identify insecure or inappropriate configurations.

Penetration Test: An authorized security exercise in which testers attempt to exploit weaknesses and demonstrate what an attacker might be able to accomplish.

Privilege Creep: The gradual accumulation of unnecessary access permissions as an employee changes roles or responsibilities.

Security Assessment: A structured examination of an organization's security controls, practices, systems, and risks.

Security Baseline: A defined set of required security settings used as a standard for configuring systems.

Security Drift: The gradual movement of systems away from their intended secure configuration over time.

Shadow IT: Technology or cloud services used within an organization without formal approval or oversight.

Vulnerability: A weakness that could potentially be exploited to compromise a system, application, account, process, or organization.

Vulnerability Management: The ongoing process of discovering, evaluating, prioritizing, correcting, and verifying security weaknesses.

Vulnerability Scanner: A tool used to identify known security weaknesses and insecure configurations.

4.37 Key Takeaways

Every organization has weaknesses. The objective is not to pretend otherwise. The objective is to identify those weaknesses before criminals do.

Effective vulnerability management begins with knowing what the organization owns and operates.

From there, businesses must identify weaknesses, understand their business impact, prioritize the most serious risks, assign responsibility, fix or mitigate problems, verify corrective action, and repeat the process.

Cybersecurity requires visibility.

4.38 Chapter Summary

Chapter 3 explained how attackers get inside. This chapter examined how organizations can find many of those same weaknesses before an attacker exploits them.

The process begins with a basic question: What do we have? Then: What is exposed? Then: What is vulnerable? Then: What matters most? And finally: Who is responsible for fixing it?

This progression turns cybersecurity from a vague concern into a manageable process.

Know what you have -> Identify what is exposed -> Find the weaknesses -> Understand the risk -> Prioritize -> Fix the problem -> Verify the fix -> Repeat.

A business does not need to eliminate every vulnerability overnight. It does need to know where its greatest risks are and make steady, deliberate progress toward reducing them.

The organizations that manage cybersecurity effectively are rarely the ones that never discover problems. They are the ones that consistently find problems, understand them, and correct them before those problems become incidents.

And that leads naturally to the next major question: Once we know where our weaknesses are, how do we build security into everyday business operations instead of treating cybersecurity as a separate technical activity?

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

Chapter 5

Building Cybersecurity Into Everyday Business Operations

Authored by Dr. Charles Edwards

Cybersecurity works best when it becomes part of the way the organization operates every day.

It should not appear only in an annual training session, a policy nobody reads, or a conversation that happens after something has already gone wrong.

A business becomes more secure when cybersecurity is part of ordinary decision-making.

When a new employee is hired, security matters. When someone changes jobs, access should be reviewed. When a vendor is given access, security matters. When software is purchased, a new office opens, or an employee leaves, security should be part of the decision.

That is what I mean by building cybersecurity into everyday operations.

The goal is not to turn every employee into a cybersecurity specialist.

The goal is to make secure behavior part of normal business behavior.

5.1 Cybersecurity Is a Business Responsibility

Cybersecurity is often treated as an IT responsibility. That is understandable because many cybersecurity controls involve technology.

But cybersecurity affects much more than technology. It affects finance, human resources, legal, operations, customer service, compliance, vendor management, executive leadership, and every employee who uses company systems.

Consider a fraudulent wire transfer. The technical team may manage email security, but the finance department may control payment approval.

Consider a former employee whose account remains active. The technical team may disable the account, but Human Resources must communicate that the employee has left.

Cybersecurity therefore requires coordination. No single department can manage it effectively alone.

5.2 Leadership Sets the Tone

Employees pay attention to what leadership treats as important.

If management says security matters but consistently ignores security procedures, employees notice.

If executives routinely ask employees to bypass approval procedures because they are in a hurry, employees learn that urgency is more important than security.

If managers share passwords, employees may believe that password sharing is acceptable.

Leadership behavior creates security culture.

A strong security culture begins when management communicates: We expect people to work securely, and we will support them in doing so.

5.3 Policies Should Be Useful, Not Decorative

Most organizations have policies. Some organizations have excellent policies. But a policy provides little value if nobody understands it or follows it.

A cybersecurity policy should answer practical questions about passwords, MFA, personal devices, software installation, sensitive information, suspicious email, cloud storage, incident reporting, vendor access, and employee departure.

A policy should help people make decisions. It should not merely exist to satisfy an audit requirement.

5.4 Avoid Writing Policies Nobody Can Follow

Security policies sometimes fail because they are unrealistic.

A useful policy must be understandable, practical, enforceable, communicated, and supported by real procedures.

Good security policies reflect how the organization actually operates.

5.5 Policies and Procedures Are Different

A policy states what the organization requires. A procedure explains how to do it.

For example, a policy may state that terminated employee access must be disabled promptly. The procedure then defines how Human Resources notifies IT, how accounts are disabled, how company devices are recovered, and how access is verified.

The policy creates the requirement. The procedure creates the action.

5.6 Start With the Most Important Policies

A small business does not need hundreds of pages of cybersecurity documentation.

It does need clear guidance in important areas.

- Acceptable Use Policy
- Password and Authentication Policy
- Access Control Policy
- Remote Work Policy
- Mobile Device Policy
- Data Handling Policy
- Backup Policy
- Incident Reporting Policy
- Vendor Access Policy
- Software Installation Policy
- Employee Offboarding Policy

5.7 Hiring Is a Cybersecurity Process

Cybersecurity begins before the employee's first day.

When a new employee is hired, the organization should determine what systems the employee needs, what information the employee should access, whether remote access is required, whether administrative privileges are necessary, what devices will be issued, and what security training is needed.

This process is often called onboarding. A good onboarding process creates the correct access from the beginning.

5.8 Do Not Give Everyone the Same Access

Employees have different responsibilities, and their access should reflect those responsibilities.

This is another application of the principle of least privilege: give people what they need and remove what they do not.

5.9 Role-Based Access

One useful approach is role-based access control. Instead of creating permissions individually for every employee, organizations define access based on job roles.

When someone joins a department, that person receives the appropriate role. When the person changes departments, the old role is removed and the new one is assigned.

This can reduce privilege creep and make access reviews easier.

5.10 Access Reviews Should Be Routine

Access should not be reviewed only when something goes wrong.

Organizations should periodically ask who has access to each system, who has administrator rights, whether all accounts are still needed, whether former employees or vendors are still listed, and whether anyone has more access than necessary.

5.11 Employee Transfers Create Hidden Risk

An employee who changes jobs inside the company may accumulate access.

Internal transfers should therefore include access review - not just adding access, but removing access that is no longer required.

5.12 Offboarding Must Be Immediate and Complete

When an employee leaves, access should be removed promptly.

The organization should consider email, VPN, cloud services, company applications, shared accounts, administrator accounts, mobile devices, physical access cards, remote-management tools, and vendor systems.

A delayed offboarding process creates unnecessary risk.

5.13 Shared Accounts Create Accountability Problems

Shared accounts may sometimes be operationally necessary, but they create accountability problems.

If several people use the same login, it becomes difficult to determine who performed an action, changed a setting, accessed a file, or downloaded information.

Individual accounts provide better accountability.

5.14 Administrator Accounts Need Special Treatment

Administrator accounts can make major changes and should receive stronger protection.

Good practices include using separate administrator accounts, requiring MFA, limiting the number of administrators, avoiding routine web browsing from administrator accounts, monitoring administrator activity, and reviewing privileges regularly.

5.15 Remote Work Changes the Security Environment

Remote work is normal for many organizations, but the organization no longer controls the physical environment.

Remote work policies should address device security, Wi-Fi use, VPN requirements, MFA, physical privacy, document handling, screen locking, software updates, and lost or stolen devices.

5.16 Home Networks Are Not Corporate Networks

Employees often connect company devices to home Wi-Fi, but home networks vary greatly in security.

Organizations can reduce dependency on those networks by using secure remote access, endpoint protection, local firewalls, patching, MFA, disk encryption, and a model that treats the local network as untrusted.

5.17 Public Wi-Fi Requires Caution

Public Wi-Fi is convenient, but it is outside the organization's control.

Useful safeguards include secure VPN connections, mobile hotspots where appropriate, disabling unnecessary sharing, and verifying the correct wireless network.

Unfamiliar networks should be treated as untrusted.

5.18 Company Devices Versus Personal Devices

Organizations must decide whether employees can use personal devices for work. This is often called Bring Your Own Device, or BYOD.

BYOD can provide flexibility, but it also raises questions about encryption, patching, screen locks, data removal, device sharing, and lost devices.

There is no single answer for every organization, but there should be a deliberate decision.

5.19 Mobile Device Management

Organizations with many mobile devices may use Mobile Device Management, or MDM.

MDM can help enforce requirements such as screen locks, encryption, device updates, approved applications, remote wipe, and security settings.

5.20 Software Installation Should Be Controlled

Every new application creates potential risk.

Organizations should know what software is installed, and employees should not automatically be able to install any software they choose.

The process should match the organization, but software installation should not be invisible.

5.21 Free Software Is Not Always Free

Employees may download free tools because they are convenient.

The organization should still consider what information the software accesses, where data is stored, whether the vendor is trustworthy, whether the tool receives security updates, and whether licensing permits business use.

The security impact matters too.

5.22 Browser Extensions Deserve Attention

Browser extensions can access surprising amounts of information.

Depending on their permissions, extensions may read web pages, modify website content, access browsing activity, read form entries, or interact with cloud applications.

Organizations should consider limiting browser extensions to approved tools.

5.23 Data Classification Helps People Make Better Decisions

Not all information needs the same level of protection.

One way to manage this is through data classification. Simple categories might include Public, Internal, Confidential, and Restricted.

The value comes from helping employees understand how information should be handled.

5.24 Employees Need to Know Where Sensitive Data Is Allowed

A classification policy is useful only if employees know what to do.

Confidential information may be allowed in approved cloud storage, encrypted company devices, and authorized business applications, while being prohibited in personal email, personal cloud drives, unapproved USB devices, or public AI tools.

Clear rules reduce accidental exposure.

5.25 Email Is Not Automatically the Right Place for Everything

Email is extremely useful, but sensitive information may remain in mailboxes for years, be forwarded, or be downloaded to personal devices.

Organizations should think carefully about what kinds of sensitive information should be sent by ordinary email.

5.26 File Sharing Needs Control

Modern businesses frequently share documents through cloud platforms.

Common problems include links that allow anyone to view a document, files shared with personal accounts, former employees retaining access, documents shared indefinitely, and entire folders shared when only one file was needed.

Sharing permissions should be reviewed periodically.

5.27 Financial Processes Need Security Controls

Finance departments are common targets because attackers want money.

Useful controls may include dual approval, independent verification, transaction limits, known callback procedures, separation of duties, and unusual-payment review.

Vendor banking changes should be verified through a previously known contact method.

5.28 Separation of Duties

Separation of duties means one person should not control every step of a sensitive transaction.

One employee may enter a payment while another approves it. One employee may create a new vendor while another verifies the bank information.

This reduces fraud and mistakes and creates protection when one account is compromised.

5.29 Vendor Management Is Part of Cybersecurity

Vendors often need access to company systems or information.

Organizations should understand what the vendor can access, why the access is needed, how the vendor authenticates, whether MFA is required, how incidents will be reported, and when access will be removed.

5.30 Vendor Access Should Expire

Temporary vendor access frequently becomes permanent.

Where practical, temporary access should have an expiration date. If access is still needed, it can be renewed.

5.31 Cybersecurity Training Should Be Continuous

Annual training is better than no training, but people forget and threats change.

Short, recurring security reminders are often more useful than one large annual presentation.

5.32 Training Should Reflect Real Work

Generic security training has limitations.

Employees learn better when examples resemble situations they actually face. Accounting may need payment-fraud examples, HR may need payroll and employee-data examples, administrators may need privileged-account examples, and executives may need business-email-compromise examples.

5.33 Employees Need a Safe Way to Report Mistakes

If employees fear punishment or embarrassment, they may stay silent after clicking a suspicious link or making another mistake.

Organizations should encourage immediate reporting because delay can help an attacker.

Fast reporting can prevent a small event from becoming a large one.

5.34 Make Reporting Easy

Employees should not have to search through a policy manual to report a suspicious message.

Provide a clear method such as a dedicated email address, help desk number, reporting button, or security contact.

5.35 Security Should Be Part of Purchasing

Organizations routinely purchase software, cloud services, equipment, consulting services, and technology platforms.

Security should be considered before purchase. Questions may include whether the product supports MFA, audit logs, role-based access, secure data handling, reliable updates, breach notification, data export, and centralized account disabling.

5.36 Security by Design

Security by design means considering security during planning rather than adding it after implementation.

Authentication, permissions, encryption, logging, backups, patching, incident response, and data protection should be considered before deployment.

5.37 Change Management Matters

Organizations make changes constantly, and those changes can accidentally weaken security.

A basic change-management process helps track what is changing, who approved it, why it is necessary, when it will occur, and how it can be reversed if something goes wrong.

5.38 Emergency Changes Need Review Too

Sometimes changes must happen quickly, but they should still be reviewed afterward.

Otherwise emergency configurations can become permanent and contribute to security drift.

5.39 Logging Is Part of Normal Operations

Security logs may show logins, failed authentication attempts, administrator changes, file access, firewall events, cloud activity, and application behavior.

Logs are valuable only if they exist, are retained, and can be reviewed when needed.

5.40 Alerts Need Owners

Security tools may generate alerts, but an alert without an owner is just a message.

Every important alerting system should identify who receives alerts, who reviews them, how quickly they are handled, and what happens after business hours.

5.41 Backups Are an Operational Process

Backups should not be thought of as a one-time setup.

Organizations should regularly verify that backups are completing, important systems are included, storage is available, copies are protected, retention is appropriate, and restoration works.

5.42 Incident Response Should Be Practiced

An incident response plan should not be created and forgotten.

Organizations should periodically discuss scenarios such as Microsoft 365 compromise, ransomware, executive account takeover, legal notification, customer communication, insurance contact, and system shutdown authority.

These exercises are often called tabletop exercises.

5.43 Cyber Insurance Does Not Replace Security

Cyber insurance can help organizations manage financial consequences, but it does not prevent a breach.

Organizations should understand their policy requirements and recognize that security controls such as MFA, backups, monitoring, endpoint protection, and training remain necessary.

5.44 Measure What Matters

Organizations need ways to determine whether cybersecurity is improving.

Useful measurements may include patching rates, MFA coverage, dormant accounts removed, open critical vulnerabilities, backup restoration success, phishing reporting rates, offboarding speed, number of privileged accounts, and response times.

5.45 Avoid Metrics That Encourage the Wrong Behavior

Metrics can be misleading.

Zero reported phishing incidents may mean employees are not reporting suspicious messages. Closing 1,000 vulnerabilities may sound impressive while the five most dangerous ones remain unresolved.

Good metrics require context.

5.46 A Business Scenario: The Payroll Change

Consider a fictional company called Riverside Professional Services with approximately 60 employees.

One Friday afternoon, payroll specialist Daniel receives an email that appears to come from a senior manager asking for a last-minute direct-deposit change.

The email looks normal and payroll closes soon, but Riverside has a policy requiring payroll banking changes to be verified through a second communication method.

Daniel calls the manager using the phone number already stored in the company directory. The manager says the email was fraudulent.

The attacker fails because of a business process, not because of antivirus software or a firewall.

5.47 Another Scenario: The Departing Contractor

A contractor works with an organization for six months and receives VPN access, cloud storage, a project-management account, and access to a shared development server.

The project ends, but nobody removes the accounts.

Nine months later, the contractor's personal email account is compromised and some old credentials still work.

This could have been prevented with a simple offboarding checklist.

5.48 Building a Security Culture

Security culture is the collection of behaviors, expectations, and attitudes surrounding cybersecurity.

A strong security culture means employees naturally think about verification, reporting, password protection, software approval, data protection, and unusual activity.

Culture develops through leadership, training, communication, procedures, reinforcement, and example.

5.49 Security Culture Should Not Be Based on Fear

Fear can produce short-term compliance, but it can also produce silence.

If employees believe every mistake will result in punishment, they may hide problems.

A more effective message is that security is part of everyone's job, questions are welcome, concerns should be reported, and mistakes should be disclosed quickly.

5.50 Practical Steps for Small and Mid-Sized Businesses

A business does not need a massive security department to improve operations.

Consistency matters more than complexity.

1. Require MFA for important accounts.
2. Maintain an employee onboarding checklist.
3. Maintain an employee offboarding checklist.
4. Review administrator access.
5. Review vendor access.
6. Establish a suspicious-message reporting process.
7. Require secondary verification for financial changes.
8. Document backup responsibilities.
9. Control software installation.
10. Train employees using realistic examples.
11. Maintain basic cybersecurity policies.
12. Review important access periodically.
13. Practice an incident-response scenario.
14. Track important security tasks.
15. Assign ownership.

5.51 Questions Every Business Owner Should Ask

Leadership should ask practical questions about employees, administrators, remote work, financial processes, data, vendors, and security culture.

About Employees

- Do new employees receive only the access they need?
- Is cybersecurity training part of onboarding?
- Are employee transfers reviewed?
- Are former employee accounts removed promptly?

About Administrators

- How many administrator accounts exist?
- Are they individually assigned?
- Is MFA required?
- Are privileges reviewed?

About Remote Work

- Are company devices encrypted?
- Are they patched?
- Is MFA required?
- Do employees know how to work safely on public networks?

About Financial Processes

- Are payment changes independently verified?
- Are large transactions subject to secondary approval?
- Are vendor banking changes confirmed separately?

About Data

- Do employees know what information is sensitive?
- Do they know where sensitive information may be stored?
- Are sharing permissions reviewed?

About Vendors

- Who has outside access?
- Why?
- When does that access expire?
- Is MFA required?

About Security Culture

- Do employees know how to report suspicious activity?
- Are reports encouraged?
- Are mistakes reported quickly?
- Does leadership follow the same rules?

5.52 Questions to Ask Your IT Provider

Businesses should understand what their provider handles and what remains the organization's responsibility.

1. How are new employee accounts created?
2. How are former employee accounts disabled?
3. Do you receive timely termination notifications?
4. Are administrator accounts separate from normal accounts?
5. Is MFA required for privileged users?
6. Do you periodically review user access?
7. Do you help us identify dormant accounts?
8. How do you control software installation?
9. Are company laptops encrypted?
10. How are remote devices protected?
11. Are mobile devices managed?
12. How do you monitor security alerts?
13. Who responds after business hours?
14. How are backups tested?
15. Do you help us conduct incident-response exercises?
16. What security policies should we maintain?
17. Do you review vendor access?
18. How are security changes documented?
19. How do you track unresolved security issues?
20. What security responsibilities remain with us?

5.53 Chapter 5 Cybersecurity Checklist

Use the following checklist to review leadership, policies, employee life cycle, access control, remote work, data protection, financial security, vendors, security operations, and security awareness.

Leadership

☐ Leadership communicates that cybersecurity is important.
☐ Executives follow security procedures.
☐ Cybersecurity responsibilities are assigned.

Policies

☐ Password and authentication requirements are documented.
☐ Remote work rules are documented.
☐ Data handling expectations are documented.
☐ Incident reporting procedures are documented.
☐ Vendor access rules are documented.

Employee Life Cycle

☐ New employees receive appropriate access.
☐ Cybersecurity training is included in onboarding.
☐ Employee transfers trigger access review.
☐ Access for departing employees is disabled promptly.
☐ Company equipment is returned.

Access Control

☐ Least privilege is used.
☐ Administrator accounts are limited.
☐ Privileged access is periodically reviewed.
☐ Dormant accounts are removed.
☐ Shared accounts are minimized.

Remote Work

☐ Company devices are patched.
☐ Disk encryption is enabled.
☐ Screen locking is required.
☐ MFA is used.
☐ Remote access is protected.

Data Protection

☐ Sensitive data categories are understood.
☐ Employees know approved storage locations.
☐ Personal email is not used for company data.
☐ Cloud sharing permissions are reviewed.

Financial Security

☐ Bank changes require independent verification.
☐ Payroll changes are verified.

☐ Significant transactions require appropriate approval.
☐ Employees know how to identify executive impersonation.

Vendors

☐ Vendor access is documented.
☐ Vendor accounts are individually assigned where possible.
☐ MFA is required where appropriate.
☐ Temporary access expires.
☐ Vendor accounts are removed when no longer needed.

Security Operations

☐ Security alerts have an owner.
☐ Backups are monitored.
☐ Restores are tested.
☐ Important changes are documented.
☐ Incident-response plans are practiced.

Security Awareness

☐ Employees receive recurring security reminders.
☐ Training reflects job responsibilities.
☐ Reporting suspicious activity is easy.
☐ Employees are encouraged to report mistakes quickly.

5.54 Key Terms

Acceptable Use Policy: A policy explaining how employees are permitted to use company technology and systems.

Bring Your Own Device (BYOD): A practice that allows employees to use personally owned devices for business purposes.

Change Management: A structured process for planning, approving, documenting, and reviewing changes to technology or systems.

Data Classification: The process of categorizing information based on sensitivity and required protection.

Least Privilege: The principle of giving users only the access necessary to perform their responsibilities.

Mobile Device Management (MDM): Technology used to manage and enforce security requirements on mobile devices.

Offboarding: The process of removing access and recovering company assets when an employee, contractor, or vendor leaves.

Onboarding: The process of creating appropriate access and preparing a new employee to work securely.

Role-Based Access Control: A method of assigning system permissions according to job roles.

Security by Design: The practice of considering cybersecurity during planning and development rather than adding it afterward.

Security Culture: The shared attitudes, expectations, and behaviors surrounding cybersecurity within an organization.

Separation of Duties: The practice of dividing sensitive responsibilities among multiple people so that no one individual controls the entire process.

5.55 Key Takeaways

Cybersecurity becomes stronger when it is built into ordinary business operations.

It should be part of hiring, employee transfers, employee departures, remote work, purchasing, financial approvals, vendor relationships, data handling, system changes, backups, and incident reporting.

Policies matter, but policies must lead to practical procedures. Technology matters, but technology cannot replace sound business processes. Employee training matters, but training should reflect real job responsibilities. Leadership matters because employees pay attention to what leaders actually do.

The goal is to make secure behavior normal.

5.56 Chapter Summary

Chapter 4 focused on finding weaknesses before attackers do.

Chapter 5 has focused on preventing many of those weaknesses from developing in the first place.

An organization becomes more secure when cybersecurity is included in everyday operations.

The process looks like this: Hire securely -> Grant appropriate access -> Protect devices and data -> Verify sensitive transactions -> Manage vendors -> Control changes -> Monitor security activity -> Train employees -> Remove access when no longer needed -> Review and improve.

Cybersecurity is not a separate activity performed by one technical department. It is part of how the business operates.

Strong organizations do not rely on employees remembering security only when something looks suspicious. They build procedures that make secure actions easier and unsafe actions harder.

And that leads naturally to the next major question: When an incident still occurs despite all of these safeguards, how should the organization respond during the first minutes, hours, and days?

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

Chapter 6

What to Do When Something Goes Wrong: Incident Response and Recovery

Authored by Dr. Charles Edwards

No organization can guarantee that a cybersecurity incident will never happen.

That is one of the most important realities in cybersecurity.

Strong security controls reduce risk. Employee training reduces risk. Multifactor authentication reduces risk. Good backups reduce risk. Monitoring reduces risk. Policies and procedures reduce risk.

But none of these eliminates risk completely.

A determined attacker may still find a way in. An employee may make a mistake. A vendor may be compromised. A new vulnerability may appear before a patch exists.

The question is not only, 'How do we prevent incidents?'

We also have to ask, 'What will we do when prevention fails?'

That is the purpose of incident response.

Incident response is the organized process used to identify, contain, investigate, eliminate, recover from, and learn from cybersecurity incidents.

A strong response process helps the organization move from confusion to coordinated action.

The first minutes and hours matter because what the organization does during that time can affect the size of the loss and the speed of recovery.

The goal is not panic. The goal is preparation.

6.1 What Is a Cybersecurity Incident?

A cybersecurity incident is an event that threatens or compromises the confidentiality, integrity, or availability of information or systems.

Examples may include ransomware, stolen credentials, unauthorized account access, malware infections, data theft, lost or stolen devices, fraudulent financial activity, website compromise, denial-of-service attacks, cloud account takeover, unauthorized access by an employee, compromised vendor access, or accidental disclosure of sensitive information.

Not every unusual event is a major incident. A single failed login may simply be a user typing the wrong password. Ten thousand failed login attempts from several countries may be something very different.

The purpose of incident response is to determine what is happening and respond appropriately.

6.2 Events, Alerts, and Incidents Are Different

Organizations often receive security information in different forms.

An event is something that happened. For example: A user logged in.

An alert is a notification that an event may require attention. For example: A user logged in from an unusual country.

An incident is a confirmed or suspected security problem requiring coordinated response. For example: A criminal used stolen credentials to access the employee's account and download confidential files.

This distinction matters. If every alert is treated as a full emergency, employees become overwhelmed. If serious alerts are ignored, incidents may go undetected.

Organizations need a process for determining which events require action.

6.3 The First Goal Is to Understand What Is Happening

When an incident is first reported, information may be incomplete.

An employee may say, “My computer is acting strange.” Another may say, “I clicked something and now I cannot log in.” A manager may report, “Customers are receiving emails from my account that I did not send.”

The security or IT team needs to quickly determine what happened, when it started, which systems are affected, which accounts may be involved, whether the problem is continuing, whether sensitive information may be at risk, and whether business operations are affected.

The organization should avoid making assumptions too early. Initial reports may be inaccurate or incomplete.

Incident response should begin with facts.

6.4 Why Speed Matters

Attackers do not stop because the organization discovered them.

If a criminal still has access, the attacker may continue to steal information, create new accounts, install malware, move to other systems, disable security tools, delete backups, or encrypt systems.

That means response time matters. An incident contained in 20 minutes may cause far less damage than one allowed to continue for three days.

This does not mean acting recklessly. Fast action still requires coordination.

The goal is: Move quickly, but move deliberately.

6.5 Incident Response Begins Before the Incident

The worst time to decide what to do during a cyberattack is after the attack has already begun.

Preparation should happen in advance.

Organizations should know who will lead the response, how to contact key people, who can authorize major decisions, how outside experts will be reached, what insurance requirements apply, who will communicate with customers, who will contact law enforcement if necessary, and how evidence will be preserved.

This preparation becomes the foundation of the incident-response plan.

6.6 The Incident-Response Team

A cybersecurity incident may involve more than the IT department.

Depending on the severity, the response team may include executive leadership, IT, cybersecurity staff, legal counsel, human resources, finance, communications, compliance, insurance representatives, outside forensic specialists, managed service providers, and law enforcement.

Each participant has a different responsibility.

The technical team may contain the attack. Legal counsel may advise on notification requirements. Communications may prepare customer messaging. Executives may authorize major operational decisions. Finance may handle emergency purchases. Human Resources may assist if employee behavior is involved.

A serious incident becomes a business event, not just a technical event.

6.7 Define Roles Before the Emergency

Organizations should decide in advance who is responsible for what.

Incident Coordinator

Organizes the overall response.

Technical Lead

Directs investigation, containment, and recovery.

Executive Decision-Maker

Approves major business decisions.

Legal Contact

Provides legal guidance and notification advice.

Communications Lead

Coordinates internal and external messaging.

Insurance Contact

Coordinates with the cyber insurance carrier.

Documentation Lead

Maintains a timeline and records decisions.

Clear roles reduce confusion. Without them, several people may attempt to control the incident at the same time. Or worse, everyone may assume someone else is responsible.

6.8 Build a Contact List Before You Need It

During a ransomware incident, the company's email system may be unavailable. During a cloud compromise, the attacker may be reading company email.

This creates a practical problem: How do you contact the response team?

Organizations should maintain an emergency contact list outside the normal production environment.

The contact list should be accessible even when normal systems are unavailable.

- executive phone numbers
- IT leadership
- legal counsel
- insurance carrier
- managed service provider
- incident-response vendor
- key cloud providers

- law enforcement contacts
- critical vendors

6.9 Alternative Communication Methods Matter

Suppose the company believes Microsoft 365 has been compromised. Should everyone continue discussing the incident through company email? Possibly not.

If attackers have access to the email environment, they may be able to read those conversations.

Organizations should plan alternative communication methods, such as telephone, secure messaging platforms, emergency conference bridges, or prearranged external communication tools.

The method should be chosen before an incident.

6.10 The Incident-Response Life Cycle

A practical incident-response process can be divided into several major phases: Preparation; Detection and Analysis; Containment; Eradication; Recovery; Lessons Learned.

These phases may overlap. Real incidents do not always proceed in a perfect sequence.

But the model helps organizations understand what needs to happen.

6.11 Phase One: Preparation

Preparation includes everything done before an incident occurs.

Examples include creating the incident-response plan, assigning roles, maintaining contact lists, establishing logging, protecting backups, preparing emergency communications, identifying outside experts, reviewing cyber insurance, documenting critical systems, and practicing tabletop exercises.

Preparation determines how effectively the organization responds later.

A company cannot create missing logs after the breach. It cannot retroactively test a backup that was never configured. It cannot easily determine system ownership if no inventory exists.

Preparation creates options.

6.12 Phase Two: Detection and Analysis

Detection begins when something unusual is noticed.

The signal may come from security software, an employee, a customer, a vendor, law enforcement, a bank, an insurance carrier, or another organization.

Examples include impossible-travel login alerts, malware detections, large data transfers, suspicious account activity, unusual administrator changes, ransom notes, unexpected password resets, or fraudulent payment requests.

The organization must then determine whether a real incident has occurred. This process is called analysis.

6.13 Employees Are Part of Detection

Security tools are important, but employees often notice incidents first.

A user may report, “My password suddenly stopped working.” Another may report, “I received MFA requests I did not initiate.” Another may say, “A customer told me they received a strange email from my account.”

Employees should understand that these observations matter. They should know how to report them quickly.

Good incident detection depends on both technology and people.

6.14 Containment: Stop the Damage From Spreading

Once an incident is confirmed or strongly suspected, the organization may need to contain it.

Containment means limiting the attacker's ability to continue causing damage.

Possible actions may include disabling a compromised account, isolating an infected computer, blocking malicious network traffic, disconnecting a server, disabling remote access, forcing password resets, revoking access tokens, removing external sharing, or temporarily shutting down a vulnerable service.

Containment decisions can be difficult. Disconnecting a critical server may stop the attacker. It may also stop business operations.

Leadership may need to balance security and continuity.

6.15 Do Not Destroy Evidence During Containment

Organizations may be tempted to immediately erase or rebuild compromised systems.

Sometimes that becomes necessary. But doing so too quickly can destroy valuable evidence.

Evidence may help determine how the attacker entered, what systems were affected, what accounts were used, what information was accessed, how long the attacker was present, and whether legal notification is required.

The technical team should consider evidence preservation before rebuilding systems. This may require outside forensic expertise.

6.16 Disconnecting a Computer Is Not the Same as Turning It Off

This distinction may matter during an investigation.

Disconnecting a computer from the network can stop communication with the attacker while preserving information in memory.

Turning the system off may destroy certain volatile evidence.

The correct action depends on the situation. Employees should generally avoid experimenting with compromised devices themselves. They should contact the designated IT or security team and follow instructions.

6.17 Eradication: Remove the Attacker

Containment stops or limits the immediate damage. Eradication removes the cause.

This may include removing malware, deleting unauthorized accounts, patching vulnerabilities, changing compromised passwords, revoking tokens, removing malicious email rules, correcting firewall settings, rebuilding systems, removing persistence mechanisms, and closing unauthorized remote access.

The organization must be careful not to stop after fixing only the obvious problem.

If one compromised account is reset but the attacker created three additional accounts, the incident is not over.

The question becomes: Have we removed the attacker completely?

6.18 Recovery: Return to Normal Operations

Recovery is the process of restoring systems and services safely.

This may involve rebuilding servers, restoring backups, reissuing credentials, reconnecting systems, validating configurations, testing applications, monitoring for recurrence, and gradually returning systems to production.

Recovery should be deliberate.

An organization should avoid reconnecting a system before confirming that the weakness has been corrected. Otherwise the attacker may simply return.

6.19 Recovery Should Follow Business Priorities

Not every system has the same importance.

Organizations should identify critical functions in advance. A possible recovery order might include identity and authentication systems, network infrastructure, financial systems, customer-facing systems, file storage, internal applications, and less critical services.

This recovery order will vary.

The key principle is: Recover based on business impact, not convenience.

6.20 Backups Are Essential to Recovery

Strong backups can dramatically improve incident recovery.

But backups must be available, intact, protected, and tested.

Organizations should know which systems can be restored, how long restoration will take, how much data may be lost, whether backup copies are isolated, and whether the backups themselves were compromised.

The existence of a backup file does not guarantee successful recovery. Restoration must be tested.

6.21 Recovery Time Objective and Recovery Point Objective

Two useful concepts in disaster recovery are Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

Recovery Time Objective asks: How quickly must a system be restored? For example, payroll must be restored within eight hours.

Recovery Point Objective asks: How much data loss can the organization tolerate? For example, we can tolerate losing no more than four hours of transaction data.

These concepts help determine backup frequency and recovery priorities. They also help leadership understand the relationship between cybersecurity and business continuity.

6.22 Ransomware Changes the Recovery Problem

Ransomware incidents can be especially complex.

The organization may face several problems at once: encrypted systems, stolen information, disabled accounts, unavailable backups, customer disruption, regulatory obligations, and ransom demands.

Leadership may be pressured to make rapid decisions. This is where preparation becomes critical.

Organizations should know in advance who has authority to make major decisions, what insurance requires, whether legal counsel must be involved, who communicates with law enforcement, and how recovery will proceed.

6.23 Should a Business Pay a Ransom?

This is one of the most difficult questions during a ransomware incident.

There is no simple answer.

Paying a ransom does not guarantee that decryption will work, that stolen data will be deleted, that attackers will not return, or that information will not be sold later.

There may also be legal, regulatory, insurance, and sanctions considerations.

The decision should not be made casually. Leadership should involve legal counsel, insurance, incident-response specialists, and law enforcement where appropriate.

The strongest long-term strategy is to build recovery capabilities that reduce dependence on the attacker.

6.24 Communication During an Incident

Communication can either help or harm an incident response.

Employees need accurate information. Customers may need updates. Vendors may need instructions. Executives need decision-making information.

Poor communication can create rumors, panic, conflicting instructions, accidental disclosure, and loss of trust.

Organizations should designate who communicates externally. Not every employee should speak publicly about an active incident.

6.25 Tell Employees What They Need to Know

During an incident, employees may need practical instructions.

Examples include: Do not connect to VPN. Do not power off your workstation. Change your password. Do not use company email. Use this temporary communication channel. Report unusual activity immediately. Do not discuss the incident publicly.

Instructions should be clear and specific.

Employees do not need every technical detail. They need to know what action to take.

6.26 Customer Communication Requires Care

If customers may be affected, communication may become necessary.

Messages should be accurate, timely, consistent, and coordinated with legal counsel.

Avoid guessing.

The employee says, “I clicked a link, but I closed it right away.”

Another employee says, “I entered my password, but then the page looked strange.”

6.27 Legal and Regulatory Notification

Depending on the type of data involved, organizations may have legal obligations to notify individuals, regulators, government agencies, business partners, or other parties.

Notification requirements vary by jurisdiction, industry, contract, and type of information.

This is why legal counsel should be involved when sensitive information may have been exposed.

Cybersecurity incidents can quickly become legal events.

6.28 Preserve a Timeline

During an incident, someone should maintain a timeline.

The timeline may include when the incident was discovered, who reported it, what systems were affected, what actions were taken, who authorized decisions, when

accounts were disabled, when systems were isolated, when outside experts were contacted, when backups were restored, and when services returned.

This record becomes valuable for investigation, insurance, legal review, lessons learned, and future improvement.

Without documentation, people rely on memory. During a stressful incident, memory is unreliable.

6.29 Keep Decision Records

Important decisions should also be documented.

Examples include shutting down a system, delaying a public statement, restoring from a particular backup, involving law enforcement, contacting customers, or approving emergency spending.

Decision records help explain what happened later. They also support accountability.

6.30 Cyber Insurance During an Incident

Organizations with cyber insurance should understand the policy before an incident.

Some policies require the insured organization to contact the carrier quickly, use approved vendors, obtain authorization before hiring outside firms, or follow specific reporting procedures.

Ignoring these requirements may affect coverage.

The insurance contact information should therefore be part of the incident-response plan.

6.31 Outside Incident-Response Specialists

Some incidents exceed the capabilities of internal IT staff.

Outside specialists may provide digital forensics, malware analysis, ransomware response, data-breach investigation, threat hunting, recovery support, and legal coordination.

Organizations should identify potential specialists before an emergency.

Trying to find qualified help during a major breach wastes valuable time.

6.32 Law Enforcement

Organizations may choose or be advised to contact law enforcement during serious cyber incidents.

Law enforcement may provide investigative assistance, threat intelligence, coordination with other victims, information about known criminal groups, or guidance related to extortion.

The decision should be coordinated with legal counsel and leadership.

6.33 Do Not Assume the Attacker Has Left

After systems appear normal, the organization should continue monitoring.

Attackers may have additional accounts, stolen credentials, persistence mechanisms, access tokens, remote-management tools, or compromised vendor accounts.

Recovery should include heightened monitoring.

The organization should look for signs that the attacker is trying to return.

6.34 Password Resets Must Be Planned

During an account compromise, organizations often reset passwords.

But password resets should be coordinated.

If all users are forced to change passwords while malware remains active on devices, the attacker may simply steal the new passwords.

The organization may need to clean or isolate affected systems, revoke active sessions, disable unauthorized accounts, reset credentials, enforce MFA, and monitor for suspicious activity.

The order matters.

6.35 Revoke Sessions and Tokens

Modern cloud platforms may allow users to remain logged in after a password change.

Attackers may possess browser sessions, refresh tokens, application tokens, or remembered-device access.

For this reason, incident response may require more than changing a password.

Active sessions may need to be revoked. Trusted devices may need to be reviewed. Applications may need to be reauthorized.

Identity recovery must be comprehensive.

6.36 Business Continuity During an Incident

A cyberattack can disrupt normal operations.

Businesses should plan how critical work will continue.

Questions include whether payroll can be processed manually, whether customers can still reach the organization, whether employees can access emergency procedures, whether orders and invoices can be processed, whether employees can communicate, and whether critical systems can operate offline.

This is where cybersecurity overlaps with business continuity.

The organization must protect both information and operations.

6.37 Manual Workarounds Can Be Valuable

Technology failures do not always need to stop every business process.

Some organizations maintain manual alternatives, such as printed emergency contact lists, paper forms, offline customer lists, alternate payment procedures, manual scheduling, and emergency communication trees.

These methods may seem old-fashioned. During a major outage, they may keep the business functioning.

6.38 A Business Scenario: The Compromised Email Account

Consider a fictional company called Evergreen Consulting Group.

Evergreen has 75 employees.

On Tuesday morning, employee Jasmine receives several MFA prompts she did not initiate. She ignores them.

The employee says, “I did not think it was important.”

She immediately reports the problem.

The IT team discovers that Jasmine's account was accessed from an unfamiliar location. The attacker created an email rule that hid messages containing the words payment, fraud, bank, and suspicious.

The account is disabled. Active sessions are revoked. The password is reset. The malicious email rule is removed.

The organization checks whether other accounts were accessed. Finance contacts customers who may have received fraudulent payment instructions. Security logs are preserved. Legal counsel is notified.

Because Jasmine reported the incident quickly, the attacker had limited time to act. No customer made a fraudulent payment.

Speed mattered.

6.39 What Evergreen Did Well

Evergreen had several advantages.

The employee knew how to report the issue. The IT team had access to login logs. The organization knew how to disable the account. The finance team had a customer contact process. Legal counsel was already identified. The response was coordinated.

The organization was not perfect. The account was still compromised.

But the damage was limited because the response process worked.

6.40 A Ransomware Scenario

Now consider another fictional company: Stonebridge Manufacturing.

Stonebridge has 140 employees.

At 6:45 AM on Monday, employees report that shared files are unavailable. Several workstations display ransom notes.

The IT team discovers that multiple servers have been encrypted. The attacker also disabled several backup jobs.

Management immediately activates the incident-response plan.

Because the company had previously tested backups, it knows that clean copies exist.

Recovery will still take time. But the company has options.

That is the value of preparation.

The company:

1. isolates affected network segments
2. disables remote access
3. contacts its incident-response provider

The organization notifies its cyber insurance carrier, legal counsel, and appropriate external responders according to its plan and obligations.

5. engages legal counsel

6. begins preserving evidence

7. moves communications to an alternate platform

8. reviews backup integrity

9. identifies which systems are business-critical

10. begins recovery planning

6.41 What If Stonebridge Had No Plan?

Imagine the same incident without preparation.

The organization contacts a random ransomware-response company found online before checking its incident-response plan, insurance requirements, or established contacts.

The incident becomes more damaging because the response itself is disorganized.

Cybersecurity preparation cannot guarantee a painless incident. But it can prevent unnecessary chaos.

6.42 Tabletop Exercises

A tabletop exercise is a discussion-based practice session.

The organization is given a scenario, such as: It is Monday at 8:15 AM. Employees cannot access the file server. Several computers display ransom notes. What do we do?

Participants then walk through the response.

Questions may include who takes control, who is contacted first, whether systems should be shut down, how the organization communicates, who contacts insurance, whether legal counsel is available, where backups are, who speaks to customers, and what happens if email is unavailable.

The goal is not to test technical skill. The goal is to identify confusion.

6.43 A Tabletop Exercise Should Reveal Weaknesses

A successful tabletop exercise often uncovers problems.

For example, emergency phone numbers may be outdated, nobody may know the insurance policy number, the backup contact may be on vacation, legal counsel may

not have been identified, the alternate communication system may never have been tested, or nobody may know who can authorize shutting down production.

These are valuable discoveries.

It is better to find them during a practice exercise than during a real attack.

6.44 Lessons Learned

After the incident is controlled, the organization should conduct a lessons-learned review.

The purpose is not to assign blame. The purpose is to improve.

Questions should include: What happened? How did the attacker get in? When did the incident begin? When did we detect it? What worked well? What caused delays? Which controls failed? Which controls succeeded? What information was missing? What should change?

The answers should lead to specific improvements.

6.45 The Incident Is Not Over When the Computers Work Again

Technical recovery may happen before organizational recovery.

After systems return, the business may still need to complete legal notifications, respond to customers, file insurance documentation, improve security controls, reset additional credentials, review vendor access, train employees, update policies, and monitor for recurrence.

Cyber incidents have long tails.

Recovery continues after technology returns to normal.

6.46 Avoid the Blame Cycle

After an incident, people often ask: Who caused this?

That question may sometimes be necessary. But the more useful question is: Why was one mistake able to cause this much damage?

Perhaps an employee clicked a phishing link. But why did that lead to administrative access? Why was MFA ineffective? Why were backups reachable? Why was lateral movement possible? Why did detection take four days?

These questions improve systems.

Blaming one individual may hide deeper weaknesses.

6.47 Cybersecurity Incidents Are Learning Opportunities

No organization wants a cyber incident.

But an incident can reveal unknown vulnerabilities, weak procedures, inadequate monitoring, missing documentation, unclear responsibilities, and training gaps.

Organizations that learn from incidents become stronger.

Organizations that simply restore systems and return to normal may experience the same problem again.

6.48 Questions Every Business Owner Should Ask

Leadership should periodically ask practical questions about preparedness, communication, legal and insurance obligations, technology, recovery, and practice.

About Preparedness

- Do we have an incident-response plan?
- When was it last updated?
- Who is in charge during an incident?
- Do we have emergency contact information outside company systems?

About Communication

- What happens if email is unavailable?
- Who communicates with employees?
- Who communicates with customers?
- Who is authorized to speak publicly?

About Legal and Insurance

- Who is our legal contact?
- What does our cyber insurance require?
- Who contacts the insurance carrier?
- Are we required to use approved response vendors?

About Technology

- Can we isolate compromised systems quickly?
- Do we retain useful security logs?
- Can we revoke cloud sessions?

- Can we reset credentials quickly?
- Are backups isolated and tested?

About Recovery

- Which systems must be restored first?
- How long can the business operate without them?
- Have we tested restoration?
- Do we have manual workarounds?

About Practice

- Have we conducted a tabletop exercise?
- Did leadership participate?
- Were identified weaknesses corrected?

6.49 Questions to Ask Your IT Provider

Ask your IT provider practical questions about ransomware response, after-hours contact, system isolation, cloud sessions, logging, forensic evidence, backups, communications, incident documentation, and recovery responsibility.

1. What should we do first if we suspect ransomware?
2. Who do we contact after business hours?
3. Can you isolate systems remotely?
4. How quickly can compromised accounts be disabled?
5. Can cloud sessions and tokens be revoked?
6. How long are security logs retained?
7. Do you preserve forensic evidence?
8. Do you provide incident-response services?
9. If not, who do you recommend?
10. Can you restore our critical systems from backup?
11. When was our last successful restoration test?
12. Are backups protected from ordinary administrator accounts?
13. Do you maintain emergency contact procedures?
14. Do you help with tabletop exercises?
15. How would you communicate with us if our email system were unavailable?
16. How are incidents documented?

17. What responsibilities belong to us during an incident?
18. Who decides when systems are safe to reconnect?
19. How do you verify that attackers have been removed?
20. What post-incident improvements do you recommend?

6.50 Chapter 6 Cybersecurity Checklist

Use this checklist to review the organization's incident-response plan, communication methods, technical response capabilities, backups and recovery, legal and insurance preparation, external support, documentation, and practice.

Incident-Response Plan

☐ We have a written incident-response plan.
☐ The plan identifies response roles.
☐ Executive leadership knows its responsibilities.
☐ Emergency contacts are current.
☐ The plan is stored outside production systems.

Communication

☐ We have an alternate communication method.
☐ Employees know how to report incidents.
☐ A communications lead is designated.
☐ Public statements are coordinated.

Technical Response

☐ Compromised accounts can be disabled quickly.
☐ Systems can be isolated.
☐ Security logs are retained.
☐ Cloud sessions can be revoked.
☐ Administrator activity can be reviewed.

Backups and Recovery

☐ Critical systems are backed up.
☐ Backup copies are protected.
☐ Restore procedures are documented.
☐ Restoration is tested.
☐ Recovery priorities are defined.

Legal and Insurance

☐ Legal counsel is identified.
☐ Cyber insurance contact information is available.
☐ Insurance notification procedures are understood.
☐ Regulatory requirements are considered.

External Support

☐ Incident-response specialists are identified.
☐ Managed service provider contacts are current.
☐ Law enforcement contacts are available where appropriate.

Documentation

☐ Incident timelines are maintained.
☐ Major decisions are documented.
☐ Evidence preservation is considered.

Practice

☐ Tabletop exercises are conducted.
☐ Leadership participates.
☐ Lessons learned are documented.
☐ Improvements are tracked to completion.

6.51 Key Terms

Business Continuity: The ability of an organization to continue critical operations during and after a disruptive event.

Containment: Actions taken to limit the spread or impact of a cybersecurity incident.

Cybersecurity Incident: An event that threatens or compromises the confidentiality, integrity, or availability of information or systems.

Detection: The process of identifying potentially malicious or unauthorized activity.

Eradication: The process of removing malicious software, unauthorized access, compromised accounts, or other causes of an incident.

Incident Response: The structured process used to identify, contain, investigate, eliminate, recover from, and learn from cybersecurity incidents.

Incident-Response Plan: A documented plan defining how an organization will respond to cybersecurity incidents.

Recovery: The process of safely restoring systems, applications, information, and business operations after an incident.

Recovery Point Objective (RPO): The maximum amount of data loss an organization is prepared to tolerate.

Recovery Time Objective (RTO): The target amount of time within which a system or business function should be restored.

Tabletop Exercise: A discussion-based practice exercise used to test how an organization would respond to a hypothetical incident.

Threat Hunting: The proactive search for evidence of malicious activity that may not have triggered automated alerts.

6.52 Key Takeaways

Cybersecurity incidents cannot always be prevented.

Organizations therefore need the ability to respond.

Effective incident response depends on preparation, clear roles, fast detection, controlled containment, evidence preservation, complete eradication, tested recovery, careful communication, legal and insurance coordination, and lessons learned.

The first minutes and hours matter. So do the days and weeks afterward.

The strongest organizations do not improvise every decision during a crisis. They prepare. They practice. They document. They learn.

6.53 Chapter Summary

Chapter 5 focused on building cybersecurity into everyday business operations.

Chapter 6 has focused on what happens when those safeguards are not enough.

Incident response should be understood as a cycle: Prepare -> Detect -> Analyze -> Contain -> Eradicate -> Recover -> Learn -> Improve.

A cybersecurity incident can become a business crisis very quickly. But preparation changes the outcome.

An organization that knows who is in charge, whom to call, how to communicate, how to isolate systems, how to preserve evidence, how to restore from backup, and how to make business decisions is in a far stronger position than one trying to invent a response under pressure.

The objective is not to guarantee that incidents never happen.

The objective is to ensure that when something goes wrong, the organization can respond quickly, limit the damage, restore operations, and emerge stronger.

And that leads naturally to the next major question: If cybersecurity incidents can affect customers, employees, vendors, regulators, and business partners, what responsibilities does an organization have for protecting sensitive information in the first place?

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

Chapter 7

Protecting Sensitive Data: Privacy, Confidentiality, and Information Handling

Authored by Dr. Charles Edwards

Every organization has information that matters.

Some of it belongs to the business. Some belongs to employees, customers, vendors, partners, or other organizations.

It may include financial records, contracts, passwords, payroll data, customer files, tax records, intellectual property, medical information, legal documents, business plans, and confidential communications.

Cybersecurity is not only about protecting computers. It is about protecting the information those computers and cloud systems hold and process.

A damaged laptop can be replaced. A server can be rebuilt. An application can be restored. But once confidential information has been copied, stolen, or publicly exposed, the organization may never be able to take it back.

That is why data protection deserves its own attention.

Leadership should be able to answer some basic questions: What information do we have? Where is it stored? Who can access it? How sensitive is it? How long should we keep it?

Those questions form the foundation of practical information protection.

7.1 Data Is a Business Asset

Organizations often think carefully about physical assets. They know how many vehicles they own, where equipment is located, who has keys, what inventory is stored, and which assets are insured.

Information should be treated with the same seriousness.

Customer databases, employee records, contracts, and intellectual property all have business value.

That value may come from confidentiality, accuracy, availability, regulatory importance, operational usefulness, or competitive advantage.

If information is lost, altered, stolen, or exposed, the business may suffer.

Therefore, data should be managed as an asset.

7.2 Not All Data Is Equal

One of the most important ideas in information protection is that not all data requires the same level of security.

A public brochure and a payroll database are both information. But they do not have the same sensitivity.

A press release is intended to be seen. An employee Social Security number is not.

A product advertisement may be public. A confidential acquisition plan may not be.

Organizations need a practical way to distinguish between different types of information.

This is where data classification becomes useful.

7.3 Data Classification

Data classification is the process of organizing information according to its sensitivity and business importance.

A simple classification model may include Public, Internal, Confidential, and Restricted.

Public

- advertisements
- public website content
- published job openings
- public press releases
- brochures

Internal

- internal procedures
- meeting notes
- ordinary internal announcements
- nonpublic operational documents

Confidential

- customer records
- employee information
- contracts
- financial documents
- internal investigations
- business plans

Restricted

- passwords
- encryption keys
- certain health information
- highly sensitive financial records
- government-controlled information
- privileged legal material

7.4 Classification Should Lead to Action

A classification label is useful only if it changes behavior.

If a document is classified as Confidential, employees should know where it can be stored, who can access it, how it can be shared, whether it can be emailed, whether it can be printed, how long it should be retained, and how it should be destroyed.

Otherwise classification becomes paperwork.

The purpose is not to put labels on documents. The purpose is to guide protection.

7.5 Know Where Sensitive Information Lives

Many organizations cannot answer a basic question: Where is our sensitive data?

It may exist in file servers, cloud storage, laptops, email, databases, mobile devices, USB drives, backup systems, collaboration platforms, customer portals, and third-party applications.

Copies may also exist in places nobody expects.

An employee downloads a spreadsheet to a laptop. Another emails it to a colleague. A third uploads it to cloud storage. A fourth prints it.

Now one piece of information exists in several locations. Each copy becomes something that must be protected.

7.6 Data Sprawl

The uncontrolled spread of information across many systems is often called data sprawl.

Data sprawl makes security difficult because the organization may lose track of how many copies exist, where they are stored, who can access them, and whether old copies are still needed.

This is especially common with email attachments, spreadsheets, shared drives, cloud storage, and personal devices.

Reducing unnecessary copies is one of the simplest ways to reduce risk.

If information does not need to exist in five places, do not store it in five places.

7.7 Collect Only What You Need

Organizations sometimes collect more information than necessary.

A business may ask for date of birth, personal address, Social Security number, driver's license information, banking details, or other sensitive information.

But does the business truly need all of it?

Every piece of information collected creates responsibility.

If the organization does not need the data, there is little benefit in accepting the risk of storing it.

A strong principle is: Do not collect sensitive information merely because you can.

Collect what is necessary. Keep it only as long as necessary. Protect it while you have it.

7.8 Data Minimization

This idea is known as data minimization.

Data minimization means limiting the collection, storage, and use of personal or sensitive information to what is actually needed.

For example, if a business only needs to verify that a customer is over 18, it may not need to permanently store the customer's complete date of birth.

If a vendor payment system already stores banking details securely, employees may not need separate spreadsheets containing the same information.

Reducing unnecessary data reduces exposure.

7.9 Data Retention

Organizations should decide how long information must be kept.

Keeping information forever may feel safe. In reality, it can create unnecessary risk.

Old data may remain in archived email, backup systems, shared folders, retired servers, employee laptops, and legacy applications.

The business may have no operational need for the information. Yet it remains exposed.

A data retention policy defines how long different types of information should be kept.

Retention decisions may be influenced by legal requirements, contractual obligations, tax requirements, regulatory rules, business needs, and litigation holds.

The objective is not to delete information irresponsibly. The objective is to avoid keeping sensitive information indefinitely without reason.

7.10 Deleting a File Is Not Always the Same as Destroying It

When a user deletes a file, the information may not immediately disappear.

It may remain in backups, cloud recycle bins, storage media, email archives, or other copies.

Organizations should understand what secure disposal means for different media.

For example, paper may require shredding, hard drives may require secure wiping or physical destruction, cloud data may require documented deletion procedures, and backup retention may need to expire before all copies are removed.

Information disposal should be intentional.

7.11 Confidentiality, Integrity, and Availability

Three concepts are often used to describe the basic goals of information security: confidentiality, integrity, and availability.

These three goals are sometimes called the CIA triad.

Consider payroll. Confidentiality means unauthorized employees should not see payroll information. Integrity means attackers should not be able to change salary or banking records. Availability means payroll staff must be able to access the system when payroll is due.

Security must protect all three.

7.12 Confidentiality Is More Than Encryption

Encryption is an important confidentiality control.

But confidentiality also depends on access permissions, physical security, employee behavior, sharing practices, authentication, and data handling.

An encrypted laptop provides little benefit if an employee emails confidential files to the wrong person.

A secure database provides little benefit if 300 employees can access information needed by only 10.

Confidentiality requires multiple controls working together.

7.13 Integrity Matters Too

Data theft receives a great deal of attention. But unauthorized modification can also be damaging.

Imagine an attacker who changes customer bank information, payroll records, vendor payment details, product formulas, financial statements, or inventory records.

The information still exists. But it is no longer trustworthy.

Organizations need controls that help protect integrity, including access control, audit logs, digital signatures, approval workflows, version history, and separation of duties.

7.14 Availability Is a Security Requirement

Information that cannot be accessed when needed may create serious operational problems.

Ransomware attacks availability. So do system failures, corrupted databases, deleted files, cloud outages, denial-of-service attacks, and damaged infrastructure.

This is why backup, redundancy, business continuity, and disaster recovery are part of cybersecurity.

Security does not mean locking information away so tightly that nobody can use it. Security means authorized people can use information safely when they need it.

7.15 Who Should Have Access?

Access should be based on business need.

Organizations should ask: Who needs this information to perform their job? Not: Who might someday want access?

The difference matters.

For example, a payroll database may be restricted to payroll staff, selected HR personnel, authorized finance leadership, and certain administrators.

Most employees do not need access. Limiting access reduces exposure.

7.16 Need-to-Know

The need-to-know principle means information should be available only to people who require it for legitimate business responsibilities.

Need-to-know works together with least privilege.

Least privilege limits what systems a user can access. Need-to-know limits what information the user should see.

These concepts help reduce unnecessary access.

7.17 Access Control Lists and Permissions

Modern systems use permissions to control access.

A folder might allow one group to read files, another to read and modify, and an administrator to manage permissions.

Permissions should be reviewed periodically.

Over time, shared folders often accumulate old users, old groups, contractors, former employees, and excessive access.

A folder that was properly secured three years ago may no longer be properly secured today.

7.18 Sharing Links Can Create Exposure

Cloud platforms make it easy to share information.

A user clicks Share and creates a link. But what does the link allow?

Possibilities may include anyone on the Internet, anyone in the company, specific users, people with the link, or invited accounts only.

Employees may choose broad sharing because it is convenient. That can create risk.

Highly sensitive information should generally use more restrictive sharing.

7.19 Public Links Should Be Used Carefully

A public sharing link may be appropriate for a public brochure.

It is usually inappropriate for employee records, contracts, customer financial information, medical information, internal investigations, or restricted business documents.

Organizations should configure sharing defaults carefully.

The safest default is often narrower access. Employees can expand access when business needs justify it.

7.20 Email and Sensitive Information

Email is one of the most common ways sensitive information moves between people.

It is also a common source of accidental disclosure.

Employees may send to the wrong recipient, attach the wrong file, forward information, use personal email, or leave sensitive messages in mailboxes indefinitely.

Organizations should determine what types of information are acceptable for ordinary email and when more secure methods are required.

7.21 Verify the Recipient Before Sending

One of the simplest information-handling habits is also one of the most useful: Check the recipient before sending sensitive information.

Autocomplete can create mistakes.

A user named “John” should not automatically receive access to every system merely because the person works for the organization.

There may be no hacker involved. It is still a data incident.

Security includes preventing accidental disclosure.

7.22 Encryption in Transit

Information moving across networks should be protected when appropriate.

Encryption in transit protects information while it moves from one location to another.

Examples include secure websites using HTTPS, encrypted VPN connections, secure file transfer, and protected email connections.

Without encryption, network traffic may be easier to intercept.

Modern business systems should use secure protocols. Old and insecure communication methods should be removed where practical.

7.23 Encryption at Rest

Encryption at rest protects stored information.

Examples include encrypted laptop drives, encrypted databases, encrypted mobile devices, encrypted backups, and encrypted cloud storage.

Encryption at rest becomes especially important if a device is lost or stolen.

A stolen encrypted laptop may expose hardware. A stolen unencrypted laptop may expose every file it contains.

7.24 Encryption Is Only as Strong as Key Protection

Encryption uses cryptographic keys.

If an attacker obtains the key, the protection may be defeated.

Encryption keys, certificates, and secrets should therefore be protected carefully.

They should not be casually stored in spreadsheets, source code, email, text files, or shared documents.

Secrets deserve stronger controls than ordinary data.

7.25 Passwords Are Sensitive Data

Passwords should never be treated like ordinary information.

Organizations should avoid storing passwords in documents, email, notebooks, spreadsheets, or shared folders.

Where possible, users should use password managers.

Applications should store passwords securely using appropriate cryptographic methods rather than keeping readable copies.

An exposed password can become the doorway to many systems.

7.26 API Keys, Tokens, and Secrets

Modern applications use more than passwords.

They may use API keys, access tokens, secret keys, certificates, and service-account credentials.

These items may provide powerful access.

Developers sometimes accidentally place them in source-code repositories or configuration files.

Organizations should manage these secrets deliberately.

A secret that provides system access should be treated like a highly sensitive credential.

7.27 Printing Sensitive Information

Digital information is not the only concern.

Printed documents may contain payroll information, employee records, customer data, legal documents, medical information, or financial records.

Printed information may be left on printers, in conference rooms, on desks, in trash cans, or in vehicles.

Paper needs protection too.

Organizations should consider secure printing, locked storage, clean-desk expectations, and shredding.

7.28 Clean Desk Practices

A clean desk policy does not mean every desk must look perfect.

It means sensitive information should not be left exposed unnecessarily.

Employees should avoid leaving confidential documents overnight, in public work areas, near visitors, or where unauthorized employees can view them.

Simple physical habits can reduce information exposure.

7.29 Screen Privacy

Sensitive information can also be exposed visually.

Employees working in airports, coffee shops, shared spaces, customer sites, or conferences should consider who can see their screens.

Screen locking, privacy filters, and physical awareness can help.

This risk is sometimes called shoulder surfing.

Not every data breach requires malware. Sometimes an unauthorized person simply sees what should have remained private.

7.30 Lost and Stolen Devices

Laptops, tablets, phones, and USB drives can be lost.

Organizations should assume that some devices will eventually disappear.

Protection may include encryption, screen locks, MFA, remote wipe, device management, limited local data, and prompt reporting.

The question should not be: Will anyone ever lose a device?

The better question is: What happens when someone does?

7.31 USB Drives and Removable Media

Removable media can create several risks.

USB drives may be lost, carry malware, contain unencrypted files, or be copied without tracking.

Organizations should decide whether removable media is allowed and under what conditions.

Highly sensitive data should not be casually copied to unprotected USB drives.

7.32 Personal Cloud Storage

Employees may use personal cloud services because they are convenient.

They may upload company information to personal Google Drive accounts, Dropbox, personal OneDrive, or other services.

This can create problems. The organization may lose control, visibility, retention, and the ability to remove access later.

Employees should know which storage platforms are approved.

7.33 Personal Email

Sending company information to personal email can create similar problems.

The business may no longer control access, retention, forwarding, deletion, or security.

Sensitive company information should remain in approved business systems unless there is a legitimate and authorized reason otherwise.

7.34 Artificial Intelligence Tools and Sensitive Data

AI tools are increasingly used for writing, summarization, coding, analysis, research, and brainstorming.

They can be valuable.

But employees should understand what information may be entered into these systems.

Sensitive data should not be copied into an AI service unless the organization has evaluated and approved that use.

Employees should not casually paste customer records, confidential contracts, passwords, proprietary source code, employee information, or restricted business data into unapproved tools.

New technology does not eliminate old data-handling responsibilities.

7.35 Data Loss Prevention

Some organizations use Data Loss Prevention, or DLP, technologies.

DLP tools may help detect or prevent sensitive information from being emailed externally, uploaded to unauthorized cloud services, copied to removable media, printed, or shared improperly.

DLP can be useful. But it is not perfect.

It should support policies and employee training. Technology alone cannot understand every business context.

7.36 Monitoring Sensitive Data Access

Organizations should consider logging access to highly sensitive systems.

For example: who opened the payroll database, who downloaded customer records, who exported financial information, who changed permissions, and who shared files externally.

These logs may help detect unusual behavior. They can also support investigations.

For its most important systems, a business should be able to answer the question: Who accessed this information?

7.37 Large Downloads Deserve Attention

Suppose an employee usually accesses 20 customer records per day. Suddenly the account downloads 75,000 records at midnight.

That does not automatically prove malicious activity. But it deserves attention.

Large or unusual data transfers may indicate account compromise, insider misuse, accidental bulk export, or legitimate business activity.

Monitoring should focus on meaningful anomalies.

7.38 Insider Risk

Not every information-security incident is caused by an outside attacker.

Insider risk may involve malicious employees, careless employees, contractors, former employees, or authorized users whose accounts are compromised.

An employee may deliberately steal information before leaving. Another may accidentally send confidential data to the wrong recipient. Another may upload files to personal storage for convenience.

The organization should design controls around behavior and access, not assumptions about intent.

7.39 Insider Threat Does Not Mean Distrusting Everyone

Organizations sometimes misunderstand insider risk.

The goal is not to treat every employee as suspicious.

The goal is to build reasonable controls.

Examples include limiting unnecessary access, monitoring high-risk activity, separating duties, reviewing downloads, controlling removable media, and removing access promptly when employees leave.

Trust and control can coexist.

7.40 Third Parties and Data

Organizations frequently share data with payroll companies, accountants, law firms, cloud providers, software vendors, consultants, and business partners.

Before sharing sensitive information, organizations should ask what data the vendor will receive, why the vendor needs it, how it will be protected, who at the vendor can

access it, how long the vendor will keep it, what happens when the contract ends, and how the vendor will notify the organization of a breach.

Vendor data handling is part of the organization's risk.

7.41 Data Processing Agreements and Contracts

Contracts should address security and data responsibilities when appropriate.

Topics may include confidentiality, security controls, breach notification, data retention, subcontractors, destruction of data, audit rights, and regulatory obligations.

The exact requirements depend on the relationship and applicable law.

The important point is that security expectations should not remain informal.

7.42 Privacy and Security Are Related but Different

Privacy and security overlap, but they are not identical.

Security focuses on protecting information from unauthorized access, alteration, destruction, or loss.

Privacy focuses more broadly on how personal information is collected, used, shared, and retained.

A company can have strong technical security and still misuse personal information.

Good information governance considers both privacy and security.

7.43 Personal Information

Organizations should understand what kinds of information may identify or relate to individuals.

Examples may include name, home address, email address, telephone number, financial information, employment information, government identifiers, health information, online identifiers, and other personal data.

The sensitivity of the information may vary. Some categories require stronger protection.

7.44 Privacy Expectations Matter

Even when a particular action is technically legal, customers and employees may have expectations about how their information will be used.

Trust matters.

Organizations should be transparent about what they collect, why they collect it, how it is used, who it is shared with, and how long it is kept.

Surprising people with unexpected uses of their information can damage trust.

7.45 Regulatory Requirements

Different industries and jurisdictions may impose requirements for protecting certain information.

Examples may involve health information, financial information, payment card data, student information, government data, or personal information.

Organizations should identify which requirements apply to them.

Cybersecurity controls should support those obligations.

Compliance should not be treated as identical to security. Passing an audit does not guarantee that an organization is secure. But compliance requirements can provide useful minimum standards.

7.46 Compliance Is the Floor, Not the Ceiling

A business may satisfy a regulatory requirement and still face significant risk.

For example, a control may be required once per year, but the threat environment may require more frequent review.

A regulation may specify minimum password standards, but the organization may choose stronger authentication because the business risk is high.

The objective should not be: What is the minimum we can do?

The better question is: What protection does the risk actually require?

7.47 Data Breaches

A data breach occurs when protected information is accessed, disclosed, stolen, or exposed without authorization.

A data breach may result from a stolen laptop, compromised account, misdirected email, cloud-sharing error, malicious insider, lost mobile device, vendor incident, insecure application, or malware infection.

Not every cybersecurity incident becomes a data breach. But many do.

Organizations need procedures for determining what information may have been affected.

7.48 Breach Investigation

When sensitive data may have been exposed, the organization may need to determine what information was involved, how many people were affected, whether data was viewed, whether it was copied, whether it was altered, how long exposure lasted, and who may have received it.

This investigation may involve logs, forensic analysis, interviews, email review, cloud activity, and third-party assistance.

Accurate information matters because notification decisions may depend on the facts.

7.49 Accidental Disclosure Is Still a Security Incident

Consider an employee who sends a spreadsheet containing 500 customer records to the wrong external recipient.

No hacker was involved. No malware was involved. The data was still disclosed without authorization.

The organization should respond seriously.

It may need to contact the recipient, request deletion, determine what information was exposed, involve legal counsel, document the incident, and evaluate notification requirements.

Accidents can create real consequences.

7.50 A Business Scenario: The Shared Folder

Consider a fictional company called Summit Benefits Group.

Summit stores employee benefit enrollment documents in a cloud platform.

The files contain names, addresses, dates of birth, insurance information, and other personal data.

A staff member needs to share one document with an outside benefits consultant.

Instead of sharing one file with one named person, the employee changes the entire folder setting to: Anyone with the link can view.

The consultant accesses the document successfully. The employee forgets to change the setting back.

The link is later forwarded to another person. Several months pass.

The organization discovers that the folder has effectively been accessible to anyone who possessed the link.

No software vulnerability existed. No hacker defeated encryption.

The problem was a sharing configuration.

7.51 What Summit Should Do

Summit should restrict the folder immediately, review access logs, determine who accessed the information, identify what data was exposed, preserve evidence, involve legal or privacy counsel where appropriate, determine whether notification is required, review similar sharing settings, train employees, and improve default sharing controls.

The incident demonstrates an important lesson: Data security depends on configuration and behavior as much as technology.

7.52 Another Scenario: The Departing Sales Manager

A sales manager leaves the company for a competitor.

During the week before departure, the manager downloads the customer list, pricing information, sales forecasts, and contract details.

The manager's account has legitimate access to all these files. Nothing is technically hacked.

But the downloading pattern is unusual.

If the organization monitors sensitive exports, the activity may be detected. If not, the business may never know what left.

This is one reason insider risk and data access monitoring matter.

7.53 Questions Every Business Owner Should Ask

About Data

- What sensitive information do we collect?
- Why do we collect it?
- Where is it stored?

- Who can access it?
- How long do we keep it?

About Classification

- Do employees know which information is sensitive?
- Do we have practical classification categories?
- Do the categories affect how information is handled?

About Sharing

- Can users create public links?
- Are sharing permissions reviewed?
- Is sensitive information sent through ordinary email?
- Can employees use personal cloud storage?

About Devices

- Are laptops encrypted?
- Are mobile devices protected?
- Can lost devices be remotely managed?
- Is removable media controlled?

About Vendors

- Which vendors receive sensitive information?
- Do contracts define security responsibilities?
- How do vendors notify us of incidents?
- What happens to our data when the relationship ends?

About Retention

- How long do we keep sensitive information?
- Is old data deleted?
- Are backups included in retention planning?

About Monitoring

- Can we tell who accessed important data?
- Can we detect large downloads?
- Are unusual exports reviewed?

7.54 Questions to Ask Your IT Provider

1. Where is our sensitive information stored?
2. Which cloud systems contain confidential data?
3. Are company laptops encrypted?
4. Are backups encrypted?
5. Can users create public sharing links?
6. Are external sharing settings reviewed?
7. How are file permissions managed?
8. Can we identify who accessed sensitive files?
9. Do we monitor large downloads or exports?
10. Are USB devices restricted?
11. How are lost devices handled?
12. Are mobile devices centrally managed?
13. How do you protect encryption keys and administrative credentials?
14. Do you help identify old or unnecessary data?
15. Do you review inactive accounts with data access?
16. How are sensitive files securely deleted?
17. How is cloud data retained and removed?
18. Do you provide data-loss-prevention capabilities?
19. How would we investigate a suspected data leak?
20. What data-protection responsibilities remain with us?

7.55 Chapter 7 Cybersecurity Checklist

Data Inventory

☐ We know what sensitive information we collect.
☐ We know where sensitive information is stored.
☐ We know which systems contain confidential data.
☐ We identify unnecessary copies.

Classification

☐ We use practical data-classification categories.
☐ Employees understand the categories.
☐ Classification affects handling and sharing.

Access

☐ Sensitive data access is limited.
☐ Permissions are reviewed periodically.
☐ Former employees lose access promptly.
☐ Vendor access is limited.

Sharing

☐ Public sharing links are controlled.
☐ External sharing is reviewed.
☐ Sensitive information is not sent casually by email.
☐ Personal cloud storage is restricted.

Devices

☐ Company laptops use encryption.
☐ Mobile devices use screen locks.
☐ Lost devices can be managed.
☐ Sensitive data on removable media is controlled.

Credentials and Secrets

☐ Passwords are not stored in documents.
☐ Password managers are used where appropriate.
☐ API keys and tokens are protected.
☐ Encryption keys are restricted.

Retention

☐ We have data-retention rules.
☐ Old data is reviewed.
☐ Unnecessary sensitive information is removed.
☐ Disposal methods are appropriate.

Monitoring

☐ Access to highly sensitive systems is logged.
☐ Large downloads are monitored where appropriate.
☐ Unusual external sharing is reviewed.

Vendors

☐ Vendors receive only necessary data.

☐ Contracts address confidentiality and security.
☐ Vendor retention is understood.
☐ Vendor access is removed when no longer needed.

Incident Readiness

☐ Accidental disclosures are reported.
☐ Data breaches are investigated.
☐ Legal and privacy contacts are identified.
☐ Evidence can be preserved.

7.56 Key Terms

CIA Triad: The three foundational information-security goals of confidentiality, integrity, and availability.

Confidentiality: The protection of information from unauthorized access or disclosure.

Data Breach: An incident in which protected information is accessed, disclosed, stolen, or exposed without authorization.

Data Classification: The process of categorizing information according to sensitivity and required protection.

Data Loss Prevention (DLP): Technology and processes designed to detect or prevent sensitive information from being transferred or shared improperly.

Data Minimization: The practice of collecting, using, and retaining only the information necessary for a legitimate purpose.

Data Retention: The period for which information is kept before it is archived or destroyed.

Data Sprawl: The uncontrolled distribution of information across multiple systems, locations, or copies.

Encryption at Rest: The protection of stored information through encryption.

Encryption in Transit: The protection of information while it moves across networks.

Integrity: The assurance that information remains accurate and has not been altered without authorization.

Need-to-Know: The principle that information should be available only to people who require it for legitimate responsibilities.

Personal Information: Information that identifies, relates to, or can be associated with an individual.

Privacy: The principles and practices governing how personal information is collected, used, shared, and retained.

Restricted Data: Highly sensitive information requiring strong protection and tightly controlled access.

Shoulder Surfing: The unauthorized viewing of sensitive information by observing a person's screen, keyboard, or documents.

7.57 Key Takeaways

Sensitive information is one of the organization's most valuable assets.

Protecting it requires more than antivirus software or firewalls.

Organizations should understand what information they collect, where it is stored, who can access it, how it is shared, how long it is kept, and how it is destroyed.

Not all data requires the same protection. Classification helps organizations apply appropriate controls.

Data minimization reduces exposure. Encryption protects stored and transmitted information. Access control limits who can see sensitive data. Monitoring helps detect unusual behavior. Retention policies reduce unnecessary storage. Vendor agreements help extend protection beyond the organization.

Privacy and security must work together.

The objective is simple: Protect sensitive information throughout its entire life cycle.

7.58 Chapter Summary

Chapter 6 focused on responding when something goes wrong.

Chapter 7 has focused on one of the most important things organizations are trying to protect: information.

A strong information-protection process looks like this: Identify the data -> Classify it -> Limit collection -> Control access -> Protect storage -> Protect transmission -> Monitor use -> Limit sharing -> Retain only as necessary -> Destroy securely.

Cybersecurity becomes much easier to manage when organizations know which information truly matters.

Not every file needs maximum security. But sensitive information should not be treated casually.

Organizations should build controls that protect confidentiality, preserve integrity, and maintain availability.

They should also understand privacy expectations and legal obligations.

Once information has been copied, stolen, or exposed, it may be impossible to fully recover control of it.

That is why data protection should happen before an incident.

And that leads naturally to the next major question: If modern businesses depend on cloud platforms, mobile devices, remote work, and outside providers, how do we protect information and systems when the traditional office network is no longer the center of everything?

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

Chapter 8

Securing the Modern Workplace: Cloud, Remote Work, Mobile Devices, and Zero Trust

Authored by Dr. Charles Edwards

Modern business no longer happens in one building, on one network, using one set of company-owned computers.

Employees work from home, customer locations, hotels, airports, and shared offices. Applications live in the cloud. Executives approve transactions from phones and tablets. Vendors connect remotely.

That flexibility has changed business for the better in many ways. It has also changed cybersecurity.

For years, many organizations relied on a simple model: inside the office network was trusted and the Internet was untrusted. A firewall separated the two.

That model no longer matches the way most organizations work.

An employee at home may reach the same cloud application as an employee sitting in headquarters. A criminal with stolen credentials may reach it too.

The question is therefore no longer simply whether a device is inside the building.

The better questions are: Who is the user? What device is being used? Is the device healthy? What information is being requested? Does this access make sense right now?

This chapter looks at how to protect a workplace where people, devices, applications, and information are spread across many locations and services.

8.1 The Office Is No Longer the Security Boundary

For decades, many organizations organized security around the physical office and the corporate network. Employees connected desktop computers to internal switches, servers were located in a server room, and the firewall controlled access to and from the Internet.

That environment still exists in some organizations, but it is no longer the whole picture. A typical employee may begin the day on a home Wi-Fi network, connect to Microsoft 365, join a video meeting, access a cloud-based customer relationship management system, open files from a shared cloud drive, and later connect from a mobile phone.

In that environment, the traditional network boundary becomes less meaningful. The employee may never connect to the corporate network at all.

Cybersecurity therefore has to follow the user, the device, the application, and the data rather than assuming that location alone creates trust.

8.2 Cloud Computing Changes Responsibility, Not the Need for Security

Cloud computing allows organizations to use computing resources that are operated by outside providers. These resources may include email, file storage, databases, servers, applications, backup systems, identity platforms, and development environments.

One of the most common misunderstandings about cloud security is the belief that moving a system to the cloud transfers all security responsibility to the cloud provider.

It does not.

The provider may secure the physical data centers, underlying hardware, and portions of the platform. The customer still has important responsibilities for how accounts, permissions, data, devices, applications, and configurations are managed.

This division of responsibility is often described as a shared responsibility model. The exact boundary varies by service, but the principle is consistent: the provider secures some parts of the environment, while the customer remains responsible for others.

8.3 The Shared Responsibility Model

Consider a company using a cloud email platform. The provider may protect the data center, maintain the underlying servers, and operate the email service. But the provider does not know whether a departing employee should still have an account. It does not know whether a particular administrator truly needs global access. It cannot decide whether a suspicious login should be approved. It cannot prevent an employee from sharing confidential information with the wrong person if the customer has permitted that sharing.

The same idea applies to cloud infrastructure. A provider may operate secure physical servers, but a customer can still create an Internet-facing database, use weak access controls, leave a storage container public, or fail to patch software running inside a virtual machine.

Cloud security therefore requires understanding what the provider handles and what the organization must handle.

8.4 Identity Becomes the New Perimeter

When applications are spread across cloud services and employees work from many locations, identity becomes one of the most important security controls.

An identity is more than a username. It represents a person, device, service, or application and determines what resources can be accessed.

If an attacker steals the right identity, the attacker may be able to access company information without ever touching the company's physical office or internal network.

This is why modern cybersecurity places so much emphasis on strong authentication, multifactor authentication, conditional access, least privilege, and monitoring of sign-in behavior.

8.5 Passwords Alone Are Not Enough

Passwords remain common because they are familiar and easy to deploy. But passwords have serious weaknesses. People reuse them. Attackers steal them through phishing. Criminals purchase them from data breaches. Malware can capture them. Employees sometimes store them insecurely.

For this reason, important systems should not depend on a password alone.

Multifactor authentication adds another layer. Even if the password is stolen, the attacker must still satisfy another requirement.

Organizations should especially prioritize MFA for email, cloud administration, remote access, financial systems, backup systems, and any application containing sensitive information.

8.6 Strong MFA Matters

Not all forms of multifactor authentication provide the same level of protection.

Simple approval prompts can be abused through MFA fatigue, as discussed earlier in this book. Text messages can also be vulnerable to certain forms of interception or account takeover.

Stronger approaches may include authentication applications using number matching, device-bound authentication, passkeys, or hardware security keys.

The important point for business leaders is not to memorize every authentication technology. It is to understand that MFA should be selected and configured thoughtfully, especially for highly privileged accounts.

8.7 Conditional Access

Modern identity platforms can evaluate more than a password.

They may consider the user's location, device, application, risk level, and requested action before granting access.

For example, an employee may be allowed to read ordinary email from a managed laptop at home, but access to financial systems from an unknown device in another country may require additional verification or be blocked entirely.

This concept is often called conditional access.

Conditional access allows security decisions to reflect context instead of treating every login as equally trustworthy.

8.8 Zero Trust: Never Automatically Trust, Always Verify

Zero trust is often summarized as “never trust, always verify.”

Zero Trust does not mean that organizations distrust their employees as people. It means that systems should not grant broad access simply because a user or device is located on a particular network.

Instead, access decisions should be based on evidence.

The organization asks whether the identity is valid, whether the device is acceptable, whether the request is consistent with policy, whether the user actually needs the resource, and whether the session remains trustworthy.

Trust becomes something that is continuously evaluated rather than permanently granted.

8.9 Zero Trust Is Not a Product

Vendors sometimes market Zero Trust as though it were a single product that can be purchased and installed.

That is misleading.

Zero Trust is better understood as an architecture and operating philosophy.

It may involve identity management, MFA, device management, segmentation, least privilege, monitoring, encryption, and application controls. No single tool provides all of these by itself.

An organization can begin applying Zero Trust principles gradually. It does not need to redesign everything at once.

8.10 Least Privilege in the Modern Workplace

Least privilege becomes even more important when employees and applications operate across many cloud systems.

A user should receive only the access needed to perform legitimate work. An application should also receive only the permissions it needs.

This principle applies to employees, administrators, contractors, vendors, service accounts, and software integrations.

The less unnecessary access exists, the less an attacker inherits when one identity is compromised.

8.11 Privileged Accounts Need Additional Protection

Administrative accounts deserve special treatment because compromise of one privileged identity can affect an entire organization.

Administrators may be able to create users, reset passwords, change security policies, disable monitoring, access confidential data, and alter cloud configurations.

Organizations should minimize the number of privileged accounts, require strong MFA, separate ordinary and administrative identities where practical, and monitor administrative activity.

Privilege should be used when needed, not as a permanent convenience.

8.12 Privileged Access Should Be Temporary When Possible

One useful modern security practice is to provide elevated privileges only for the period in which they are needed.

An administrator may work with an ordinary account during the day and activate administrative privileges only when performing a specific task.

This reduces the amount of time powerful permissions are exposed.

Some organizations use technologies known as privileged access management or just-in-time access to support this approach.

8.13 Device Trust Matters

Identity is only one part of the access decision. The device also matters.

Imagine two login attempts using the same employee account. The first comes from a company-managed laptop with current security updates, disk encryption, endpoint protection, and a known configuration. The second comes from an unknown personal computer with outdated software.

Even though the username is the same, the risk is not.

Modern security systems can use device health as part of the access decision.

8.14 Managed Devices

A managed device is one the organization can configure, monitor, and protect through centralized tools.

Management may enforce requirements such as disk encryption, screen locking, operating system updates, endpoint protection, approved software, and remote wipe.

Central management also gives the organization visibility. It can answer questions such as whether a device is still compliant, when it last checked in, and whether required protections are active.

This becomes especially important when employees work outside the office.

8.15 Bring Your Own Device

Bring Your Own Device, or BYOD, allows employees to use personally owned devices for business purposes.

BYOD may reduce cost and increase convenience, but it creates important questions.

The organization may not fully control the device. Family members may also use it. Security updates may be delayed. Personal and business information may be mixed together. If the employee leaves, the company may need to remove business data without affecting personal information.

There is no universal rule that BYOD is always good or always bad. The decision should be intentional and supported by clear controls.

8.16 Mobile Device Management and Application Management

Organizations may use mobile device management to control company-owned devices or approved personal devices.

Some environments also use mobile application management, which focuses on protecting the business applications and data rather than controlling the entire personal device.

These technologies can help enforce encryption, screen locks, approved applications, data separation, remote wipe, and other requirements.

The correct approach depends on the organization's risk, workforce, and privacy expectations.

8.17 Remote Work Requires More Than a VPN

Virtual private networks became a common symbol of remote-work security. A VPN creates an encrypted connection between the user's device and another network or service.

VPNs remain useful in many environments, but they are not a complete security strategy.

If a criminal steals valid VPN credentials, the encrypted tunnel may simply provide the attacker with secure access to the network.

Organizations therefore need to combine remote access with MFA, device security, least privilege, monitoring, patching, and segmentation.

8.18 VPN Risk and Concentration of Access

A remote-access gateway can become a particularly valuable target because many employees depend on it.

If the system contains a serious vulnerability or weak authentication, attackers may gain a direct path into the organization.

Remote-access systems should therefore be treated as high-value infrastructure. They should be patched quickly, monitored carefully, protected with strong authentication, and exposed only as necessary.

An organization should know exactly which remote-access services are reachable from the Internet.

8.19 Zero Trust Network Access

Some organizations are replacing or supplementing traditional VPN access with approaches often called Zero Trust Network Access, or ZTNA.

Instead of placing the remote user broadly onto the internal network, ZTNA attempts to connect the user only to the specific application or service the user is authorized to reach.

This can reduce lateral movement because access is more narrowly scoped.

ZTNA is not automatically better in every environment, but the principle is valuable: do not provide more network access than the user actually needs.

8.20 Home Networks Should Be Treated as Untrusted

A company cannot realistically secure every employee's home network to the same standard as a corporate office.

Home routers vary in quality, update status, configuration, and age. Smart televisions, cameras, gaming systems, and other household devices may share the same network.

The better approach is to make company devices resilient even when connected to an untrusted network.

That means encryption, endpoint protection, local firewalls, secure authentication, patching, and safe remote-access methods should remain active regardless of location.

8.21 Public Wi-Fi

Public Wi-Fi should also be treated as untrusted.

Modern encryption and secure applications reduce some of the historical risks of public networks, but employees should still avoid assuming that a network is safe merely because it has the name of a hotel, airport, or conference.

Attackers can create convincing wireless network names. Public networks may also expose users to local threats.

Where appropriate, employees can use trusted mobile hotspots, secure VPN connections, or other approved access methods.

8.22 Cloud Storage and Collaboration

Cloud collaboration has made it easy for teams to work together from anywhere.

Employees can edit shared documents, collaborate in real time, and send links instead of large email attachments.

This convenience also changes the risk. Information may be shared outside the organization, links may remain active indefinitely, guests may retain access after a project ends, and sensitive files may be synchronized to multiple devices.

Organizations should configure sharing defaults carefully and review external collaboration periodically.

8.23 Guest Accounts and External Collaboration

Many cloud platforms allow outside users to participate as guests.

This can be extremely useful for clients, consultants, vendors, and partners.

But guest access should not become permanent by accident.

Organizations should know which external identities exist, what they can access, who sponsored the access, and whether the access is still required.

Periodic guest-access reviews can remove forgotten external users before they become a long-term risk.

8.24 SaaS Applications Multiply Quickly

Software as a Service applications can often be purchased or activated quickly.

A department may adopt a project-management tool, a marketing platform, an AI application, an analytics service, or a file-sharing system without significant involvement from IT.

This speed is valuable, but it can produce a large collection of applications that nobody centrally manages.

Each application may contain company data, user accounts, integrations, API tokens, and permissions.

Organizations need visibility into which SaaS applications are actually being used.

8.25 Shadow SaaS

When departments or employees adopt cloud applications without formal approval or oversight, the result is sometimes called shadow SaaS or shadow IT.

The security concern is not simply that the application was unauthorized. The deeper problem is loss of visibility.

The organization may not know what data has been uploaded, whether MFA is enabled, whether the vendor is trustworthy, whether former employees still have access, or how to retrieve the data if the service is discontinued.

Security teams should work with employees to understand why such tools are attractive and provide secure alternatives rather than relying only on prohibition.

8.26 Cloud Misconfiguration

Many cloud security incidents occur not because an attacker defeated the cloud provider, but because the customer's configuration was too permissive.

Examples include public storage, overly broad administrator roles, anonymous sharing, exposed databases, weak API permissions, and disabled logging.

Cloud platforms are powerful. That power creates many settings.

Organizations should establish secure configuration baselines and periodically review them.

8.27 Logging in Cloud Environments

Cloud systems produce valuable security logs. They may show sign-ins, administrator actions, file sharing, application consent, data downloads, configuration changes, and suspicious activity.

But logging is useful only if it is enabled, retained, and reviewed.

Organizations should understand what logs are available in each critical cloud service and how long those logs are retained.

During an incident, these records may be the only way to reconstruct what happened.

8.28 Application Permissions and OAuth Consent

Modern cloud applications often connect to one another through delegated permissions.

For example, an employee may authorize an application to read calendar information, access files, send email, or manage contacts.

These integrations can improve productivity, but they can also create risk if employees authorize untrusted applications or grant more permissions than necessary.

Organizations should monitor third-party application consent and restrict high-risk permissions.

8.29 API Security

Cloud environments depend heavily on application programming interfaces, or APIs.

APIs allow applications and services to exchange information automatically.

Because APIs often provide direct access to business data and functions, they must be authenticated, authorized, monitored, and protected like other important systems.

An exposed API key or overly permissive API can provide attackers with a powerful route into data without requiring an ordinary user login.

8.30 Endpoint Security Still Matters

The growth of cloud computing does not make endpoint security less important.

Employees still use laptops, desktops, phones, and tablets to reach cloud applications.

If one of those devices is compromised, attackers may steal sessions, passwords, files, or authentication tokens.

Organizations should continue to protect endpoints through patching, endpoint detection and response, anti-malware controls, disk encryption, local firewalls, and device management.

8.31 Endpoint Detection and Response

Traditional antivirus software focuses heavily on identifying known malicious files.

Endpoint Detection and Response, often called EDR, provides broader visibility into suspicious behavior on computers and servers.

EDR may detect unusual processes, credential theft, malicious scripts, ransomware behavior, or attacker movement.

For many organizations, EDR has become an important layer because modern attacks do not always rely on easily recognized malware.

8.32 Patch Management for Remote Devices

Remote work can make patch management more difficult.

Devices may spend weeks away from the office. Employees may postpone restarts. Some systems may connect only occasionally.

Organizations should design patching so that important updates continue even when devices are remote.

A laptop that has not entered the building for three months should not become invisible to security management.

8.33 Browser Security

Web browsers have become one of the primary business platforms.

Employees use browsers to access email, banking, cloud storage, customer systems, administrative consoles, and SaaS applications.

That makes browser security important.

Organizations should keep browsers updated, control risky extensions, use secure configuration settings, and consider separating highly privileged administrative activity from ordinary browsing.

8.34 Mobile Phones Are Business Devices

Smartphones often contain email, messaging applications, authentication apps, business documents, and access to cloud systems.

They should therefore be treated as part of the security environment.

A lost phone may expose notifications, saved sessions, files, or authentication prompts if it is not properly protected.

Screen locks, encryption, biometric protection, remote management, and prompt reporting of lost devices can reduce risk.

8.35 SIM-Swapping and Phone-Based Recovery

Organizations should also recognize that phone numbers can become part of the security chain.

Some services use text messages or phone calls for authentication and account recovery.

Attackers may attempt SIM-swapping or social engineering against mobile carriers to take control of a victim's phone number.

For highly sensitive accounts, stronger authentication and recovery methods may be preferable to relying on a phone number alone.

8.36 Secure Access From Anywhere

The goal of modern workplace security is not to force employees back into one physical location simply because the office network feels easier to secure.

The goal is to create secure access regardless of location.

That requires consistent identity controls, managed devices, protected applications, strong authentication, narrow permissions, encryption, and monitoring.

Security should follow the work rather than depend entirely on where the work happens.

8.37 A Business Scenario: The Traveling Executive

Consider a fictional company called Meridian Advisory Partners.

One of Meridian's executives, Karen, travels frequently. She uses a company-managed laptop and a company phone.

While staying at a hotel, Karen receives a message that appears to come from the company's cloud email provider. It claims her session has expired and asks her to sign in again.

She enters her username and password into a convincing fake page.

The attacker now knows the password.

However, Meridian requires phishing-resistant MFA for executives and administrators. The attacker cannot complete authentication using the stolen password alone.

The identity platform also notices that the attempted login originates from a device that is not registered with the company.

The login is blocked and an alert is generated.

Karen's password is reset, and the event is investigated.

The phishing attempt succeeded in stealing a password, but multiple layers prevented account compromise.

8.38 What Meridian Did Right

Meridian did not rely on one control.

The company used a managed device, strong MFA, device-based access conditions, monitoring, and a process for investigating risky sign-ins.

This is a good example of defense in depth applied to the modern workplace.

Even when one layer failed, the entire security model did not collapse.

8.39 A Second Scenario: The Forgotten Guest Account

Now consider a design firm called Northline Creative.

Northline uses a cloud collaboration platform to work with clients and outside contractors.

A contractor is invited as a guest during a six-week project. The contractor receives access to a project folder and several internal discussion channels.

The project ends, but nobody removes the guest account.

Eighteen months later, the contractor's personal account is compromised in an unrelated phishing attack.

The criminal now inherits whatever access the forgotten guest account still has.

This is not a failure of the cloud provider. It is a life-cycle management failure.

External accounts require the same discipline as employee accounts: sponsorship, review, expiration, and removal.

8.40 Building a Practical Zero Trust Roadmap

Organizations sometimes delay Zero Trust because the concept appears too large.

A more practical approach is to implement the principles gradually.

Begin with identity. Require MFA for important accounts and reduce unnecessary privilege.

Then improve device visibility. Identify which devices are managed and enforce minimum security requirements.

Next, reduce broad network access. Segment sensitive resources and avoid giving remote users more access than necessary.

Improve monitoring so that unusual sign-ins, device changes, and administrator activity are visible.

Finally, continue tightening application permissions, external sharing, and third-party access.

Zero Trust is a direction of travel, not a one-day installation.

8.41 Questions Every Business Owner Should Ask

About Identity

- Is MFA required for email and important cloud systems?
- Do administrators use stronger authentication?
- Are risky sign-ins monitored?
- Are dormant accounts removed?

About Devices

- Which devices are company managed?
- Are laptops encrypted and patched?
- Can lost devices be remotely managed?
- Can unmanaged devices reach sensitive applications?

About Cloud Services

- Which cloud services contain company data?
- Who has administrator access?
- Are external sharing settings reviewed?
- Are security logs enabled and retained?

About Remote Work

- How do employees connect securely from home?
- Do remote users receive more network access than they need?
- Are remote-access systems patched and monitored?
- Do we know which services are Internet-facing?

About External Collaboration

- How many guest accounts exist?
- Who sponsors them?
- Do they expire?
- Are vendor and contractor accounts reviewed?

About Zero Trust

- Are we relying on network location as a substitute for authentication?
- Do we evaluate device health?
- Can access be limited by application and role?

- Do we continuously review privilege?

8.42 Questions to Ask Your IT Provider

1. Which of our business systems are cloud-based?
2. Who is responsible for securing each cloud service?
3. Is MFA enforced for all privileged accounts?
4. What types of MFA are being used?
5. Can we block access from unmanaged devices?
6. How are company devices centrally managed?
7. Are laptops encrypted and monitored?
8. How are remote devices patched?
9. What remote-access services are exposed to the Internet?
10. Do we use traditional VPN, Zero Trust Network Access, or both?
11. How is external guest access reviewed?
12. Can users create public sharing links?
13. Are third-party application permissions monitored?
14. How are API keys and service credentials protected?
15. What cloud logs are retained and for how long?
16. Can we detect suspicious sign-ins or impossible travel?
17. How quickly can a lost device be disabled or wiped?
18. How do you identify shadow IT or unapproved SaaS applications?
19. Which Zero Trust principles have we already implemented?
20. What are the next three practical improvements you recommend?

8.43 Chapter 8 Cybersecurity Checklist

Identity

☐ MFA is enabled for important systems.
☐ Privileged accounts use stronger protection.
☐ Dormant accounts are removed.
☐ Risky sign-ins are monitored.

Devices

☐ Company laptops are managed.

☐ Disk encryption is enabled.
☐ Security updates are enforced.
☐ Endpoint protection is active.
☐ Lost devices can be remotely managed.

Remote Work

☐ Remote access requires MFA.
☐ Remote-access systems are patched.
☐ Home and public networks are treated as untrusted.
☐ Remote users receive only necessary access.

Cloud Security

☐ Cloud administrators are limited.
☐ External sharing is reviewed.
☐ Security logging is enabled.
☐ Public storage is restricted.
☐ Cloud configurations are periodically reviewed.

External Collaboration

☐ Guest accounts are inventoried.
☐ Guest access has business ownership.
☐ Temporary access expires.
☐ Vendor accounts are reviewed.

Applications and APIs

☐ Third-party app permissions are reviewed.
☐ High-risk OAuth consent is restricted.
☐ API keys are protected.
☐ Service accounts use least privilege.

Zero Trust

☐ Access decisions consider identity.
☐ Device health is considered where practical.
☐ Least privilege is enforced.
☐ Sensitive resources are segmented.
☐ Access is reviewed continuously.

8.44 Key Terms

Bring Your Own Device (BYOD): The use of personally owned devices for business purposes.

Cloud Computing: The delivery of computing services such as applications, servers, storage, databases, and identity services through remotely operated infrastructure.

Conditional Access: An access-control approach that evaluates context such as user identity, device, location, application, and risk before granting access.

Endpoint Detection and Response (EDR): Technology that monitors endpoint activity and helps detect, investigate, and respond to suspicious behavior.

Identity: A digital representation of a user, device, service, or application that can be authenticated and authorized.

Managed Device: A device that an organization can configure, monitor, and enforce security requirements on through centralized management.

Mobile Device Management (MDM): Technology used to manage and enforce security settings on mobile devices and often laptops.

OAuth Consent: Authorization granted by a user or administrator allowing an application to access specified cloud resources or data.

Shared Responsibility Model: The division of security responsibilities between a cloud provider and its customer.

Shadow IT: Technology or cloud services used without formal organizational approval or visibility.

Software as a Service (SaaS): A cloud model in which users access applications operated by a service provider, typically through a browser or application.

Zero Trust: A security approach that avoids granting broad trust based solely on location and instead continuously verifies identity, device, context, and need.

Zero Trust Network Access (ZTNA): A method of providing remote access to specific applications or resources without broadly placing the user onto the internal network.

8.45 Key Takeaways

The modern workplace is distributed. Employees, applications, devices, data, and vendors may all operate outside the traditional office network.

This means cybersecurity can no longer depend primarily on physical location or the assumption that anything inside the network is trustworthy.

Identity must be protected. Devices must be managed. Cloud configurations must be reviewed. Remote access must be limited. External collaboration must be controlled. Security logs must be retained and monitored.

Zero Trust provides a useful way to think about this environment: verify identity, evaluate context, limit access, and continue monitoring.

The objective is not to create friction for employees. The objective is to make secure access possible from wherever legitimate work occurs.

8.46 Chapter Summary

Chapter 7 focused on protecting sensitive information. Chapter 8 has focused on protecting that information and the systems around it when work happens across cloud platforms, mobile devices, home networks, remote-access systems, and third-party applications.

A modern access decision should consider more than a username and password.

It should consider identity, authentication strength, device health, location, application sensitivity, business need, and observed behavior.

The old model was simple: inside the network meant trusted, outside meant untrusted.

The modern model is more careful: every access request should earn the level of trust appropriate to the situation.

A practical modern security approach looks like this: Protect identity -> Require strong authentication -> Manage devices -> Limit privilege -> Secure cloud configurations -> Control sharing -> Restrict remote access -> Monitor behavior -> Review continuously.

Organizations do not need to replace every system at once. They do need to recognize that the workplace has changed.

Security must change with it.

And that leads to the next major question: If a cloud service, remote-work system, device, or other critical technology becomes unavailable despite these protections, how does the organization keep essential work moving and recover safely?

Chapter 9

Backup, Business Continuity, and Disaster Recovery

Authored by Dr. Charles Edwards

Chapter Contents

9.1 When Prevention Is Not Enough

Imagine arriving at work on Monday and discovering that the shared files will not open. The appointment system is unavailable. The telephone still works, but employees cannot see who is scheduled or what customers were promised. Someone says, “We have backups.” That is reassuring, but it does not tell anyone when work can resume.

The organization needs answers to several different questions. Are the backup copies usable? Can staff reach them without the failed systems? Is replacement equipment available? Which service should return first? What should employees tell customers while the technical team investigates? These questions connect backup, business continuity, and disaster recovery.

A backup is a recoverable copy of information. Business continuity is the ability to keep essential activities operating during disruption, sometimes through temporary arrangements. Disaster recovery is the organized restoration of the technology that supports those activities. The three support one another, but none substitutes for the others. A restored database is useful only when people can safely use the application and complete their work. [1]

Consider a small training company. Its backup contains enrollment records. Its continuity arrangement allows staff to record new registrations on controlled temporary forms. Its recovery procedure explains how to rebuild the enrollment application, restore the records, verify access, and enter the registrations collected during the outage. Each part solves a different problem.

Chapter 8 explained how to protect a workplace spread across devices, cloud services, and remote locations. This chapter asks what happens when one of those protections fails, a provider becomes unavailable, or ordinary equipment breaks. Recovery planning applies to accidental deletion, fire, power loss, failed updates, and human mistakes as well as cyberattacks.

The goal is not to promise that disruption will never happen. It is to decide in advance what matters, what can wait, who will act, and what evidence will show that recovery succeeded. Those decisions are easier to make before customers are waiting and employees are improvising.

9.2 Begin with the Work That Must Continue

A business impact analysis, often shortened to BIA, examines what happens when a business activity stops. Begin with activities such as processing orders, teaching classes, paying employees, or scheduling appointments. Then identify the people, information, applications, equipment, and outside services that each activity needs. NIST uses this analysis to establish recovery priorities and resource requirements. [1]

Asking which server is most important may produce a technical answer that misses the business problem. Asking what must happen before employees can serve a customer reveals the dependencies. An order application might need a database, an identity service, network connectivity, a payment provider, and a printer. Restoring only the application does not restore the whole activity.

Impact changes with time. A payroll interruption early in the week may be manageable; the same interruption just before the processing deadline may prevent timely payment. A college can postpone a routine administrative report, but an unavailable learning platform during a final examination creates a more urgent disruption. Priorities should reflect these time-sensitive conditions.

For each essential activity, identify its owner and describe the minimum acceptable service. A business may need to accept urgent orders before it restores historical reporting. A department may need to communicate revised submission instructions before it restores its entire document archive. Minimum service should be specific enough that staff can recognize when it has been achieved.

Also identify how long a temporary arrangement can work. Paper forms may support a few hours of orders but become unmanageable after several days. A backup Internet connection may support messaging but lack the capacity for every employee to join video meetings. A workaround has limits, and those limits belong in the plan.

The result should be a short, understandable account of priorities, dependencies, and consequences. Business owners decide how much disruption is tolerable; technical staff determine what resources and procedures can meet that expectation. If the proposed recovery capability is insufficient, leadership must fund improvements, change the process, or explicitly accept the remaining exposure.

9.3 Set Two Different Recovery Targets

Recovery time objective, or RTO, describes the targeted time for restoring a system or service after disruption. Recovery point objective, or RPO, describes how far back in time the recovered data may need to go. In plain language, RTO asks, “How long can this service be unavailable?” RPO asks, “How much recent work can we afford to lose?” These are planning objectives, not guarantees. [1]

Suppose a training company's registration service fails at noon. The agreed RTO is four hours, so the target is usable service by 4 p.m. The RPO is one hour, so a suitable recovery point should preserve information through at least 11 a.m. Restoring a copy from midnight might meet the time target while losing far more information than the business accepted.

Now assume the latest usable copy is from 11:30 a.m. and service returns at 3:15 p.m. In this example, the outage lasts three hours and fifteen minutes, while the recovery point leaves thirty minutes of recent activity potentially missing. Both objectives are met, but employees still need to reconcile any registrations collected during the interruption. An objective being met does not mean that every consequence has disappeared.

A backup schedule alone does not prove the RPO. An hourly job might fail, finish late, omit a database, or produce a copy that cannot be restored. Measure the age of the latest usable recovery point. Similarly, the time required to copy files is only part of RTO. Diagnosis, authorization, rebuilding, restoration, security checks, and user validation all consume time.

Define the start and finish of the recovery measurement before testing. For this chapter's examples, the clock starts when service becomes unavailable and stops when the authorized business owner accepts usable service. Different organizations may define service objectives differently; an explicit definition prevents misleading comparisons.

Shorter targets normally require more preparation and expense. The sensible choice is a target justified by the activity's needs and demonstrated through testing, with room for delays rather than an optimistic estimate made under ideal conditions.

9.4 Choose Backup Methods You Can Restore

A full backup copies all selected data. An incremental backup generally copies changes since the preceding backup in its sequence. A differential backup generally copies changes since the last full backup. Products may implement these methods differently, so the actual restore procedure matters more than the label on the purchase order.

In a conventional example, a full backup runs on Sunday and incremental backups run each evening. A Thursday recovery may need Sunday's full backup and the required incremental sequence through Wednesday. With differential backups, the equivalent recovery normally needs Sunday's full backup and Wednesday's differential. An incomplete dependency chain can make a recovery point unusable. [1]

The tradeoff is not simply storage space. Small, frequent backups may reduce transfer demands but create more restore dependencies. A larger backup can consume more time and capacity while simplifying recovery. Some systems construct a new full recovery point from existing copies; others use transaction logs or continuous change capture. Require a demonstration of the method the organization will actually use.

Backup scope also matters. Copying office documents does not preserve application settings, access permissions, database consistency, software installers, or the instructions needed to rebuild a service. Record these supporting requirements. Protect recovery keys and emergency access information so they remain available during an outage without placing unprotected secrets in the plan.

An application-consistent backup captures data in a state the application can reliably recover. Copying the files of a busy database without coordinating with the database may not achieve that result. Use a supported procedure and test by opening the restored application and checking meaningful records, not merely by verifying that files exist.

Retention determines how far back recovery can go. Keeping only the newest copy leaves little room to recover from corruption discovered days later. Preserve enough history for the organization's detection and recovery needs, while considering storage capacity and information-retention obligations. Document what is included, what is excluded, how long versions remain, and who reviews failures. A successful backup report is useful evidence of a completed operation; a successful restore provides stronger evidence that the operation served its purpose.

9.5 Keep a Recoverable Copy Outside the Failure

Several copies are helpful only when the same event cannot destroy them all. Three drives in one office may all be lost in a fire. Copies managed through the same compromised administrator account may all be deleted. Physical separation and administrative separation address different kinds of failure.

The familiar 3-2-1 backup guideline calls for three copies of data, on two types of storage media, with one copy offsite. Treat it as a starting point for examining independence, not as a certification of safety. Ask whether a building incident, account compromise, storage failure, or provider outage could affect every copy at once.

Offline backups are disconnected from the systems that could alter them. Immutable backups use controls designed to prevent modification or deletion for a defined retention period. These concepts differ: an immutable copy may remain online, while a disconnected copy may be writable when reconnected. Confirm the actual retention controls, administrative exceptions, and access dependencies rather than assuming a marketing term provides complete protection.

CISA recommends maintaining offline, encrypted backups of critical data and regularly testing their availability and integrity. Encryption helps protect exposed backup media, but recovery also requires access to the correct decryption keys. Losing the only key can make a perfectly intact backup unusable. [2]

Separate backup administration from ordinary employee access. Limit who can change retention, delete recovery points, or disable protection. Protect those privileged identities with strong authentication, and monitor significant changes. Emergency recovery access should be documented, controlled, and tested before the primary identity system fails.

Storage capacity deserves attention too. A backup destination that fills unexpectedly may stop accepting new recovery points. An overly short retention setting may remove the last known-good copy before corruption is discovered. Someone must receive failure alerts, investigate them, and confirm that protection resumed.

For a small business, the practical question is straightforward: if every ordinary work account and every office computer became unavailable today, could an authorized person still obtain a usable copy? If the answer depends on one employee, one password, or one device, that dependency needs attention before the next disruption.

9.6 Cloud Storage and Synchronization Have Limits

Synchronization keeps locations aligned. If a user changes a file, the change appears elsewhere. This is convenient for collaboration, but an unwanted deletion or damaged file can also propagate. Version history and recycle bins may help, yet their retention, scope, and administrative controls must be understood before they are relied upon for recovery.

A replica is another current copy of a system or its data. It can help when hardware fails, but it may also receive corrupted or malicious changes. A snapshot records a point-in-time view; its usefulness depends on whether it survives loss of the underlying storage and whether an attacker can remove it. None of these terms alone establishes an independent, usable backup.

Consider a company that stores all project work in a cloud collaboration service. The provider may keep the platform available while an employee accidentally deletes a shared folder. That is a customer-data recovery problem even though the provider's infrastructure is healthy. The company needs to know who can restore the folder, which versions remain, and whether permissions and metadata return correctly.

The opposite situation is also possible: the data remains intact, but the provider is unavailable. An exported copy may allow staff to read records without restoring the original application's workflows. A spreadsheet export of appointments is helpful, but it does not automatically recreate reminders, access restrictions, or integration with other systems.

For each cloud service, identify what the service can restore and what the customer must arrange. Check the recovery window, supported export formats, deletion behavior, account-recovery process, and expected retrieval time. Product terms change, so confirm these details for the actual service and agreement instead of assuming every cloud platform behaves alike.

Finally, examine the route to the backup. If backup access requires the same unavailable cloud identity service, a separate storage location may still be inaccessible. If a large restore requires downloading over a slow office connection, transfer time may defeat the recovery target. Cloud recovery planning therefore includes identity, connectivity, application reconstruction, and people who know how to use the recovered information.

9.7 Keep Essential Work Moving During an Outage

Business continuity begins while recovery is still underway. Staff need to know what work can proceed, what must pause, and who may authorize temporary arrangements. The plan should identify an incident leader and a backup person, as well as owners for technology recovery, employee communication, and essential business activities. Small organizations may combine roles, but the responsibilities still need names. [1]

Return to the training company. Its registration system is unavailable, so staff use an approved temporary form with a unique reference number and the time of each request. They collect only the information needed to handle the registration. They do not ask customers to send payment-card details to personal email accounts. The workaround preserves service while retaining basic security boundaries.

Temporary records create a later obligation. Assign a person to enter them into the restored system, identify duplicates, confirm totals, and retain or dispose of the temporary copies according to policy. Without reconciliation, continuity efforts can introduce missing transactions or conflicting records that survive long after the outage.

Communication also requires an alternative route. A plan stored only in an unavailable shared drive cannot guide the response. Keep a controlled, accessible copy of essential procedures and current contact information. Establish how staff will verify instructions if company email is unavailable or suspected to be compromised. An alternate channel is useful only when employees know which messages to trust.

Customer updates should state what is known, what service remains available, and when the next update will occur. Avoid promising a restoration time that has not been supported by evidence. A scheduled update can reduce repeated inquiries even when there is no new technical milestone to announce.

Continuity has a human limit. Employees may be dealing with a power outage, transportation problems, or their own family needs. Document substitutes for key roles and make temporary procedures simple enough for another trained person to follow. A plan that depends on a single expert working continuously is fragile, even when the technology is well protected.

9.8 Restore Safely and Preserve Useful Evidence

Recovery after an ordinary equipment failure differs from recovery after a suspected intrusion. In a cyber incident, copying data back too quickly may reintroduce compromised software or reconnect a clean system to an attacker-controlled environment. Coordinate restoration with the incident-response process so urgency does not undermine containment. CISA's recovery guidance emphasizes prioritizing critical services and avoiding reinfection. [3]

Before rebuilding, document the affected systems and important observations. Qualified responders should decide which logs, disk images, memory captures, or other evidence to preserve, considering operational urgency. Record who collected evidence, when it was collected, where it was stored, and who handled it. A business backup is not automatically a forensic acquisition: it may omit deleted data, volatile memory, or other relevant artifacts.

Choose a recovery point using evidence about the incident, not simply the most recent timestamp. An attacker may have entered long before the visible outage. An older copy is not automatically trustworthy either. Review available logs, inspect the restored material, and use appropriate malware checks while recognizing that no single scan proves absence of compromise.

Prepare a controlled recovery environment. Rebuild compromised systems from trusted installation sources where appropriate, correct the entry point, and address exposed accounts, credentials, or tokens. Restore supporting services in a workable order, including secure administration, networking, identity, and application dependencies. The exact order depends on the environment; a documented dependency map prevents guesswork.

Validate both technical operation and business use. An application may start while its permissions are wrong or its records are incomplete. Ask an authorized user to complete a representative transaction. Confirm expected access restrictions, monitoring, and backups before returning the service to ordinary operation. Document who approved that decision and any remaining limitations.

Continue watching the recovered environment for unexpected activity. Recovery is not complete merely because employees can log in again. The organization must establish that essential work can proceed, temporary records have an owner for reconciliation, and the conditions that caused the disruption have been addressed or explicitly managed.

9.9 Test the Plan Against Measurable Results

A tabletop exercise is a guided discussion of a disruption. Participants work through decisions such as who declares the incident, how employees receive instructions, and what happens if the primary administrator cannot be reached. It reveals unclear responsibilities and missing information without interrupting production. It does not demonstrate that a backup can actually restore a working service. [1]

A restoration test provides technical evidence. Begin with a small, isolated recovery, then progress to a representative application and its dependencies. Define success before starting: which records must appear, what operation a user must complete, how permissions should behave, and what time target applies. Record failures as findings to correct, not as reasons to hide the result.

Suppose the training company begins a controlled test at 9 a.m. Preparing the recovery environment takes forty-five minutes, restoring data takes ninety minutes, validating the application takes forty-five minutes, and obtaining business approval takes fifteen minutes. Assuming these steps are sequential, the total is 195 minutes, or three hours and fifteen minutes. A four-hour target leaves forty-five minutes of margin.

That result describes the conditions tested. It may exclude waiting for replacement equipment, rebuilding identity, or dealing with an unavailable provider. State those exclusions clearly. Otherwise, a successful laboratory restore can be mistaken for proof that the entire business can recover within the same period during a real incident.

Test the less convenient cases as well: an older recovery point, an unavailable primary administrator, an inaccessible office, or a restore that needs emergency credentials. These scenarios examine whether the protection survives the failures it was purchased to address. Use authorized, isolated environments and disposable test data when practicing destructive events.

Keep a record of the date, recovery point used, measured durations, validation results, exceptions, and approving owner. Assign each corrective action a responsible person and due date. Revisit the plan after major application changes, provider changes, staffing changes, and incidents. The value of testing comes from improving the next attempt, not simply from producing a report that says an exercise took place.

9.10 A Practical Recovery Exercise

Use this fictional scenario for individual practice or a classroom discussion. A training center loses its registration application at noon. Its RTO is four hours and its RPO is one hour. The latest verified recovery point is 10:30 a.m.; an 11:30 a.m. backup job failed. Staff can accept requests on approved temporary forms. The recovery team expects usable service by 3:15 p.m.

First, explain whether each target is met. The expected outage is three hours and fifteen minutes, which meets the four-hour RTO if validation and business approval finish by 3:15 p.m. The available recovery point is ninety minutes before the outage, which exceeds the one-hour RPO by thirty minutes. The failed job cannot be counted as protection merely because it was scheduled.

Next, describe what happens to the missing registrations. Staff should identify transactions after 10:30 a.m. using trustworthy available records, then reconcile them with registrations collected during the outage. They should check for duplicates and unresolved gaps. Reconstructing records can reduce the practical loss, but it does not change the timestamp of the backup that was available when recovery began.

For a safe hands-on extension, create a disposable folder containing a fictional registration file and a separate backup copy. Record the backup time and a SHA-256 hash of the original file. Change only the working copy, then restore the backup into a new folder. Compare the restored hash with the recorded original and inspect the contents. Matching hashes demonstrate byte-for-byte consistency with that original; they do not prove the information was correct or malware-free when backed up.

Finish by writing a one-page recovery plan for the fictional service. Include the business owner, dependencies, RTO, RPO, backup location, controlled recovery access, temporary process, validation criteria, and method for reconciling records. Identify the most serious untested assumption. A strong answer should explain how the center would know it had recovered, rather than merely naming a product or stating that backups exist.

9.11 Chapter Review and Further Reading

Backups preserve recovery options. Business continuity keeps essential activities moving. Disaster recovery restores the technology those activities need. A useful plan connects all three through clear ownership, realistic objectives, protected recovery copies, tested procedures, and evidence that the service is usable again.

Remember the distinction between the two recovery targets: RTO concerns interruption time; RPO concerns the age of the recoverable data. Faster restoration does not compensate for an unacceptable recovery point. Likewise, a recent backup does not compensate for a process that takes too long to restore service.

Before accepting the statement “We are backed up,” ask the responsible team to describe its last successful restore, the recovery point used, the time required, and the business checks performed. Then ask what would change if the office, primary administrator, or normal identity service were unavailable. The answers reveal whether recovery is a demonstrated capability or an untested expectation.

For review, explain why synchronization alone may not protect against deletion; describe a failure that could affect both production and its backup; and identify the evidence needed before declaring a restored application ready. These questions apply equally to a small business, a college department, or a larger organization.

Further reading

[1] NIST. Special Publication 800-34 Revision 1, Contingency Planning Guide for Federal Information Systems (2010, updated November 2010). Foundational guidance on business impact analysis, recovery objectives, backup methods, exercises, and plan maintenance. Its federal context should be distinguished from voluntary business use.

https://csrc.nist.gov/pubs/sp/800/34/r1/upd1/finalX

[2] CISA and partner agencies. #StopRansomware Guide. Practical guidance on protecting and testing recovery copies.

https://www.cisa.gov/stopransomware/ransomware-guideX

[3] CISA. I've Been Hit By Ransomware! Recovery guidance addressing service priorities and safe restoration.

https://www.cisa.gov/stopransomware/ive-been-hit-ransomwareX

Sources checked September 6, 2026. All training-company examples, times, and exercises in this chapter are fictional and provided for instruction.

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

Chapter 10

Cybersecurity for Government Contractors

Authored by Dr. Charles Edwards

Chapter Contents

10.1 The Business Problem: Protecting a Customer's Mission

Winning government work can be an important step for a growing business. It can also introduce responsibilities that are easy to overlook when attention is focused on price, delivery dates, and staffing. A company may receive engineering drawings, nonpublic project information, or other material that needs specific protection. The computers holding that information may belong to the company, but the consequences of exposure can extend to the government and its mission.

Consider a small manufacturer preparing to supply parts through a larger defense contractor. The owner believes cybersecurity is covered because the company has antivirus software, a firewall, and an IT service provider. Then the customer asks which systems will handle protected information, what security requirements have been implemented, and what evidence supports the company's assessment. A list of purchased products does not answer those questions.

Government-contracting cybersecurity connects information protection with contractual promises. The organization must understand the information it will receive or create, identify the requirements that apply, and demonstrate that its working environment supports those requirements. Sales, contracts, leadership, and IT all have a role. Giving the entire responsibility to the person who repairs laptops leaves important decisions unowned.

The requirements are not identical for every government customer. Civilian-agency contracts, defense contracts, and work involving classified information can follow different rules. This chapter introduces the common concepts and gives particular attention to defense contracting because the Cybersecurity Maturity Model Certification program, or CMMC, creates frequent questions for smaller businesses.

Chapter 9 explained how to preserve operations and recover information. This chapter adds another question: what protections and reporting commitments has the organization agreed to provide for government-related information? Recovery remains important, but it is one part of a larger responsibility.

The guidance here is current as checked on September 6, 2026. It explains the decisions business leaders need to make; it does not determine the obligations of an unseen contract. For an actual award, review the solicitation, incorporated clauses, amendments, and agency instructions with the people responsible for contracts and security.

10.2 Understand the Information Before Buying a Solution

Federal contract information, usually called FCI, is certain nonpublic information provided by or generated for the government under a contract to deliver a product or service. The definition excludes information the government makes public and simple transactional information needed to process payments. A publicly posted solicitation is therefore different from nonpublic information received during performance. FCI is a defined category, not a label for every document a contractor owns. [1]

Controlled unclassified information, or CUI, is government-related information that requires or permits safeguarding or dissemination controls under law, regulation, or government-wide policy. It is unclassified, but that does not make it public. The National Archives maintains the CUI Registry, which identifies categories and their authorities. Contractors should consult the responsible agency's implementing guidance as well as the Registry. [2]

The practical distinction is that information may need protection even when it is not classified. Engineering information can be CUI when the relevant authority and circumstances apply; an ordinary drawing is not automatically CUI simply because a company considers it valuable. Likewise, company-confidential information does not become CUI merely because an employee stamps that label on it.

Ask the government customer, or the prime contractor through the proper contractual channel, what information is expected, what category applies, and how it must be handled. Record the answer. If markings are missing or instructions conflict, seek clarification rather than assuming that the absence of a banner authorizes unrestricted sharing.

Look beyond the original attachment. Employees may copy information into quotations, working drawings, meeting notes, email, printouts, support tickets, or backup systems. A control that protects the original folder may fail when staff move its contents somewhere else. Protection must follow the actual work.

Classified information belongs to a separate security regime. A CMMC assessment or a CUI-capable environment does not authorize a business to handle classified material. Business leaders do not need to memorize every information category, but they do need a reliable method for identifying the information entrusted to them and keeping it within appropriate boundaries.

10.3 Read the Contract as a Set of Responsibilities

The Federal Acquisition Regulation, or FAR, provides government acquisition rules. The Defense Federal Acquisition Regulation Supplement, or DFARS, adds defense-specific requirements. For a business leader, the essential point is that a clause in an applicable contract can create duties that a general security checklist does not capture.

The basic safeguarding clause, commonly encountered as FAR 52.204-21, addresses systems handling FCI. Its fifteen requirements include limiting access, authenticating users, controlling physical access, correcting flaws, and protecting against malicious code. Other requirements may still apply. A company should not interpret this baseline as a complete cybersecurity program for every kind of government information. [1]

DFARS 252.204-7012 addresses safeguarding covered defense information and reporting specified cyber incidents. Covered defense information is a defined term that includes qualifying controlled technical information and other qualifying CUI in the contractual context. The clause also addresses external cloud services, evidence preservation, and subcontract responsibilities. Its reporting duties are discussed in Section 10.8. [3]

CMMC addresses assessment of the required protections. When applicable, the contract identifies the required status and related obligations. DFARS 252.204-7021 connects that status with systems used to process, store, or transmit covered information and with continuing affirmations. A company should distinguish implementing security, assessing security, and recording the required status: completing one does not automatically complete the others. [4]

Clause numbering is also changing during the acquisition-rule overhaul. For example, the official FAR overhaul guide places basic safeguarding at 52.240-93. Older numbers remain visible in existing contracts and published references. Read the clause title, text, date, applicable deviation, and amendments together; a number alone is not a reliable account of the obligation. [7]

Create a short responsibility register for each relevant contract. Record the information involved, applicable clause or instruction, required assessment, reporting route, responsible employee, and review date. Where a requirement is unclear, send a specific question through the contracting channel. This gives leadership an actionable record instead of a collection of acronyms and assumptions.

10.4 CMMC: What the Current Program Requires

CMMC stands for Cybersecurity Maturity Model Certification. It provides a framework for assessing whether contractors have implemented required protections. It does not replace daily security work, and its name does not mean that every contractor currently needs a third-party certification. The applicable requirement depends on the information, procurement, and current implementation instructions.

As of September 6, 2026, official guidance states that the July 13 suspension of Phase II leaves Phase I self-assessment requirements in place. The planned November 2026 transition is suspended. During the suspension, the implementation memorandum limits procurement designations to Level 1 (Self) or Level 2 (Self), while preserving underlying safeguarding duties. It also directs amendments and modifications for affected solicitations and contracts. Obtain the actual amendment or modification rather than treating a headline as a rewritten contract. [5]

Level 1 concerns basic protection of FCI. Current program guidance calls for an annual self-assessment and annual affirmation against the fifteen FAR safeguarding requirements. Plans of action and milestones cannot be used to obtain Level 1 status while leaving requirements unfinished. [6]

Level 2 addresses broader protection of CUI. Current guidance describes a self-assessment every three years against the 110 requirements of NIST Special Publication 800-171 Revision 2, with annual affirmation. Limited conditional status is possible under defined rules, with eligible deficiencies closed within 180 days. That is not permission to place any missing safeguard on a future-work list. [6]

Older materials describe Level 2 third-party assessments and government-led Level 3 certification. Those remain part of the regulatory framework, but the suspension changes current implementation. Government-led security assessments can still occur. Do not confuse a pause in a certification rollout with cancellation of information-protection responsibilities. [5]

NIST has published Revision 3 of SP 800-171, while current CMMC guidance continues to reference Revision 2. A newer publication does not by itself change the baseline incorporated into a particular assessment or contract. Identify the required revision explicitly and watch for official changes before using a different checklist. [6, 8]

10.5 Define Where Protected Information Can Go

Scope means the people, systems, services, and locations included in the protection and assessment effort. It should reflect reality. If protected files travel through ordinary email, personal laptops, a shared printer, and a cloud backup account, drawing a box around one server does not make those other paths disappear.

Start by following one representative item from receipt to disposal. Who receives it? Where is it downloaded? Which application opens it? Who can print, forward, synchronize, or copy it? What happens when someone requests technical support? Where do backup copies reside? This exercise often reveals more than a network diagram that shows only major equipment.

A smaller, deliberately separated environment is sometimes called an enclave. It can reduce the number of systems handling CUI when both technical controls and working procedures enforce the separation. For example, a manufacturer might limit protected design work to managed workstations and an approved collaboration service, while ordinary marketing work remains elsewhere.

Separation must be usable. If employees routinely email protected files to themselves because the approved application is inconvenient, the intended boundary has failed. Consider clipboard use, downloads, printing, removable media, remote sessions, and external sharing when designing workflows. Explain the permitted path clearly and test that ordinary tasks can be completed within it.

Supporting systems also matter. Identity services, administrative tools, monitoring platforms, and backup services may provide security functions for the protected environment. NIST SP 800-171 Revision 2 includes components handling CUI and components providing protection for them. Assessment scoping requires attention to those dependencies, not just to where the main documents are saved. [8]

Document the boundary and keep it current. A new application, service provider, office, or subcontractor can change the information flow. Review such changes before introducing protected information. The objective is not to make the assessment boundary artificially small; it is to create an accurate, manageable environment whose protections can be explained and demonstrated.

10.6 Make Security Claims Match Daily Practice

A written policy explains what the organization expects. A procedure explains how people carry it out. Technical settings enforce parts of the procedure. Evidence shows what actually happened. These pieces should agree. A policy stating that access is removed immediately has little value if a former employee's account remains active for weeks.

A system security plan, or SSP, describes the system boundary, operating environment, and how security requirements are implemented. It should explain the actual organization, including responsibilities shared with providers. A generic document that names products the company does not use creates confusion when employees are asked to demonstrate their work. [8]

Consider access control. The SSP might explain that the manager approves access to a protected project, IT assigns the approved role, and the project owner periodically reviews membership. Supporting evidence might include an approval record, current permissions, and a completed review. Together they explain who should have access, how that decision is applied, and whether someone checks it.

Use the same discipline for updates, security alerts, incident response, and staff training. An installed monitoring product is not proof that anyone investigates alerts. A training subscription is not proof that employees completed the relevant instruction. Select evidence that demonstrates the requirement rather than accumulating screenshots without context.

A plan of action and milestones, often shortened to POA&M, records deficiencies, corrective work, responsible people, and deadlines. It is useful for managing improvement, but assessment rules determine whether a particular deficiency can support conditional status. A planned purchase is not an implemented safeguard. Keep current capability separate from future commitments. [6]

The Supplier Performance Risk System, or SPRS, records relevant assessment information and affirmations. A score in that system is not a product certification or proof that every future activity is covered. Before submitting or affirming a result, leadership should understand the scope, assessment method, evidence, and unresolved issues. Maintain a review calendar so the organization continues operating as described after the assessment is complete. [4]

10.7 Providers and Subcontractors Share the Work

An IT provider can configure systems, monitor alerts, maintain backups, and help prepare evidence. It cannot make the contractor's responsibilities disappear. The business must still know what information it handles, what it promised its customer, and which tasks the provider has actually agreed to perform.

Begin with a responsibility discussion. If a suspicious sign-in occurs, who investigates it? Who contacts the business? Who can export logs? Who decides whether the event meets a reporting threshold? Who submits the report? Write down these assignments and the escalation route. A service agreement that says only “security monitoring included” may leave important work undefined.

For the situation described in paragraph (b)(2)(ii)(D) of DFARS 252.204-7012, an external cloud service handling covered defense information must meet security requirements equivalent to the FedRAMP Moderate baseline and satisfy specified incident-reporting and investigative-support obligations. A general claim that a service is secure does not establish those conditions. [3]

Review the exact service offering and its documented boundary. A provider's authorization or assurance may cover one service rather than every product it sells. Ask which protections the provider supplies, which settings the customer must configure, and what evidence will be available. Also determine how support staff, integrations, and recovery copies interact with the protected information.

Subcontractors need the right requirements before they receive the information. The FAR safeguarding clause and applicable DFARS clauses include flowdown duties, with scope and exceptions that depend on the clause. Do not send protected material to a supplier and assume that an ordinary nondisclosure agreement supplies every required safeguard. [1, 3, 4]

Finally, plan the end of the relationship. Establish how information is returned or retained, how access is removed, and how the contractor can obtain records needed for continued operation or an investigation. A provider relationship should make responsibility clearer. If neither party can explain who owns an important task, resolve that gap before depending on the service for government work.

10.8 Prepare for Reporting Before an Incident

Chapter 9 emphasized restoring usable service. Government contracting adds reporting and evidence-preservation duties that may continue even after operations resume. A recovered system does not erase the need to review a qualifying incident or notify the required parties.

When applicable, DFARS 252.204-7012 requires rapid reporting of cyber incidents affecting a covered contractor information system or its covered defense information, or affecting the contractor's ability to perform designated operationally critical support. The clause defines rapid reporting as within 72 hours of discovery. This is not a blanket statement that every technical alert at every government contractor must be reported under that clause; the defined trigger and contractual context matter. [3]

Do not build a process that waits for a completed forensic investigation before evaluating the reporting obligation. Investigations can take much longer than the reporting window. The response team needs a prompt route to security, contracts, and legal personnel who can assess the facts, document decisions, and arrange the required submission.

The same clause requires preservation and protection of images of known affected systems and relevant monitoring or packet-capture data for at least 90 days from submission of the incident report. The period runs from reporting; it is not simply a requirement to have ninety days of historical logs. The clause also addresses access for forensic analysis and handling of discovered malicious software. [3]

Prepare access to the designated reporting mechanism in advance, including the credentials or certificates currently required. Confirm who can report if the primary contact is unavailable. Subcontractors should also understand the requirement to pass the government-assigned incident report number to the prime contractor or next higher tier as soon as practicable. Other contractual notifications may apply as well. [3]

During recovery, preserve relevant evidence before wiping or rebuilding systems when feasible under the response plan. Record the timeline and decisions. Coordinate with qualified responders so containment, business recovery, and evidence collection support one another. Test this process through a discussion exercise: a reporting plan that depends on an unavailable mailbox or one employee's memory is not ready for an actual incident.

10.9 A Realistic Scenario: The Drawing That Escaped

The following example is fictional. Riverside Precision, a twenty-person manufacturer, receives a protected technical drawing from its prime contractor. The drawing's CUI status and handling requirements have been identified for the work. Riverside has an approved location for project files, but employees have not practiced the complete workflow.

An estimator downloads the drawing to a personal laptop to finish a quotation at home. The laptop synchronizes its desktop to a personal cloud account. When the estimator needs help, a screenshot containing part of the drawing is pasted into an ordinary support ticket. Later, a suspicious sign-in appears on the personal cloud account.

What went wrong? The company treated protection as a property of the original shared folder. It did not account for the employee's download, synchronization, or support process. Its intended boundary was smaller than its real information flow. The IT provider maintained the shared folder but had not agreed to manage the personal account or monitor the laptop.

The suspicious sign-in does not, by itself, establish every fact about a disclosure. It does require prompt investigation and evaluation of applicable reporting duties. Riverside should preserve available evidence, identify what information was present, contain unauthorized access, and involve the people responsible for its incident and contractual response. Quietly deleting the cloud copy would not resolve the exposure or preserve the facts needed to understand it.

What should have happened? Before accepting the drawing, Riverside should have established a workable approved path for estimating, remote work, and support. Employees should have known which devices and services were permitted. Technical restrictions and staff instruction should have reinforced that path, while the provider agreement explained responsibility for monitoring and escalation.

The lesson for leadership is practical: ask an employee to walk through an ordinary task from beginning to end. Watch where information moves when time is short or something goes wrong. This is not about blaming an employee for finding a convenient shortcut. It is about designing a process that people can follow while keeping the organization's security promises true.

10.10 Turn Understanding into a Practical Plan

Questions to ask your IT team or provider

Ask who confirmed the information categories, where protected information travels, which systems and providers are in scope, and which contract establishes the requirements. Request a current diagram and responsibility record.

Request evidence of approved access, completed access reviews, and removal of departed users. Ask who checks assessment results before leadership affirms them, who tracks remediation, and who would evaluate and report a qualifying weekend incident.

TRG Practical Takeaway

Begin with the contract and the information. Establish an accurate boundary, assign responsibility, and compare actual practices with applicable requirements. Use the findings to prioritize changes and make truthful assessment claims. A smaller business can make progress systematically without assuming that a product purchase or a purchased policy package completes the work.

Before You Move On

- Confirm information categories and contract requirements.
- Document the protected environment and its dependencies.
- Assign provider, subcontractor, and internal responsibilities.
- Verify assessment evidence and track required affirmations.
- Test incident escalation, reporting access, and recovery.

Key terms in plain English

FCI is qualifying nonpublic federal contract information. CUI is unclassified government-related information subject to authorized safeguarding or dissemination controls. An SSP explains the system and its protections. A POA&M tracks corrective work. SPRS records assessment information and affirmations. CMMC is the framework used to assess required protections; its applicable assessment path must be checked against current instructions.

Chapter summary

Government-contracting cybersecurity joins ordinary security practice with specific information-handling and contractual responsibilities. Know what you protect, demonstrate how you protect it, and keep those protections working. The next chapter

examines artificial intelligence, including the risk of submitting protected business or government information to a tool that has not been approved for that use.

10.11 Sources and Guidance for Further Reading

Current-status note

Sources were checked on September 6, 2026. The CMMC discussion reflects the July 2026 implementation suspension, not the earlier planned Phase II schedule. Review current official instructions and your actual contract before making an assessment, procurement, or reporting decision. This chapter concerns unclassified contractor information; it does not provide authorization for classified work.

Further reading

[1] Acquisition.gov. FAR 52.204-21, Basic Safeguarding of Covered Contractor Information Systems. Definitions, fifteen safeguards, and subcontract provisions.

Read official sourceX

[2] National Archives. About Controlled Unclassified Information and the CUI Registry. Definitions, categories, and agency guidance.

Read official sourceX

[3] Acquisition.gov. DFARS 252.204-7012, Safeguarding Covered Defense Information and Cyber Incident Reporting. See paragraphs (b) through (g) and (m).

Read official sourceX

[4] Acquisition.gov. DFARS 252.204-7021, Contractor Compliance With the Cybersecurity Maturity Model Certification Level Requirements. Read with current implementation guidance and modifications.

Read official sourceX

[5] Defense acquisition implementation memorandum. Implementing Suspension of CMMC Phase II (July 2026). Includes procedures for affected solicitations and contracts.

Read official sourceX

[6] Office of the Chief Information Officer. About CMMC. Current implementation, assessment, affirmation, and remediation summary.

Read official sourceX

[7] Acquisition.gov. FAR Overhaul, Part 40. See Section 40.303 for basic safeguarding in the overhaul guidance.

Read official sourceX

[8] NIST. SP 800-171, Revisions 2 and 3. Read the revision required by the applicable assessment and contract.

Read official sourceX

Chapter 11

Artificial Intelligence, Cybersecurity, and Business Risk

Authored by Dr. Charles Edwards

Artificial intelligence has moved from an emerging technology to an everyday business tool. Employees use AI to draft emails, summarize documents, write reports, analyze spreadsheets, create software, prepare presentations, research competitors, answer customer questions, and improve productivity.

For many organizations, the question is no longer whether employees will use AI. They already are. The more useful questions are whether leadership knows how AI is being used, what information is being submitted, what systems the AI can reach, and what decisions people are making based on AI-generated output.

AI can create tremendous business value. It can also create new cybersecurity risk when adoption moves faster than governance, data handling, access control, and human verification.

I do not believe the answer is to ban every new technology. The better approach is to understand what the technology does, what information it receives, what access it has, what could go wrong, and what reasonable safeguards should be in place.

Chapter Contents

11.1 The Business Problem: AI Adoption Is Moving Faster Than Governance

Consider a 60-person engineering company. The company has not formally purchased a generative AI platform, so leadership assumes artificial intelligence is not yet a significant cybersecurity issue. But employees have already begun using it. A project manager copies customer material into a public chatbot for a summary. A salesperson uploads a draft proposal for rewriting. A developer pastes internal source code into an AI coding assistant. An administrator uses AI to rewrite a spreadsheet containing employee information.

None of these employees intends to create a security problem. They are trying to work more efficiently. The difficulty is that the organization may no longer know where its information is going.

For years, organizations have tried to control where sensitive information is stored. Generative AI introduces another destination for organizational information. An employee may copy information into an AI system in seconds without thinking of the action as a data transfer. From the employee's perspective, the person is simply asking a question. From a cybersecurity perspective, information has been submitted to an external service.

The leadership question is therefore not simply whether employees are using AI. Leadership must know which systems are being used, what information is being submitted, what happens to that information, who can access it, whether it is retained, whether it may be used for model improvement, and whether the service meets contractual or regulatory requirements.

11.2 Treat AI as Another System That Handles Information

Artificial intelligence often feels different from traditional software because the interaction resembles a conversation. That simplicity can hide what is happening underneath. A user may be sending information to a cloud service operated by another company, and that service may connect to external data sources, plugins, APIs, search systems, or other services.

The safest management approach is to stop treating AI as magic and treat it as an information-processing system. If an employee would not upload a confidential customer file to an unknown public website, the employee should not automatically assume it is acceptable to paste the same information into an unapproved AI chatbot.

This principle becomes especially important for customer records, employee information, financial data, passwords or credentials, source code, legal documents,

medical information, contract data, intellectual property, security configurations, incident information, and government-related information.

11.3 Public AI Tools and Enterprise AI Services Are Not Necessarily the Same

One of the most important distinctions for leadership is the difference between a publicly available consumer AI service and an AI service purchased and managed for organizational use. Employees often see the same product name and assume the security characteristics are identical. They may not be.

An enterprise service may provide centralized identity management, administrative controls, contractual privacy commitments, data-retention options, audit capabilities, access restrictions, and other safeguards that are not available in the same way through an individual consumer account.

This does not mean every enterprise AI service is automatically appropriate for every type of information. Organizations must evaluate the actual service, account type, configuration, contract, and use case.

11.4 The Most Immediate Risk Is Often Sensitive Information Disclosure

When people hear “AI cybersecurity,” they may imagine an attacker compromising a sophisticated machine-learning model. Those risks exist. But the most immediate problem for many ordinary businesses is much simpler: an employee gives sensitive information to the wrong AI system.

Imagine an estimator working for a government contractor. The estimator receives a technical drawing from a customer and wants to prepare a quick summary. The employee copies part of the drawing into a public generative AI tool. The AI produces an excellent summary and saves time. The unresolved question is whether the employee was authorized to provide that information to that service.

Artificial intelligence does not create an exception to existing promises about protecting information. It creates another place where those promises can be broken.

11.5 AI Can Expose Information Without an Employee Uploading a File

Sensitive-information exposure is not limited to manually pasting text into a chatbot. AI assistants increasingly connect to email, calendars, cloud storage, internal documentation, customer systems, source-code repositories, and databases.

This capability can make AI dramatically more useful, but it also makes permissions more important. If an AI system can access information on behalf of a user, the organization must ask whether existing permissions are already too broad.

AI can amplify weaknesses that already exist. Poor permissions, weak data classification, uncontrolled application access, weak vendor management, and inadequate logging all become more consequential when an AI assistant can search and act across information at high speed.

11.6 AI Answers Can Be Convincing and Wrong

Generative AI systems can produce polished, professional, and convincing responses. Those responses can still be wrong. An incorrect answer may contain invented facts, nonexistent citations, inaccurate calculations, faulty technical recommendations, or insecure code.

The danger is not simply that AI can make mistakes. The danger is that AI may present a mistake with confidence. Human verification is therefore essential when the consequence of error is significant.

If AI drafts a routine internal message, an unnoticed error may have little consequence. If AI recommends a firewall configuration, interprets a contract requirement, generates production software, evaluates a security incident, or recommends a financial transaction, the consequence may be much greater. The more important the decision, the stronger the verification should be.

11.7 Artificial Intelligence Does Not Remove Accountability

A dangerous organizational mindset is beginning to appear: “The AI told me to do it.” That statement does not transfer responsibility.

If a manager uses AI to prepare a decision, the manager remains responsible for the decision. If a developer uses AI to generate software, the organization remains responsible for what it deploys. If a security analyst uses AI to recommend containment actions, the security team remains responsible for verifying that those actions are appropriate.

AI should be understood as a tool that can support human work. It should not become an invisible decision-maker whose recommendations are accepted simply because they were generated by an advanced system.

11.8 Prompt Injection: When Instructions Become an Attack Surface

Traditional cybersecurity professionals are familiar with malicious input. Generative AI introduces a related problem known as prompt injection. A prompt injection attack attempts to manipulate the instructions an AI system follows.

The attack may be direct, such as telling an AI system to ignore previous instructions, or indirect, such as placing instructions inside a document, webpage, or email that an AI assistant later retrieves and processes.

The risk becomes more serious when the AI system has tools or permissions. An AI system that can only generate text has one level of risk. An AI system that can read email, modify records, call APIs, create accounts, execute software, or initiate business processes has another.

The more authority an AI system receives, the more carefully its inputs, permissions, and actions must be controlled.

11.9 AI Agents Raise the Stakes

AI agents move beyond answering questions and can perform actions. An agent might read support requests, search internal systems, create tickets, draft responses, update customer records, or assist security operations by investigating alerts.

These capabilities can create enormous efficiency, but they also change the risk model. When an AI system can take action, the principle of least privilege applies in the same way it does to human administrators and service accounts.

If an agent only needs to read a record, it should not automatically receive permission to delete it. If it only needs to draft an email, it may not need authority to send without approval. High-impact actions should often require confirmation, separation of duties, or human approval.

11.10 The AI Supply Chain Matters

Organizations rarely build an entire artificial intelligence system themselves. They rely on models, cloud providers, software libraries, datasets, plugins, APIs, external tools, and other vendors. That creates an AI supply chain.

A business may trust the application it purchased without realizing that the application connects to additional AI providers. A developer may download an AI model without understanding its origin. A vendor may change the model underlying its service. A plugin may gain access to business systems.

The management discipline is familiar: know what you depend upon, what access the provider receives, what information leaves your environment, how updates are handled, how incidents are reported, and how the relationship can be terminated.

11.11 Shadow AI Is the New Shadow IT

Shadow AI is AI technology used within an organization without appropriate visibility, approval, or governance. An employee finds a free AI application online, creates a personal account, and begins using it for company work. No one in IT knows the application exists, no security review has occurred, and no one knows what information is being submitted.

Attempting to solve the problem only through prohibition may fail. Employees often turn to unauthorized tools because those tools solve a real business problem. The better response is to understand why employees want the capability and provide an approved alternative where practical.

Security becomes stronger when the secure path is also the usable path.

11.12 AI Is Making Social Engineering More Convincing

Social engineering has always depended on convincing a person to trust something that should not be trusted. Artificial intelligence makes that easier.

Employees were once taught to look for poor grammar, awkward wording, and generic greetings as phishing clues. Those clues are becoming less dependable. Generative AI can create polished, industry-specific, context-aware messages in seconds.

The protection must therefore come from business process rather than from the employee's ability to identify unusual grammar. High-risk transactions should require independent verification, and urgency should never eliminate normal authorization procedures.

11.13 Voice Cloning Changes the Meaning of “I Heard His Voice”

Artificial intelligence can generate convincing synthetic audio. Public recordings from presentations, webinars, interviews, podcasts, and social media can provide material for imitation.

A familiar voice is therefore not the same thing as verified identity. Finance employees receiving unusual payment instructions should verify requests using a known number or established approval system. Employees should be supported, not punished, for slowing down suspicious transactions.

11.14 Deepfake Video Raises the Same Problem

Synthetic video extends the same challenge. An attacker may create or manipulate video that appears to show another person speaking or participating in a meeting.

Video conferences often contain weak lighting, compression artifacts, delays, and small images, all of which can make manipulation harder to recognize. The correct response is not to distrust every video call, but to make important business actions depend on more than one signal.

Identity verification and transaction approval are different controls. They should remain different.

11.15 AI Does Not Invent Social Engineering — It Scales It

Artificial intelligence did not invent phishing, impersonation, business email compromise, fraudulent invoices, or telephone scams. It can reduce the effort required to make those attacks more convincing and more personalized.

The old defenses remain important: question unusual requests, verify payment changes, avoid entering credentials through suspicious links, report suspicious activity, and use multifactor authentication. AI raises the quality and scale of the attack; it does not make foundational defenses obsolete.

11.16 AI Can Also Increase the Speed of Cyberattacks

Generative AI can help explain technical concepts, generate scripts, modify code, summarize documentation, analyze errors, translate text, and automate repetitive tasks. Those same capabilities that help legitimate defenders can help malicious actors work more efficiently.

This does not mean AI automatically turns an inexperienced person into an elite attacker. Cybersecurity still requires knowledge, access, infrastructure, and judgment. But AI can reduce the time required for some tasks.

If attackers become faster, defenders must improve detection and response speed. Logging, alerting, endpoint detection, identity monitoring, and incident-response readiness become even more important.

11.17 Artificial Intelligence Can Help Defenders Too

AI can provide significant cybersecurity benefits. Security teams often face too much information rather than too little. Authentication logs, firewall events, endpoint alerts,

email-security events, cloud logs, vulnerability findings, and application events can overwhelm analysts.

AI can help summarize alerts, identify patterns, explain technical findings, correlate information, draft incident timelines, analyze suspicious code, and search documentation.

But security logs and incident data may themselves be sensitive. Analysts should not assume that incident details can be submitted to an unapproved public AI system merely because the purpose is cybersecurity.

11.18 AI-Generated Code Must Still Be Reviewed

AI can help developers generate functions, tests, documentation, and design ideas. It can also produce code that appears correct while containing subtle security weaknesses.

An AI-generated program may use an outdated library, fail to validate input, mishandle authentication, expose secrets, create injection vulnerabilities, use weak cryptography, or recommend an untrustworthy dependency.

AI-generated code should therefore be treated as untrusted code until it has been reviewed and tested appropriately. Code review, dependency management, secret scanning, testing, and secure development practices remain necessary.

11.19 Never Paste Secrets Into an AI Prompt

Passwords, private keys, API keys, authentication tokens, database credentials, cloud access keys, recovery codes, and other secrets should not be casually pasted into an AI system.

Developers troubleshooting a problem may be tempted to paste entire configuration files into a chatbot. If the file contains a production credential, the request for help has created an exposure.

If a secret is accidentally exposed, revocation and replacement may be required rather than simply deleting the conversation and assuming the problem has disappeared.

11.20 AI Hallucinations Become Dangerous When They Turn Into Actions

AI can produce information that sounds correct but is unsupported or wrong. The business consequence depends on what happens next.

If an AI system incorrectly recommends a security setting and an administrator applies it to production, the output becomes a configuration change. If it invents a software package and a developer installs a malicious package with that name, the output becomes a supply-chain event. If it misreads a contract requirement and management acts on it, the output becomes a governance failure.

The impact of an AI mistake depends on what happens after the answer is produced. The more consequential the action, the stronger the validation should be.

11.21 Human-in-the-Loop Is More Than a Buzzword

Organizations frequently say that they will keep a human in the loop. That matters only if the human is actually able to evaluate what the AI is doing.

If one employee is expected to approve thousands of recommendations every day, the process may degrade into automatic clicking. A meaningful reviewer must understand the decision, have enough information and time, possess authority to reject it, and be supported when slowing down an unsafe action.

Human involvement should not become a decorative control.

11.22 AI Permissions Should Follow Least Privilege

Least privilege applies directly to artificial intelligence. A scheduling assistant may need calendar access but not payroll records. A customer-service agent may need to read order status but not issue unlimited refunds. A security AI may need telemetry without receiving unrestricted administrator authority.

Excessive access increases the blast radius. If the AI makes a mistake, is manipulated, or is compromised, unnecessary permission determines how much damage can occur.

This is particularly important for agentic AI because the question changes from “What can the AI say?” to “What can the AI do?”

11.23 Connectors Can Quietly Expand AI Access

AI tools become more useful when they connect to Microsoft 365, Google Workspace, SharePoint, cloud drives, CRM platforms, ticketing systems, source-code repositories, databases, and internal knowledge bases.

Every connector creates a trust relationship. Leadership should know what the connector can access, whether it can write or delete data, whether activity is logged, who approved it, whether permissions can be restricted, and how the connection can be disabled during an incident.

AI governance cannot focus only on the model. Identities, connectors, APIs, data sources, and permissions may matter just as much.

11.24 Retrieval-Augmented Generation Does Not Automatically Make AI Safe

Retrieval-augmented generation, often shortened to RAG, allows an AI system to retrieve information from organizational documents or databases and use it as context when generating an answer.

This can make AI dramatically more useful, but it also means the AI now has access to organizational information. The system must enforce permissions correctly, protect indexes and embeddings, and consider whether retrieved documents contain malicious instructions.

RAG can improve accuracy and usefulness. It does not eliminate the need for access control, data classification, logging, validation, and security testing.

11.25 AI Governance Should Start With an Inventory

Organizations cannot govern AI they do not know they are using. The first practical step is visibility.

Leadership should identify which AI tools employees use, which AI functions are embedded in existing applications, which departments use them, what information is submitted, what systems they can access, who administers them, and what business purpose each serves.

AI governance begins by discovering reality, not by writing a policy that assumes reality is different.

11.26 Classify AI Use by Risk

Not every AI use deserves the same level of control. Asking an approved AI system for ideas for a company picnic title is different from asking an AI system to determine which employees should be terminated.

Organizations should consider the sensitivity of information, importance of the decision, impact on people, ability of the AI to perform actions, contractual or regulatory requirements, external communications, and how easily an error can be reversed.

A risk-based approach allows low-risk use to remain simple while higher-risk use receives stronger review.

11.27 A Practical AI Acceptable-Use Policy

An AI acceptable-use policy should be understandable enough that employees can actually use it. Employees need clear answers about approved tools, prohibited information, customer and employee data, source code, government information, personal AI accounts, external communications, code review, automated decisions, and incident reporting.

A practical policy may distinguish public information, ordinary internal information, confidential or regulated information, and secrets such as passwords or authentication tokens. The exact categories will vary, but the rules should be clear enough for routine decisions.

11.28 Do Not Write an AI Policy That Forces Employees to Hide

A policy can be too strict to succeed. If an organization prohibits all AI while employees know it can save hours of work, some employees may continue quietly. Visible AI use then becomes shadow AI.

There are legitimate reasons to prohibit specific tools or uses, but governance works better when leadership distinguishes responsible use from unacceptable use and provides an approved path where practical.

The goal is controlled adoption, not uncontrolled adoption and not reflexive prohibition.

11.29 AI Vendor Assessment Requires Familiar Questions

The questions organizations should ask AI vendors resemble ordinary vendor-risk questions: what information the provider receives, where it is processed, how long it is retained, whether it is used for training or service improvement, what administrative controls exist, whether single sign-on and MFA are supported, what logs are available, how subprocessors are used, how incidents are reported, and what happens when the contract ends.

When AI systems connect to business applications, additional questions arise: what permissions the integration requires, whether it can write data, invoke tools, or perform actions, whether administrators can restrict connectors, and whether activity can be reconstructed after an incident.

The technology is new. Many of the management questions are familiar.

11.30 NIST Provides a Useful Governance Model

The National Institute of Standards and Technology published the AI Risk Management Framework and a Generative Artificial Intelligence Profile to help organizations identify and manage AI risk. The framework organizes work around four broad functions: Govern, Map, Measure, and Manage.

For a small or mid-sized organization, these ideas can be translated into ordinary business questions: Who owns this AI use? What are we using it for? What information does it touch? What could go wrong? How will we know? What safeguards exist? Who reviews the system over time?

That is AI governance without unnecessary jargon.

11.31 Governance Is Not the Same as Creating an AI Committee

A committee may be useful, but meetings alone do not create governance. Governance means responsibility is assigned, decisions are documented, high-risk uses are reviewed, employees know what is permitted, approved systems are identified, incidents can be reported, vendors are evaluated, and permissions are controlled.

A small company does not need the same governance structure as a multinational corporation. The principles still apply: someone must be responsible, know what tools are being used, decide what information can be submitted, monitor changes, and respond when something goes wrong.

11.32 AI Risk Changes Over Time

An AI system that was acceptable when purchased may create a different risk six months later. The provider may add new features, Internet access, connectors, automation, new retention terms, or a different model provider.

AI assessment therefore cannot be a one-time procurement exercise. Organizations should periodically ask whether the technology, their use, the information, permissions, integrations, or known risks have changed.

AI governance should follow the system throughout its life cycle.

11.33 A Realistic Scenario: The Helpful Employee

Northstar Consulting is a seventy-person professional-services company supporting commercial and government customers. Leadership has discussed AI but has not adopted an official platform or policy. Employees, however, are already using AI.

An account manager named Melissa is preparing a customer presentation. The customer has provided internal documents about an upcoming acquisition. Melissa uses a public AI chatbot through a personal account, pastes six pages of customer material into it, and asks for an executive summary. The output is excellent, the presentation succeeds, and no obvious security alert occurs.

Yet Northstar may still have a cybersecurity problem because customer information was transferred to an external service that had never been evaluated or approved.

- What went wrong? The organization had no approved AI platform, no information-handling guidance, and no easy approval process.
- What should have happened? Northstar should have acknowledged likely AI use, provided an approved path, and clearly defined what information could be submitted.
- Leadership lesson: Lack of formal AI adoption does not mean lack of AI use. Visibility comes first.

11.34 A Realistic Scenario: The AI-Generated Security Fix

A systems administrator named Kevin is troubleshooting a remote-access server. He pastes part of the configuration into an AI assistant and asks for a quick fix to certificate errors. The AI recommends temporarily disabling certificate verification. Kevin applies the change, the connection problem disappears, and help-desk complaints stop.

Two weeks later the security team discovers that the change weakened server identity verification. Kevin did not intend to disable security; he accepted a technically convenient answer without evaluating the broader consequence.

AI-generated technical guidance should be treated as a suggestion, not authority. High-impact changes affecting authentication, encryption, firewalls, privileges, or monitoring should receive appropriate review and be compared with trusted documentation.

11.35 A Realistic Scenario: The Deepfake Payment Request

The controller of a mid-sized company receives a video call from what appears to be the chief executive. The executive requests a $275,000 transfer for a confidential acquisition and insists normal procedures be bypassed. The face, voice, and mannerisms appear familiar. The controller processes the payment. Later, the real executive denies making the request.

The most important control question is not how sophisticated the deepfake was. It is why one unusual conversation could authorize a $275,000 transfer.

Large or unusual financial transactions should require established approval procedures regardless of whether the request arrives by email, telephone, video, text message, or in person.

11.36 Build Verification Into the Business Process

Security cannot depend entirely on employees detecting deception. Attackers will sometimes succeed in sounding convincing, and employees will sometimes make mistakes.

The stronger approach is to design business processes so that one mistake does not automatically produce a catastrophic result. Vendor bank-account changes can require confirmation through a known contact. Large transfers can require two approvals. Privileged password resets can require additional verification. Production firewall changes can require peer review.

AI strengthens the argument for these controls because deception is becoming cheaper and easier to produce.

11.37 Questions Every Business Leader Should Ask About AI

- Do we know which AI systems employees currently use and which are formally approved?
- Are employees using personal AI accounts for company work?
- What information are employees allowed to submit, and what information is prohibited?
- Can customer, employee, source-code, regulated, or government-related information be entered?
- Are enterprise AI accounts centrally managed and protected with MFA?
- What happens to prompts and uploads, how long are they retained, and are they used for model improvement?

- Can AI agents send email, modify files, create or delete records, approve transactions, or call external APIs?
- Which outputs require human verification?
- How would an employee report accidental disclosure through an AI tool?
- Who is responsible for managing AI risk?

11.38 Questions to Ask Your IT Team or Service Provider

1. Which generative AI applications are currently visible in our environment?
2. Which AI tools have been formally approved?
3. Can we identify employees using unauthorized AI applications?
4. Can approved AI systems use our existing identity provider and MFA?
5. Can user accounts be disabled centrally?
6. What AI-related activity can we log?
7. Which AI tools have access to Microsoft 365, Google Workspace, SharePoint, CRM systems, source-code repositories, or internal platforms?
8. Do any AI integrations have write or delete authority?
9. Can those permissions be reduced?
10. Are AI service accounts using least privilege?
11. Are prompts or outputs retained, and where?
12. Can retention be reduced?
13. Can administrators restrict external connectors?
14. How would we revoke an AI integration during an incident?
15. Can data-loss prevention controls detect sensitive information being submitted to AI services?
16. Are developers permitted to submit proprietary source code to external AI systems?
17. How do we protect API keys and credentials used by AI applications?
18. How do we review AI-generated code before deployment?
19. Are agent actions logged sufficiently for investigation?
20. Can high-risk AI actions require approval?
21. Do incident-response procedures include AI-related data exposure?
22. What are the three highest-priority improvements we should make to our AI security posture?

11.39 A Practical AI Cybersecurity Checklist

- Governance: We know which AI systems are being used, approved systems are documented, responsibility is assigned, and employees have practical guidance.

- Information protection: Employees know what information may be entered; confidential, regulated, government, and credential information receives appropriate restrictions.
- Identity and access: Enterprise accounts are centrally managed, MFA is enabled where available, departed users are removed promptly, and AI service accounts use least privilege.
- Agents and automation: AI agents have clearly defined authority, high-impact actions require appropriate approval, and activity is logged.
- Vendor management: Providers have been evaluated, data-retention and model-improvement practices are understood, and contractual protections are documented.
- Software development: AI-generated code receives review and testing; dependencies and secrets are checked.
- Social engineering: High-risk financial requests require independent verification, and familiar voice or video is not treated as sufficient proof.
- Incident response: Employees know how to report accidental AI disclosure, logs can be preserved where available, and connectors or tokens can be disabled or revoked.

11.40 Key Terms in Plain English

- Artificial Intelligence (AI): Technology that performs tasks associated with human intelligence, such as recognizing patterns, generating content, making predictions, or assisting with decisions.
- Generative AI: AI designed to generate new content such as text, images, code, audio, or video.
- Large Language Model (LLM): A type of AI model trained on large amounts of language data and designed to understand and generate text.
- Prompt: Information or instructions provided to a generative AI system.
- Prompt Injection: An attempt to manipulate an AI system through instructions designed to interfere with its intended behavior.
- Indirect Prompt Injection: A prompt-injection attack delivered through content the AI retrieves or processes rather than through a direct user instruction.
- AI Agent: An AI-enabled system that can perform tasks or actions through tools, applications, APIs, or other systems.
- Connector: An integration that permits an AI system to interact with another application or data source.
- Hallucination: A commonly used term for AI-generated information that appears plausible but is incorrect, unsupported, or fabricated.

- Human-in-the-Loop: A design in which a human reviews, approves, corrects, or controls AI-generated decisions or actions.
- Least Privilege: Providing only the minimum access required to perform a legitimate function.
- Shadow AI: AI technology used without appropriate organizational visibility, approval, or governance.
- Retrieval-Augmented Generation (RAG): A technique that retrieves external or organizational information and uses it as context for AI-generated answers.
- Deepfake: Synthetic or manipulated media designed to make a person appear to say or do something that did not occur.
- Voice Cloning: AI-based generation of speech intended to imitate a particular person's voice.
- AI Governance: Policies, responsibilities, controls, processes, and oversight used to manage AI use and risk.
- Sensitive Information Disclosure: Exposure of confidential, regulated, proprietary, credential, customer, employee, or other protected information to an unauthorized person or system.

11.41 TRG Practical Takeaway

Artificial intelligence does not require businesses to abandon the cybersecurity principles they already understand. It makes many of those principles more important.

Know what information you have. Know where it goes. Control access. Use least privilege. Verify important transactions. Review vendors. Protect credentials. Monitor activity. Train employees. Prepare for incidents. Maintain accountability.

For most organizations, the most important first step is determining which AI tools are already being used and establishing simple rules for what information may be placed into them. From there, governance can improve gradually.

11.42 Before You Move On

23. Which artificial intelligence systems are our employees using?
24. What information are they permitted to submit?
25. Which AI systems have access to our internal applications or data?
26. Which AI-generated decisions or actions require human verification?
27. Who is responsible for managing AI risk?

11.43 Key Takeaways

Artificial intelligence is already part of everyday business. Organizations should assume employees are experimenting with AI unless they have strong evidence otherwise.

The greatest immediate risk may not be a sophisticated attack against an AI model. It may be an employee placing sensitive information into an inappropriate service. AI can also make phishing, impersonation, synthetic audio, fraudulent communications, and other social-engineering attacks more convincing.

Generative AI can improve cybersecurity operations, software development, documentation, research, and productivity. But its output should not automatically be treated as verified truth. AI-generated code requires review, AI-generated security recommendations require validation, and AI-generated business decisions require accountability.

AI agents deserve particular attention because they can move from generating information to performing actions. The more authority an AI system receives, the more important least privilege, logging, approval controls, testing, and monitoring become.

Organizations do not need to prohibit artificial intelligence in order to manage its risk. They need visibility, approved tools, understandable policies, appropriate access controls, vendor oversight, human accountability, and incident readiness.

11.44 Chapter Summary

Artificial intelligence is changing both business productivity and cybersecurity. Employees can use AI to summarize documents, write software, draft communications, analyze information, and automate work. Attackers can use similar capabilities to improve phishing, impersonation, research, automation, and deception.

The most practical organizational approach is neither uncontrolled adoption nor reflexive prohibition. It is governed adoption.

Organizations should know which systems are being used, understand what information those systems receive, evaluate service providers, control access, restrict high-impact actions, verify important outputs, and maintain human accountability.

Artificial intelligence should be treated as another powerful information-processing capability within the organization's existing cybersecurity program. Security principles do not disappear because AI is involved. They become more important.

11.45 Sources and Guidance for Further Reading

The guidance in this chapter should be reviewed periodically because artificial intelligence technology, security guidance, and organizational use are changing rapidly.

National Institute of Standards and Technology (NIST), Artificial Intelligence Risk Management Framework (AI RMF) and NIST AI 600-1, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. https://www.nist.gov/itl/ai-risk-management-framework

OWASP GenAI Security Project, OWASP GenAI / LLM application security guidance and Top 10 resources. https://genai.owasp.org/

Organizations developing or acquiring AI-enabled software should also consult current secure software development guidance from NIST and relevant sector-specific requirements.

Transition to Chapter 12

Artificial intelligence demonstrates an important theme that has appeared throughout this book. Technology changes. Attack methods change. Business practices change. A technical incident can quickly become an organization-wide crisis.

A suspicious login can become an account compromise. A fraudulent email can become a financial loss. A stolen laptop can expose company information. Ransomware can interrupt operations. A vendor may report a breach. Sensitive information may be placed accidentally into an unauthorized cloud or AI service.

At that point, leadership must coordinate legal obligations, insurance, communications, financial decisions, vendors, customers, business continuity, and recovery while the technical team handles containment and investigation.

That broader leadership challenge is the focus of Chapter 12: Managing a Cyber Crisis: Legal, Insurance, Communications, and Business Decisions.

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

Chapter 12

Managing a Cyber Crisis: Legal, Insurance, Communications, and Business Decisions

Authored by Dr. Charles Edwards

Every organization hopes its cybersecurity protections will prevent a serious incident.

Some incidents will still happen.

An employee will click something dangerous. A password will be stolen. A laptop will disappear. A cloud account will be compromised. A vendor will report a breach. Malware will execute. A business email account will be used for fraud. Sensitive information will be sent to the wrong person. An unauthorized AI service will receive protected information. A security tool will detect activity that nobody immediately understands.

The question is not whether every cybersecurity incident can be prevented. No organization can guarantee that.

The more useful question is: When something goes wrong, will we recognize it quickly, make good decisions, limit the damage, preserve what matters, restore operations, and learn from what happened?

That is cyber-crisis management built on effective incident response.

Chapter 6 focused on the operational response process. This chapter looks at the broader business crisis that can develop around that technical response.

Leadership may need to coordinate legal obligations, insurance, communications, financial decisions, vendors, customers, business continuity, and recovery while the technical team handles containment and investigation.

A serious cybersecurity incident therefore tests much more than the IT department.

It tests the organization as a whole.

A strong response does not begin when the attacker arrives. It begins before the incident.

12.1 The Business Problem: Incidents Create Pressure Before They Create Clarity

Imagine arriving at work on Monday morning and discovering that employees cannot open company files.

Several computers display a ransom note. Email still works for some employees but not others. The accounting department says its shared drive is unavailable. A server administrator believes backups may also have been affected. Customers are beginning to call.

Someone asks whether the company should shut down every computer. Another person wants to pay the ransom immediately. The insurance broker is requesting information. The company's president asks: “What exactly happened?”

Nobody knows yet.

That uncertainty is normal during the early stages of a cybersecurity incident. The first information available is often incomplete. Some of it may be wrong. Employees may be frightened. Managers may be under pressure. Customers may want answers before the organization has answers to give. Executives may feel compelled to make immediate decisions. Attackers may still be active.

This combination of uncertainty and urgency makes incident response difficult.

The organization must act without pretending to know more than it knows. That requires preparation.

A company that waits until the first major incident to decide who has authority, whom to call, where the backups are, how systems can be isolated, or how customers will be notified has created unnecessary risk.

Incident response is therefore partly about technology. It is also about decision-making under pressure.

12.2 An Event Is Not Always an Incident

Organizations generate thousands or millions of technology events.

A user enters the wrong password. A firewall blocks an Internet connection. An endpoint-protection system stops a suspicious file. A cloud administrator changes a configuration. A laptop loses its network connection. A user deletes a document. A vulnerability scanner identifies outdated software.

These events do not automatically represent cybersecurity incidents.

An event is something that occurred.

An alert is a signal suggesting that an event or group of events may deserve attention.

A cybersecurity incident is an occurrence that actually or potentially jeopardizes information, systems, operations, or security objectives sufficiently to require coordinated response.

The distinction matters because organizations cannot treat every alert as a crisis.

At the same time, they cannot ignore warning signs simply because complete proof is not yet available.

Incident response therefore requires triage.

What happened? Which system is involved? Which account is involved? Is the activity continuing? Could sensitive information be affected? Could operations be interrupted? How confident are we? What must happen next?

Good incident response improves those answers over time.

12.3 Detection Is the Moment the Organization Begins to Learn

Many incidents do not announce themselves dramatically.

There may be no ransom note. No system may suddenly shut down.

The first sign could be something small.

An employee receives an unexpected MFA notification. A customer asks why the company sent a strange invoice. A security tool reports a suspicious process. An administrator notices a login from an unusual location. A payroll employee discovers that a direct-deposit account was changed. A user reports that files are disappearing. A cloud provider reports suspicious account activity. A vendor announces that its environment was compromised. Someone notices an unfamiliar forwarding rule in an executive mailbox.

Each may be the beginning of a larger investigation.

This is why employees are part of incident detection.

Sophisticated monitoring tools are valuable, but employees often see business abnormalities before technology identifies them.

A finance employee may recognize that a payment request is unusual. A receptionist may recognize a strange call. A salesperson may notice customer complaints. A system administrator may recognize a configuration change that does not make sense.

Employees need a simple way to report those concerns.

If reporting requires finding a twenty-page procedure, employees may not report quickly.

A simple instruction is often better:

If something involving company systems, accounts, money, or sensitive information seems wrong, report it immediately.

12.4 The First Ten Minutes Matter - But Panic Does Not Help

When employees believe an incident is occurring, their first instinct may be to fix it.

Sometimes that instinct creates additional problems.

An employee may delete suspicious email. An administrator may erase logs. A manager may restart a compromised server. A technician may reimage a laptop before investigators understand what happened. Someone may contact an attacker directly. Someone may post information about the event on social media. An employee may begin deleting files in an attempt to "clean things up."

These actions may destroy evidence or make the incident harder to understand.

That does not mean employees should do nothing.

It means they should know which immediate actions are appropriate.

For example, if ransomware is actively spreading from a workstation, isolating that device from the network may be appropriate. If an employee believes credentials were entered into a fraudulent website, the account may need to be secured quickly. If a fraudulent payment is underway, the bank may need immediate notification.

Incident response is about controlled urgency.

Fast action can be important.

Uncoordinated action can make things worse.

12.5 Build an Incident-Response Team Before You Need One

A small organization may not have a dedicated cybersecurity incident-response team.

That is acceptable.

It still needs to know who will participate.

The team may include people from leadership, IT, cybersecurity, legal counsel, insurance, finance, Human Resources, communications, operations, and external service providers.

Not every incident requires everyone.

A lost laptop may involve IT and management. A ransomware incident may involve nearly the entire group. A fraudulent wire transfer may immediately involve finance, leadership, banking contacts, IT, insurance, and legal counsel.

The important issue is not the size of the team.

It is knowing who owns which decisions.

Someone must coordinate the response. Someone must make technical containment decisions. Someone must communicate with leadership. Someone must document events. Someone must decide when outside help is required. Someone must address customer, regulatory, contractual, or law-enforcement communication.

Without defined responsibility, organizations often lose valuable time asking: “Who is handling this?”

12.6 Create an Incident Contact List That Works During an Incident

Organizations frequently maintain emergency contact information inside systems that may become unavailable during an attack.

That creates a problem.

If ransomware disables the network, the company may lose access to the document containing the ransomware-response contact list.

Incident contacts should therefore be available through a method that remains accessible when normal systems are unavailable.

The list may include executive leadership, IT staff, managed service provider, cybersecurity incident-response provider, cyber insurance carrier or breach hotline, outside legal counsel, banking contacts, critical technology vendors, cloud providers, communications personnel, law enforcement contacts when appropriate, and important contractual contacts.

The list should include alternate communication methods.

If company email is compromised, employees should not rely exclusively on company email to coordinate the response.

An organization may establish a separate emergency communication method for serious incidents.

The exact solution matters less than the principle:

Do not make your incident-response process completely dependent on the systems that may be under attack.

12.7 Know Who Has Authority to Disconnect Systems

Containment decisions can have major business consequences.

Suppose investigators believe a critical server is compromised.

Should it be disconnected immediately?

From a cybersecurity standpoint, isolation may prevent additional damage.

From an operations standpoint, disconnecting the server may halt production.

Someone needs authority to make that decision.

Similarly: Who can disable an executive account? Who can block remote access? Who can shut down Internet connectivity? Who can isolate an entire office? Who can disable a cloud application? Who can take a public website offline? Who can stop financial transactions?

The answer should not be invented during an emergency.

Organizations should establish decision authority in advance.

For routine technical containment, IT or cybersecurity personnel may have delegated authority.

For actions that could stop the business, executive approval may be appropriate unless immediate action is necessary to prevent catastrophic harm.

The incident-response plan should make those boundaries clear.

12.8 Preserve Evidence Before Destroying It

A cybersecurity incident creates information that may later help explain what happened.

That information may include security logs, email headers, authentication records, endpoint telemetry, firewall logs, cloud audit logs, malicious files, memory captures, disk images, screenshots, network traffic, user reports, telephone records, and timestamps.

Organizations should preserve relevant evidence when practical.

This does not mean every business must conduct a full forensic investigation internally.

It means employees should avoid unnecessarily destroying information that investigators may need.

For serious incidents, professional digital-forensics or incident-response assistance may be appropriate.

12.9 Documentation Matters More Than People Expect

Cybersecurity incidents can become confusing very quickly.

Without documentation, the team may lose track of what happened and why decisions were made.

A response log should record important events.

For example:

8:07 a.m. - Finance reports inaccessible shared drive.

8:14 a.m. - IT confirms encrypted files on FIN-PC-07.

8:19 a.m. - Device isolated from network.

8:31 a.m. - Similar activity detected on two additional endpoints.

8:44 a.m. - Cybersecurity provider contacted.

9:02 a.m. - Domain administrator credentials reset.

9:26 a.m. - Cloud backups confirmed unaffected.

10:15 a.m. - Executive response team convened.

The log does not need literary quality.

It needs accuracy.

Documentation helps investigators reconstruct the incident. It helps leadership understand decisions. It may help with insurance, legal analysis, regulatory reporting, lessons learned, and future improvements.

In an incident, memory is not enough.

12.10 Containment Is About Stopping the Damage From Growing

Once an organization confirms or strongly suspects malicious activity, containment becomes a priority.

Containment asks: How do we prevent the attacker or problem from causing additional harm while preserving our ability to understand and recover from the incident?

Possible containment actions include isolating an endpoint, disabling a compromised account, resetting credentials, revoking authentication tokens, blocking malicious IP addresses or domains, disabling remote access, removing unauthorized forwarding rules, restricting network segments, disabling compromised applications, or temporarily stopping selected business processes.

The correct action depends on the incident.

Containment can also create tradeoffs.

That is why significant incidents may require experienced responders.

12.11 Do Not Assume the First Compromised Device Is the Only One

Attackers frequently move.

A compromised laptop may be only the first visible symptom.

An attacker may have already stolen credentials, accessed email, created forwarding rules, connected to cloud applications, moved laterally to servers, created additional accounts, installed remote-access software, stolen information, or established persistence elsewhere.

Incident investigation therefore asks a broader question:

How far did this go?

Scope is one of the most important concepts in incident response.

An organization cannot confidently recover until it understands what it is recovering from.

12.12 Eradication Means Removing the Cause, Not Just the Symptom

After containment, the organization must remove the mechanisms that allowed the incident to continue.

This is sometimes called eradication.

Changing a password may be necessary, but additional work may also be required.

Malicious inbox rules may need removal. Active sessions may need revocation. Unauthorized applications may need to be removed. MFA methods may need review. Other accounts may require investigation.

Eradication asks:

What allowed this incident to persist?

That mechanism must be addressed.

12.13 Recovery Is More Than Turning Systems Back On

Organizations naturally want to restore operations quickly.

But recovery performed too quickly can reintroduce the problem.

Before returning affected systems to normal operation, the organization should have reasonable confidence that malicious access has been removed, credentials have been secured, important vulnerabilities have been addressed, restored systems are trustworthy, security monitoring is active, and backups used for restoration are safe.

Recovery may happen gradually.

Critical services may return first.

Lower-priority systems may return later.

Recovery should connect directly with business-continuity planning.

12.14 A Backup Is Valuable Only If It Can Be Restored

Ransomware incidents frequently turn attention toward backups.

Leadership asks: “Do we have backups?”

The next questions are: When was the last successful backup? What information is included? Are the backups isolated from the production environment? Can attackers delete them? Are they encrypted? How long are they retained? Have restoration procedures been tested? How long will restoration take? Who knows how to perform it?

Backups create recovery options only when they are usable.

Testing reveals problems before an emergency.

Recovery exercises are therefore a cybersecurity control.

12.15 Do Not Confuse Restoration With Resolution

Suppose a company restores its systems from backup and operations resume.

That does not necessarily mean the incident is over.

Important questions may remain.

Did information leave the organization before systems were encrypted? Were customer records accessed? Were credentials stolen? Did the attacker reach cloud systems? Do reporting obligations apply? Could another persistence mechanism remain? What vulnerability allowed initial access? Were vendor systems involved?

Recovery and incident closure are related.

They are not identical.

12.16 Communication Can Become a Second Incident

Poor communication can make a cybersecurity incident worse.

During an incident, different audiences may need different information.

Employees need instructions. Leadership needs situational awareness. IT staff need technical details. Customers may need reassurance or notification. Insurance providers may need specific facts. Attorneys may need evidence. Regulators or government customers may require particular reports.

The organization should coordinate communication.

Communication should distinguish what is known from what is suspected.

Accuracy creates credibility.

12.17 Establish an Internal Communication Rhythm

Leadership does not need a telephone call every time an investigator discovers a new log entry.

It does need regular updates.

A regular briefing schedule helps separate investigation from executive decision-making.

Each briefing can answer: What do we know? What has changed? What is affected? What actions have been taken? What business impact exists? What decisions are needed? What do we still not know?

That structure reduces confusion.

12.18 Cyber Insurance Should Be Understood Before an Incident

Many organizations purchase cyber insurance.

The policy may provide valuable resources following an incident.

But businesses should understand how the policy works before they need it.

Questions may include who must be contacted first, whether approved incident-response firms must be used, what expenses require prior authorization, what evidence must be preserved, and what reporting timelines exist.

The time to learn how the policy works is not while ransomware is spreading.

12.19 Legal Counsel May Be Important Earlier Than Expected

Cybersecurity incidents can create legal questions quickly.

Does notification law apply? Does a contract require reporting? What communications should be preserved? What information may be disclosed? What obligations exist to customers or government entities?

The answers depend on the organization, jurisdiction, contract, and incident.

Strong incident response brings appropriate disciplines together.

12.20 Law Enforcement Can Be Part of the Response

Some cybersecurity incidents involve criminal activity.

Organizations may decide to contact law enforcement.

For serious cybercrime, organizations in the United States may interact with the FBI, the U.S. Secret Service, or local law enforcement, depending on the event. CISA may also provide cybersecurity assistance and coordination.

The organization's own response responsibilities remain.

12.21 Incident Response Is Now Integrated With Cybersecurity Risk Management

Modern incident-response guidance increasingly emphasizes that preparation, detection, response, recovery, governance, asset knowledge, and protection are interconnected.

NIST SP 800-61 Rev. 3 reflects this approach by integrating incident-response considerations across the six functions of the NIST Cybersecurity Framework 2.0: Govern, Identify, Protect, Detect, Respond, and Recover.

The goal is not simply to survive one incident.

It is to become harder to compromise and better prepared for the next one.

12.22 Ransomware Creates Both Technical and Business Decisions

Ransomware incidents create immediate pressure because they often affect both technology and business operations at the same time.

Files may be encrypted. Servers may become unavailable. Employees may be unable to work. Customers may be affected. Backups may be threatened. Attackers may claim they have stolen information. The organization may receive a ransom demand.

At that moment, leadership may ask: Should we pay?

That question is understandable. It is also more complicated than it first appears.

Payment does not guarantee that systems will be restored. It does not guarantee that stolen information will be deleted. It does not guarantee that the attacker will provide a working decryption tool. It does not guarantee that the organization will avoid future targeting.

The decision therefore should not be made casually or by one person acting under pressure.

The first priority should usually be understanding and containing the incident.

12.23 Double Extortion Changes the Ransomware Problem

Traditional ransomware focused primarily on encryption.

Modern ransomware groups may also steal information before encryption.

They then threaten to publish or sell the information if payment is not made.

This is often called double extortion.

Recovery from encryption does not erase the data-breach question.

12.24 Email Account Compromise Can Be Financially Devastating

Not every major incident involves malware.

A compromised email account can produce serious losses.

A finance manager's Microsoft 365 credentials are stolen. The attacker quietly observes email for several weeks, learns how invoices are processed, creates a hidden forwarding rule, and waits for a real invoice.

When the invoice arrives, the attacker sends a convincing follow-up message changing the payment instructions.

The finance department sends the payment to the criminal's account.

No virus is detected. No system is encrypted. No ransom note appears.

Yet the company may lose hundreds of thousands of dollars.

This is a cybersecurity incident.

12.25 Treat Fraudulent Payment Incidents as Time-Sensitive Emergencies

Organizations should have a predefined procedure for suspected wire or ACH fraud.

The procedure should answer: Who contacts the bank? Which telephone number should be used? Who contacts the vendor? Who contacts law enforcement? Who preserves the email evidence? Who investigates account compromise? Who informs leadership? Who contacts insurance?

Speed matters.

12.26 Lost and Stolen Devices Require More Than Replacement

A lost laptop is often treated as a hardware problem.

Replace the laptop. Restore the user's files. Move on.

But the cybersecurity questions are different.

Was the device encrypted? Was the user logged in? Were credentials stored? Did the device contain locally stored sensitive information? Could cloud sessions still be active? Can the device be remotely locked or wiped?

The incident classification should therefore depend on exposure, not merely replacement cost.

12.27 Mobile Phones Can Create Similar Exposure

Phones and tablets now contain email, messaging, documents, authentication applications, banking applications, cloud storage, and business contacts.

A lost mobile device may therefore affect far more than telephone service.

For bring-your-own-device environments, incident procedures should clearly explain what authority the organization has over business data on personal devices.

12.28 Vendor Incidents Can Become Your Incidents

Organizations depend on cloud providers, payroll systems, software companies, managed service providers, payment processors, legal firms, accounting firms, hosting providers, and other vendors.

A cybersecurity incident at one of those companies may affect your organization even when your own systems were not directly attacked.

Third-party incidents should therefore be included in incident-response planning.

12.29 Managed Service Provider Compromise Requires Special Attention

A managed service provider may have privileged access to many customer environments.

That makes the provider both valuable and potentially high-impact.

The organization remains responsible for understanding the impact on its own environment.

Outsourcing IT does not outsource all incident accountability.

12.30 Cloud Incidents May Leave Different Evidence

Traditional investigations often focus on physical computers and local network logs.

Cloud incidents may require different evidence.

Useful records may include cloud sign-in logs, audit trails, administrative changes, mailbox rules, OAuth application grants, session information, sharing activity, file-access history, API activity, and conditional-access events.

Logging and retention decisions are therefore part of incident readiness.

12.31 AI-Related Incidents Need Their Own Response Questions

Artificial intelligence introduces new incident scenarios.

An employee may paste confidential information into an unapproved AI tool. An AI agent may perform an unauthorized action. A connector may expose data beyond intended permissions. A prompt-injection attack may influence an AI system. A developer may accidentally expose an API key while using an AI assistant.

The organization should treat AI-related exposure as a data-handling and access-control problem, not as something outside normal cybersecurity governance.

12.32 Do Not Hide Mistakes That Involve AI or Cloud Services

Employees may hesitate to report accidental disclosure.

That delay can make the incident harder to manage.

Organizations should encourage rapid reporting.

The response should focus first on reducing harm.

12.33 A Realistic Scenario: The Monday Morning Ransomware Incident

Consider a regional architecture and engineering firm named Harbor Design Group.

The company has 110 employees.

It uses Microsoft 365, a local file server, several cloud applications, and a managed IT provider.

At 7:42 a.m. on Monday, an employee reports that project files will not open.

At 7:48 a.m., another employee sees strange file extensions.

At 7:54 a.m., a ransom note appears.

By noon, the organization understands that this is not merely an encryption event. It is also a possible data-exposure event.

What Went Right?

Employees reported the problem quickly.

The company had an IT provider that could respond.

Systems were isolated.

Backups existed.

The insurer was contacted.

Leadership became involved.

What Went Wrong?

The administrator account used in the attack had excessive access.

Remote access had not been reviewed closely.

Cloud and network logs had insufficient retention periods.

The company had never rehearsed ransomware response.

What Should Have Happened?

The organization should have established an incident-response plan before the event.

Privileged accounts should have received stronger controls.

Remote access should have been monitored.

Critical logs should have been retained.

Backup restoration should have been tested.

12.34 A Realistic Scenario: The Compromised Executive Mailbox

A manufacturing company receives an email from its chief financial officer asking the accounts-payable manager to update banking information for an important supplier.

The message is part of a legitimate email chain.

Two weeks later, the supplier calls asking why payment has not been received.

The payment was sent to an attacker.

The lesson is simple: Cybersecurity controls and financial controls must support each other.

12.35 A Realistic Scenario: The Stolen Laptop

An employee leaves a company laptop in a vehicle overnight.

The vehicle is broken into.

The laptop is stolen.

The laptop is encrypted. The user account requires MFA. The device can be remotely disabled.

The lesson is that security controls determine the consequence of an incident.

12.36 A Realistic Scenario: The Vendor Breach

A payroll provider sends an urgent notice explaining that unauthorized access may have exposed employee data.

Leadership contacts legal counsel, insurance, and the provider.

Employees are warned about likely phishing.

The lesson is that vendor incidents should already be part of the incident-response plan.

12.37 Tabletop Exercises Turn Plans Into Practice

An incident-response plan can look excellent on paper and fail during a real event.

A tabletop exercise helps expose those weaknesses.

A tabletop is a discussion-based simulation.

Participants are given a realistic scenario and explain what they would do.

The exercise reveals gaps before a real attacker does.

12.38 Tabletop Exercises Should Include Leadership

Incident-response exercises are sometimes treated as IT drills.

That misses much of the value.

Executives should participate in scenarios involving ransomware, financial fraud, cloud compromise, vendor breaches, sensitive-data exposure, and prolonged system outages.

The purpose is to test whether the organization can make coordinated business decisions under pressure.

12.39 Test Specific Decisions, Not Just General Awareness

A weak tabletop exercise asks: “What would you do during ransomware?”

A stronger exercise asks: “At 9:10 a.m., investigators recommend disconnecting the ERP server. Doing so will stop shipping for the rest of the day. Who has authority to approve that?”

Specific questions produce specific learning.

12.40 Conduct an After-Action Review

After a serious incident or exercise, the organization should review what happened.

The purpose is not to assign blame.

The purpose is to improve the system.

The results should become tracked actions.

Otherwise, the meeting becomes a conversation that produces no improvement.

12.41 Measure Improvement After Incidents

Organizations should not treat an incident as complete when the final meeting ends.

Improvements should be assigned, prioritized, and tracked.

Each action should have an owner.

That turns lessons into risk reduction.

12.42 Questions Every Business Leader Should Ask

Leadership should be able to answer questions about responsibility, contacts, critical systems, backups, communication, decision authority, reporting, and recovery before a major incident.

If those answers are unclear, incident planning is incomplete.

12.43 Questions to Ask Your IT Team or Service Provider

1. What happens when we report a suspected cybersecurity incident?
2. Who answers after normal business hours?
3. What is the escalation process?
4. Which incidents can you handle internally?
5. When do you bring in outside incident-response specialists?
6. Which security logs do we currently retain?
7. How long are they retained?
8. Can logs be exported quickly?
9. Can compromised devices be isolated remotely?
10. Can cloud sessions be revoked?
11. Can privileged accounts be disabled quickly?
12. Can you identify unusual mailbox rules?
13. Can you identify suspicious OAuth applications?
14. Do we have endpoint detection and response?
15. Which systems are covered?
16. Are critical servers monitored?
17. How are backups protected from ransomware?
18. When was the last full restore test?
19. How long would it take to restore critical systems?
20. Do we have an inventory of privileged accounts?
21. Do we know which vendor accounts have remote access?
22. How would we respond if your MSP account were compromised?
23. Can you provide an incident timeline after an event?
24. How do you preserve forensic evidence?

25. What are the three most important incident-response improvements we should make this year?

12.44 Chapter 12 Cyber Crisis Management Checklist

Preparation:

- An incident-response owner has been identified.
- Key roles and responsibilities are documented.
- Emergency contact information is current.
- Alternate communication methods exist.
- Cyber insurance contact procedures are known.
- Legal and banking contacts are documented.
- Critical vendors are identified.

Detection:

- Employees know how to report suspicious activity.
- Security alerts are monitored.
- Cloud activity is logged.
- Endpoint monitoring is active.
- Suspicious financial activity is escalated.

Containment:

- Compromised devices can be isolated.
- Accounts can be disabled rapidly.
- Sessions and tokens can be revoked.
- Remote-access systems can be restricted.
- Decision authority for major containment actions is documented.

Evidence:

- Relevant logs can be preserved.
- Suspicious email can be retained.
- Important systems have sufficient logging.

- Investigators know where cloud evidence is located.
- Employees understand not to destroy evidence unnecessarily.

Recovery:

- Backups are protected.
- Restoration procedures are documented.
- Restoration is tested.
- Critical systems have recovery priorities.
- Recovered systems receive appropriate security validation.

Communication:

- Leadership receives regular incident updates.
- Employees know where to obtain instructions.
- Customer communication is coordinated.
- Public statements require approval.
- Known facts are distinguished from assumptions.

Third Parties:

- Vendor incidents are included in the response plan.
- MSP compromise procedures exist.
- Vendor emergency contacts are documented.
- Provider access can be revoked.

Financial Fraud:

- Bank fraud-response contacts are current.
- Payment changes require independent verification.
- High-value transactions use appropriate approval controls.
- Suspected fraud is escalated immediately.

AI and Cloud Incidents:

- AI-related data disclosure can be reported.
- AI tokens and connectors can be revoked.

- Cloud audit logs are retained.
- Unauthorized applications can be removed.

Improvement:

- Tabletop exercises are conducted periodically.
- After-action reviews are performed.
- Improvement items receive owners.
- Incident lessons are incorporated into policy and controls.

12.45 Key Terms in Plain English

Cybersecurity Event: Any observable occurrence involving a system, account, network, application, or information.

Alert: A notification suggesting that an event may require investigation.

Cybersecurity Incident: An event or group of events that threatens or affects information, systems, operations, or security objectives sufficiently to require coordinated response.

Triage: The process of determining severity, scope, priority, and appropriate next actions.

Containment: Actions taken to stop or limit additional harm.

Eradication: Removal of malicious access, persistence, vulnerabilities, or other causes that allow an incident to continue.

Recovery: The process of restoring systems and business operations safely.

Scope: The extent of systems, accounts, data, people, and operations affected by an incident.

Digital Evidence: Information that may help explain what happened during a cybersecurity incident.

Chain of Custody: Documentation showing how evidence was collected, handled, transferred, and preserved.

Endpoint Detection and Response (EDR): Technology used to monitor, investigate, and respond to suspicious activity on endpoints.

Business Email Compromise (BEC): Fraud involving compromised or impersonated business email accounts, often used to redirect payments or obtain sensitive information.

Ransomware: Malware or related extortion activity designed to disrupt access to information or systems, often accompanied by payment demands.

Double Extortion: A ransomware tactic in which attackers both encrypt information and threaten to release stolen data.

Tabletop Exercise: A discussion-based simulation used to test plans, roles, decisions, and communication.

After-Action Review: A structured review performed after an incident or exercise to identify lessons and improvements.

Out-of-Band Communication: A communication method that does not rely on the potentially compromised primary system.

Incident Timeline: A chronological record of important events, observations, and response actions.

12.46 TRG Practical Takeaway

A cybersecurity incident will rarely unfold exactly as expected.

That is why organizations should prepare around principles rather than scripts alone.

Know who is responsible. Know who to call. Know which systems matter most. Know how to communicate when normal systems are unavailable. Know how to isolate compromised devices and accounts. Know how backups will be restored. Know how financial fraud will be handled. Know how vendor incidents will be escalated. Know which decisions require executive authority.

Most importantly, practice.

An incident-response plan that has never been tested is still a theory.

A simple tabletop exercise can expose weaknesses before a real attacker does.

The objective is not perfect response.

The objective is organized response.

12.47 Before You Move On

Before continuing to the next chapter, leadership should be able to answer seven questions:

1. Who leads our cybersecurity incident response?
2. How do employees report suspected incidents?
3. How do we communicate if company email is unavailable?
4. Can we isolate compromised systems and accounts quickly?
5. Have we tested restoration from backup?
6. Do we know our cyber insurance, legal, banking, and vendor emergency contacts?
7. Have we conducted an incident-response exercise during the past year?

If several answers are unclear, the organization has identified its next set of priorities.

12.48 Key Takeaways

Cybersecurity incidents create uncertainty, urgency, and business pressure.

The organization should prepare before the incident.

Employees are an important source of detection.

Fast reporting matters.

Fast action should still be coordinated.

Containment limits damage.

Eradication removes the mechanisms that allowed the incident to continue.

Recovery restores trusted operations.

Evidence should be preserved where practical.

Documentation matters.

Backups should be tested, not merely assumed.

Financial fraud can be a cybersecurity incident even when malware is absent.

Vendor incidents can become customer incidents.

Cloud and AI incidents require the same disciplined response principles as traditional cybersecurity events.

Leadership, legal counsel, insurance, finance, IT, cybersecurity, vendors, and communications may all play roles depending on severity.

Tabletop exercises reveal gaps.

After-action reviews turn experience into improvement.

Incident response is not a side activity.

It is part of cybersecurity risk management.

12.49 Chapter Summary

No organization can guarantee that every cybersecurity incident will be prevented.

The practical objective is resilience.

A resilient organization detects problems, coordinates decisions, limits damage, preserves important evidence, communicates carefully, restores operations, and learns from what happened.

Incident response begins long before the first alert.

It begins with governance, asset knowledge, logging, backups, access control, vendor management, communication planning, and clear responsibility.

During an incident, organizations must resist the temptation to confuse speed with disorder.

The best response is controlled, documented, informed, and coordinated.

After the incident, the organization should ask not only: “How did the attacker get in?”

It should also ask: “What made the damage possible?” “What slowed our response?” “What worked well?” “What should we change?”

Every incident is expensive.

A mature organization extracts value from that experience by becoming better prepared for the next one.

Transition to Chapter 13

Managing a cyber crisis answers a broader leadership question:

How do we coordinate technical response, business continuity, legal obligations, insurance, communications, and recovery under pressure?

Once that capability exists, another question remains.

How does leadership know whether the organization's cybersecurity program is actually improving?

Buying security products is not enough. Writing policies is not enough. Completing training is not enough. Passing one assessment is not enough.

Organizations need a practical way to measure progress, identify weaknesses, prioritize improvements, and turn cybersecurity into an ongoing management process.

That brings us to Chapter 13: Measuring Cybersecurity and Building a Practical Improvement Plan.

Chapter 13 - Measuring Cybersecurity and Building a Practical Improvement Plan

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

Chapter 13

Measuring Cybersecurity and Building a Practical Improvement Plan

Authored by Dr. Charles Edwards

Cybersecurity improvement can be difficult to see.

A business may purchase new security tools. Employees may complete training. Multifactor authentication may be deployed. Backups may be improved. A managed service provider may install additional monitoring. Policies may be updated. A vulnerability scan may identify fewer problems than last year.

These are all potentially positive developments.

But leadership still needs to ask a harder question:

Are we actually becoming more secure?

That question cannot be answered by counting products, policies, or completed activities, and it cannot be answered simply because the organization has not experienced a major incident.

An organization can avoid a serious breach for years while still carrying substantial weaknesses. Another organization may experience an incident despite having a strong program.

Cybersecurity measurement therefore requires more than asking whether something bad happened.

The goal is to understand whether risk is being managed better.

That means examining whether the organization can govern cybersecurity, identify what matters, protect important assets, detect suspicious activity, respond effectively, and recover when something goes wrong.

Those ideas align closely with the six functions of the NIST Cybersecurity Framework 2.0: Govern, Identify, Protect, Detect, Respond, and Recover.

For a small or mid-sized business, the objective is not hundreds of measurements. It is a small set of meaningful indicators that help leadership decide what should improve next.

13.1 The Business Problem: Activity Is Not the Same as Progress

Imagine a company's annual cybersecurity meeting.

The IT provider presents a list of accomplishments: thirty laptops were replaced, a new firewall was installed, antivirus licenses were renewed, six security policies were

updated, employees completed annual awareness training, a penetration test was performed, and a backup appliance was upgraded.

The presentation contains many completed tasks. Leadership feels reassured.

Then someone asks: How many administrator accounts still do not use MFA? Nobody knows.

How quickly can we disable a compromised account? Nobody has measured it.

When did we last restore our most important system from backup? The answer is unclear.

How many former employees still have access to cloud applications? Nobody has checked recently.

How long do we retain Microsoft 365 security logs? Nobody in the meeting knows.

Which three cybersecurity risks concern us most? The answers differ depending on who is asked.

The company has been busy.

But activity is not the same as progress.

A cybersecurity program should not be judged solely by how much work was performed. It should be judged by whether important risks are being reduced.

13.2 Start With Outcomes, Not Products

A useful cybersecurity measurement program begins by asking what outcome the organization wants.

Instead of “We installed MFA,” ask: “Are important accounts protected against password-only compromise?”

Instead of “We purchased backup software,” ask: “Can we restore critical information within the time the business requires?”

Instead of “We bought endpoint detection software,” ask: “Can suspicious endpoint activity be detected, investigated, and contained quickly?”

Instead of “We completed employee security training,” ask: “Do employees recognize and report suspicious behavior?”

Instead of “We hired an IT provider,” ask: “Are security responsibilities clearly assigned and being performed?”

Products and activities support outcomes. They are not the outcome themselves.

13.3 A Metric Should Help Someone Make a Decision

Organizations sometimes collect measurements simply because the information is available.

A firewall can report millions of blocked connections. But that number alone may not tell leadership whether risk increased or what should be done next.

A useful metric should help someone decide something.

Examples include percentage of privileged accounts protected by phishing-resistant MFA, percentage of critical systems successfully restored during the last backup test, average time between employee termination and account disablement, number of critical vulnerabilities older than 30 days, and percentage of critical vendors reviewed during the past year.

Good metrics create visibility. Visibility supports decisions. Decisions create improvement.

13.4 Do Not Measure Everything

One of the fastest ways to make a cybersecurity scorecard useless is to include too much information.

A practical executive scorecard might include ten to fifteen measures such as MFA coverage, critical vulnerability remediation, backup restoration success, endpoint monitoring coverage, security logging coverage, phishing-reporting performance, privileged-account review, vendor review completion, incident-response exercise status, and critical-system recovery testing.

The exact list depends on the organization.

The goal is not quantity. The goal is usefulness.

13.5 Measure Coverage Before Measuring Sophistication

Organizations sometimes focus on advanced cybersecurity technologies while basic controls remain incomplete.

One of the most useful forms of cybersecurity measurement is coverage.

Examples include MFA coverage among users; encryption coverage among laptops; EDR coverage among managed endpoints; backup coverage for critical systems; vulnerability-scanning coverage for critical systems; security-awareness training completion; and the percentage of privileged accounts reviewed.

Coverage metrics are simple. That simplicity is valuable because they reveal incomplete implementation.

13.6 Measure Whether Controls Actually Work

Coverage alone is not enough.

Backup coverage can be 100 percent while restoration still fails.

Organizations should complement coverage with effectiveness testing.

Can backups actually be restored? Can MFA be bypassed through legacy authentication? Are endpoint alerts being reviewed? Are critical vulnerabilities actually being corrected? Can the organization execute its incident-response plan? Do employees report realistic phishing attempts?

Cybersecurity controls should not merely exist. They should work.

13.7 Current State and Target State

One of the most practical improvement methods is to compare two conditions.

Current state: Where are we today?

Target state: Where do we want to be?

For example, MFA may currently protect 72 percent of user accounts, while the target is 100 percent plus stronger authentication for privileged accounts.

Backups may run nightly today, while the target is quarterly successful restoration testing of critical systems.

Incident response may exist informally today, while the target includes documented roles, emergency contacts, escalation procedures, and annual tabletop exercises.

The gap becomes the improvement plan.

13.8 Improvement Should Be Prioritized by Risk

Not every gap deserves immediate attention.

Ask how important the affected asset is, how likely the weakness is to be exploited, how severe the consequence could be, whether sensitive or regulated information is involved, whether the weakness could stop business operations, whether many systems or users are affected, whether there is an easy improvement with significant benefit, and whether a contract requires action.

Group improvements into immediate priorities, near-term priorities, and longer-term improvements.

Everything cannot be first.

13.9 Use Simple Risk Language

A cybersecurity risk register does not need to look like a mathematical research paper.

Example:

Risk: Executive email compromise may result in fraudulent payments.
Likelihood: High.
Impact: High.
Existing controls: MFA, email filtering, payment verification.
Weakness: Some executives still use weaker authentication and vendor banking changes are occasionally approved through email only.
Priority: Immediate.
Owner: CFO and IT Director.
Action: Strengthen executive authentication and require independent verification of banking changes.
Target date: Within 30 days.

Cybersecurity risk management becomes useful when it connects a technical weakness to a business consequence and an accountable action.

13.10 Assign Owners to Improvements

Every significant improvement should have an owner, a target outcome, a target date, and a way to confirm completion.

Accountability turns recommendations into action.

13.11 Measure Trend, Not Just a Single Number

One measurement provides a snapshot. Repeated measurements show direction.

If critical vulnerabilities decline from 27 to 15 to 8 to 5, progress is visible.

If MFA coverage rises from 63 percent to 78 percent to 94 percent to 100 percent, progress is visible.

Trend matters.

13.12 Be Careful With Red, Yellow, and Green

Executive dashboards often use colors.

A useful interpretation might be:

Green: Control meets target and has been tested.
Yellow: Control exists but has gaps.
Red: Control is absent, ineffective, or materially below target.

A color should represent a meaningful outcome, not merely an activity.

13.13 Cybersecurity Maturity Is About Consistency

Cybersecurity maturity does not simply mean having advanced technology.

It reflects how consistently cybersecurity risk is understood, governed, managed, repeated, and improved.

NIST CSF 2.0 describes four Tiers: Tier 1 Partial, Tier 2 Risk Informed, Tier 3 Repeatable, and Tier 4 Adaptive.

The important idea is consistency.

13.14 Tier 1: Partial

At a very informal organization, cybersecurity may be reactive. Problems are handled when they occur. Priorities depend on whoever complains most loudly. Security responsibilities are unclear.

Improvement often begins with basic governance: who is responsible, what assets matter, what risks matter most, which controls already exist, and what needs to happen regularly.

13.15 Tier 2: Risk Informed

At this stage, leadership has greater awareness. Cybersecurity risks influence some decisions. Policies and practices may be approved, and responsibilities are clearer, but processes may still vary between departments.

The next step is making important practices repeatable across the business.

13.16 Tier 3: Repeatable

A repeatable organization has established cybersecurity risk-management practices.

Responsibilities are defined. Policies and procedures are maintained. Risk assessments occur periodically. Lessons from incidents influence improvements. Important vendor risks are managed. Metrics are reviewed.

The major difference is reliability.

13.17 Tier 4: Adaptive

At the adaptive level, the organization continually learns and adjusts.

Threat intelligence, incidents, technology changes, business changes, measurements, and lessons learned influence security practices.

The objective is not maturity for its own sake. The objective is effective risk management.

13.18 Create a Practical Cybersecurity Scorecard

A useful scorecard should fit on one or two pages.

Possible measures include:

Governance: cybersecurity owner assigned, annual risk review completed, policies reviewed, cyber insurance reviewed.

Identity: user MFA coverage, privileged MFA coverage, dormant accounts, access reviews.

Devices: managed endpoint coverage, endpoint encryption, EDR coverage, unsupported systems.

Vulnerabilities: open critical vulnerabilities, critical vulnerabilities older than 30 days, high-risk Internet-facing vulnerabilities.

Backup and Recovery: critical systems backed up, last restoration test, restore success rate.

Detection: critical systems forwarding security logs, security-alert monitoring, high-priority alerts older than SLA.

Incident Response: plan reviewed, tabletop exercise status, emergency contact list validated.

Vendors: critical vendors inventoried and reviewed; vendor accounts reviewed.

13.19 Avoid the Cybersecurity Score Trap

Some products reduce cybersecurity posture to a single number.

Single scores can be convenient. They can also be misleading.

Leadership should ask what the score represents, what was measured, what was not measured, how items were weighted, whether the score reflects business priorities, and whether the weaknesses behind the number are visible.

A summary score can be useful. It should never replace judgment.

13.20 Security Metrics Should Include People and Process

Useful measures may include time required to disable terminated employees, percentage of managers completing access reviews, percentage of high-value payment changes independently verified, employee phishing-reporting rate, time required to report suspected incidents, percentage of critical vendors with completed reviews, percentage of new systems receiving security review before purchase, and number of overdue cybersecurity improvement actions.

These measurements reveal whether cybersecurity is embedded in business operations.

13.21 Measure Detection Speed

Useful measurements might include average time between security alert and initial review, average time between confirmed compromise and containment, time between employee report and technical investigation, time required to disable a compromised account, and time required to isolate an endpoint.

The operational question is simple: How quickly can we recognize and control a problem?

13.22 Measure Recovery Readiness

Recovery metrics should focus on business outcomes.

How long did the last restoration test take? What percentage of critical systems were restored successfully? Did restoration meet business recovery objectives? Were the restored systems usable? Were recovery instructions accurate?

Recovery measurement must connect technology with business requirements.

13.23 Measure Vulnerability Age, Not Just Quantity

Useful vulnerability metrics include critical vulnerabilities, Internet-facing critical vulnerabilities, critical vulnerabilities older than the organization's remediation target, known exploited vulnerabilities affecting the environment, percentage of critical assets scanned, and average age of high-risk vulnerabilities.

The goal is prioritization, not simply producing a larger number.

13.24 Measure Identity Security

Possible identity indicators include MFA coverage, privileged MFA coverage, number of privileged accounts, dormant accounts, shared administrator accounts, inactive vendor accounts, time required to disable terminated employees, number of users with excessive privileges, frequency of access review, and percentage of risky sign-ins investigated.

The core question is: Do the right people have the right access for the right reasons?

13.25 Measure Third-Party Risk

Possible indicators include number of critical vendors, percentage with security review completed, percentage with current contracts defining cybersecurity responsibilities, percentage with documented incident-notification requirements, number of vendor accounts with privileged access, number of dormant vendor accounts, percentage of critical vendors reviewed annually, and number of critical services without a practical replacement or contingency.

13.26 Measure AI Governance

Possible indicators include approved AI systems inventoried, known unauthorized AI applications, percentage of employees acknowledging AI-use policy, number of AI

systems connected to sensitive data, percentage of high-risk AI systems receiving formal review, number of AI connectors with write access, AI-related security incidents, and percentage of enterprise AI accounts under centralized identity management.

AI risk should become part of normal governance.

13.27 A Metric Can Create Bad Behavior

Measurements influence behavior.

If a security team is judged solely by the number of vulnerabilities closed, it may focus on easy low-risk findings instead of difficult critical issues.

Metrics should therefore be interpreted carefully.

The question is not: How can we make the number look better?

The question is: Does this measurement reflect the outcome we actually want?

13.28 Combine Numbers With Explanation

A useful executive cybersecurity report includes both measurement and context.

Metrics without explanation can create false confidence. Explanation without measurement can become vague.

The strongest reporting combines both.

13.29 Report Cybersecurity in Business Language

A good cybersecurity report answers:

What is the risk?
Why does it matter?
What are we doing?
Who owns it?
When will it be addressed?
What decision does leadership need to make?

That is management information.

13.30 Create a 30-Day Cybersecurity Improvement Plan

A cybersecurity assessment becomes valuable when it produces action.

Days 1-5: Establish ownership and visibility. Identify the cybersecurity owner, critical systems, critical data, privileged accounts, key vendors, cyber insurance contacts, and incident-response contacts.

Days 6-10: Protect identity. Verify MFA coverage, strengthen privileged accounts, remove dormant users, review terminated employee access, review vendor accounts, and disable unnecessary legacy authentication.

Days 11-15: Confirm device and vulnerability protection. Verify managed endpoints, encryption, endpoint protection, unsupported systems, critical vulnerabilities, and Internet-facing weaknesses.

Days 16-20: Verify backup and recovery. Identify critical systems, confirm backup coverage and isolation, perform a meaningful restoration test, and document recovery gaps.

Days 21-25: Improve detection and response. Confirm logging, alert-monitoring responsibilities, incident contacts, account disablement, device isolation, and cyber insurance escalation.

Days 26-30: Review vendors and build the next plan. Identify critical vendors, review privileged vendor access, confirm incident-notification requirements, assign owners, and establish 90-day priorities.

The organization does not become perfectly secure in thirty days. It becomes more organized.

13.31 Build a 90-Day Improvement Roadmap

Possible priorities include completing MFA deployment, replacing unsupported systems, segmenting sensitive environments, improving backup isolation, extending cloud-log retention, implementing stronger endpoint monitoring, formalizing vendor review, conducting an incident-response tabletop, improving employee offboarding, strengthening payment verification, establishing an AI acceptable-use program, and remediating high-risk vulnerabilities.

Each project should connect to a documented risk.

13.32 Build an Annual Improvement Cycle

A practical annual cycle might include:

Quarter 1: Risk review and priority setting.
Quarter 2: Control testing and vendor review.
Quarter 3: Incident-response exercise and recovery testing.
Quarter 4: Leadership review, budget planning, and next-year priorities.

The important point is repetition.

13.33 Improvement Should Follow Business Change

Organizations acquire companies, open offices, hire remote employees, adopt cloud applications, move data, add vendors, develop software, use AI, win government contracts, and launch new services.

Each change may alter cybersecurity risk.

The improvement program should therefore ask periodically: What changed in the business?

Cybersecurity should follow the business.

13.34 A Realistic Scenario: The Dashboard That Looked Green

Consider a 90-person accounting company named Meridian Financial Services.

Its quarterly cybersecurity dashboard is almost entirely green: antivirus deployment, training completion, firewall status, backup jobs, password policy.

Then a ransomware attack occurs.

Investigators discover that backups had never been restored, several administrator accounts did not use MFA, endpoint protection was installed but alerts were not actively monitored, former vendor accounts remained active, and the incident-response plan had never been tested.

The dashboard was technically accurate.

The organization had measured the wrong things.

The lesson is simple: installation is not the same as protection, and completion is not the same as readiness.

13.35 A Realistic Scenario: Forty Findings and No Priorities

Another organization completes a cybersecurity assessment containing forty-seven recommendations.

Every item is marked important. No owners are assigned. No dates are established.

Six months later, only three recommendations have been completed.

The assessment identified weaknesses. It did not create an improvement process.

A shorter prioritized plan often produces more security improvement than a longer unprioritized report.

13.36 A Realistic Scenario: The Five Metrics That Changed the Conversation

A manufacturing company begins reporting five cybersecurity measures monthly:

1. Privileged MFA coverage.
2. Critical vulnerabilities older than 30 days.
3. Backup restoration success.
4. Time required to disable terminated users.
5. Critical vendor reviews completed.

The first report is uncomfortable. Six months later, each measure has improved substantially.

Nothing magical happened.

The organization measured what mattered and acted on what it saw.

13.37 Ask Whether Cybersecurity Investments Reduced Risk

Cybersecurity budgets should eventually connect to outcomes.

If leadership spends $40,000 on a new identity platform, it should later be able to identify measurable improvements such as higher MFA coverage, elimination of shared administrator accounts, centralized access reviews, automated termination disablement, and risky sign-in detection.

Cybersecurity spending should be tied to risk reduction where practical.

13.38 Do Not Wait for Perfect Measurement

Organizations may delay measurement because their data is incomplete.

Begin with what can be measured reasonably.

A useful imperfect measurement today may be more valuable than a perfect measurement system that never gets built.

13.39 Questions Every Business Leader Should Ask

What are our top five cybersecurity risks?
Which risks improved this quarter?
Which risks became worse?
Which critical controls are incomplete?
Which critical controls have actually been tested?
Are privileged accounts strongly protected?
Can we restore critical systems?
Can we detect suspicious activity?
Can we respond outside normal business hours?
Which vendors create the greatest dependency?
Are cybersecurity improvement actions overdue?
What changed in the business that may have changed our risk?
What cybersecurity investment should receive the next dollar?

13.40 Questions to Ask Your IT Team or Service Provider

1. What cybersecurity metrics do you recommend we review monthly?
2. Which metrics measure outcomes rather than simply activity?
3. What percentage of users have MFA?
4. What percentage of privileged accounts have stronger authentication?
5. How many critical vulnerabilities are currently open?
6. How old are they?
7. Which Internet-facing vulnerabilities require immediate attention?
8. What percentage of endpoints are fully managed?
9. What percentage have EDR?
10. What percentage are encrypted?
11. How many unsupported systems remain?
12. When was our last successful backup restoration?
13. How long did it take?

14. Which critical systems are not included in recovery testing?
15. How quickly can you disable a compromised account?
16. How quickly can you isolate an endpoint?
17. Which security logs are monitored?
18. How long are important logs retained?
19. When did we last test incident response?
20. How many critical vendors have been reviewed?
21. Which vendor accounts still have privileged access?
22. Which cybersecurity improvement items are overdue?
23. Which three risks should leadership prioritize during the next 90 days?
24. How will we measure whether those improvements worked?
25. What should this dashboard look like one year from now?

13.41 Chapter 13 Cybersecurity Measurement Checklist

Governance:
- A cybersecurity owner is assigned.
- Leadership reviews cybersecurity periodically.
- Top risks are documented.
- Major improvements have owners and target dates.
- Cybersecurity is considered during major business changes.

Identity:
- MFA coverage is measured.
- Privileged MFA coverage is measured separately.
- Dormant accounts are tracked.
- Terminated-user disablement time is measured.
- Privileged access is reviewed periodically.

Devices:
- Managed endpoint coverage is known.
- Encryption coverage is known.
- EDR coverage is known.
- Unsupported systems are tracked.

Vulnerabilities:
- Critical vulnerabilities are tracked.
- Vulnerability age is measured.

- Internet-facing weaknesses receive priority.
- Critical assets are included in scanning.

Backup and Recovery:
- Critical backup coverage is known.
- Restoration success is measured.
- Restoration time is measured.
- Recovery tests occur periodically.

Detection and Response:
- Alert-review responsibility is documented.
- Detection and containment times are reviewed.
- Incident-response exercises are tracked.
- Emergency contacts are validated.

Vendors:
- Critical vendors are inventoried.
- Vendor reviews are measured.
- Privileged vendor accounts are reviewed.
- Incident-notification requirements are documented.

AI:
- Approved AI systems are inventoried.
- High-risk AI uses are reviewed.
- AI integrations are tracked.
- AI access to sensitive systems is known.

Improvement:
- Current and target states are documented.
- Improvement priorities are risk-based.
- Owners are assigned.
- Target dates are established.
- Progress is reviewed regularly.

13.42 Key Terms in Plain English

Metric: A measurement used to understand performance, risk, coverage, or progress.

Key Performance Indicator (KPI): A measurement used to track progress toward an important objective.

Key Risk Indicator (KRI): A measurement that helps show the level or direction of risk.

Coverage: The percentage or portion of the intended environment to which a control applies.

Effectiveness: How well a control achieves its intended security outcome.

Current Profile: A description of cybersecurity outcomes an organization is currently achieving.

Target Profile: A description of cybersecurity outcomes an organization wants to achieve.

Gap: The difference between current and desired cybersecurity capability.

Cybersecurity Maturity: The degree to which cybersecurity risk-management practices are established, repeatable, measured, and improved.

Risk Register: A maintained record of important risks, their potential impact, controls, owners, and planned actions.

Risk Owner: The person responsible for ensuring that a particular risk is appropriately managed.

Trend: The direction a measurement moves over time.

Baseline: A starting measurement used for future comparison.

Target: A desired future measurement or outcome.

Dashboard: A summarized presentation of important measurements used to support decisions.

Remediation: Action taken to correct or reduce a cybersecurity weakness.

13.43 TRG Practical Takeaway

Do not try to measure cybersecurity with one number.

Measure the things that reveal whether important risks are being controlled.

Start small. Know your most important assets. Know your major risks. Measure basic security coverage. Test whether important controls work. Track vulnerabilities by risk and age. Measure identity protection. Test recovery. Measure incident readiness. Review vendor exposure. Assign owners to improvements. Track progress.

The purpose of cybersecurity measurement is not to create a beautiful dashboard.

The purpose is to help the organization make better decisions.

13.44 Before You Move On

Leadership should be able to answer:

1. What are our five most important cybersecurity risks?
2. What measurements tell us whether those risks are improving?
3. Which important controls are still incomplete?
4. Which controls have been tested recently?
5. Which cybersecurity improvement projects have owners and dates?
6. What should improve during the next 30 days?
7. What should improve during the next 90 days?

If the organization can answer those questions clearly, cybersecurity is becoming a managed business process rather than a collection of technical activities.

13.45 Key Takeaways

Cybersecurity activity is not automatically cybersecurity progress.

Measure outcomes. Measure coverage. Measure effectiveness. Measure trends.

Prioritize gaps by business risk.

Use current and target states to create an improvement roadmap.

Do not overwhelm leadership with hundreds of measurements.

A small set of meaningful indicators is better than a large set of irrelevant statistics.

Metrics should support decisions.

Cybersecurity maturity reflects consistency and discipline, not merely sophisticated technology.

Every improvement should have an owner. Every important target should have a date. Every significant control should eventually be tested.

13.46 Chapter Summary

Cybersecurity becomes manageable when leadership can see it.

Measurement creates that visibility.

The organization should understand where it currently stands, determine where it needs to go, identify meaningful gaps, prioritize them according to risk, assign responsibility, and monitor improvement.

That does not require a complex measurement program.

A small business can begin with a simple dashboard showing identity protection, endpoint coverage, vulnerabilities, backup restoration, incident readiness, and vendor risk.

The dashboard will evolve.

The organization will evolve.

Cybersecurity should evolve with it.

The ultimate objective is not a perfect score.

It is a stronger organization.

13.47 From Measurement to Action

At this point, the reader has moved through the major cybersecurity issues addressed throughout this book.

We have discussed what organizations are protecting. We have examined identity, people, technology, operations, data, vendors, cloud services, government-contracting responsibilities, artificial intelligence, and incident response.

Now those ideas must become action.

A cybersecurity program improves when leadership takes the next reasonable step.

Not every possible step.

The next one.

For one company, that may mean enabling MFA. For another, it may mean testing backups. For another, it may mean removing old vendor accounts. For another, it may mean conducting its first incident-response exercise. For another, it may mean identifying where sensitive information is stored.

The right action depends on risk.

The important thing is to begin.

The 30-Day Cybersecurity Action Plan

Week 1 - Know Your Environment:
Identify critical systems, information, accounts, vendors, and business processes.

Week 2 - Reduce Immediate Exposure:
Strengthen MFA, remove unnecessary access, patch critical vulnerabilities, and verify endpoint protection.

Week 3 - Verify Resilience:
Test backups, review security logging, validate incident contacts, and confirm recovery procedures.

Week 4 - Establish Ongoing Management:
Review critical vendors, assign improvement owners, establish a leadership scorecard, and set the next 90-day priorities.

At the end of thirty days, the organization should not expect cybersecurity risk to disappear.

It should expect something more realistic and more valuable: better visibility, clearer responsibility, fewer obvious weaknesses, stronger recovery readiness, and a practical plan for continued improvement.

That is how cybersecurity becomes sustainable.

That is how risk becomes manageable.

And that is the central message of Cybersecurity Without the Jargon.

CYBERSECURITY WITHOUT THE JARGON

A Practical Guide to Protecting Your Business, Your Data, and Your People

Chapter 14

Cybersecurity as a Leadership Responsibility: Putting It All Together

Authored by Dr. Charles Edwards

Cybersecurity began as a technology problem in the minds of many business leaders.

By the end of this book, it should look very different.

Cybersecurity is still about technology.

Firewalls matter. Endpoint protection matters. Cloud configurations matter. Encryption matters. Multifactor authentication matters. Logging matters. Backups matter. Secure applications matter.

But none of those things, individually or together, fully defines cybersecurity.

Cybersecurity is ultimately about the ability of an organization to continue doing what it exists to do while protecting the people, information, systems, customers, partners, and trust upon which that mission depends.

That makes cybersecurity a leadership responsibility.

It does not mean the chief executive should configure firewalls or that a business owner must understand every security protocol.

It means leadership has to understand the risks well enough to ask the right questions, assign responsibility, set priorities, provide appropriate resources, and make informed decisions.

That is where this book has been heading from the beginning.

The central lesson is not that every business needs more technology.

The central lesson is that every business needs a practical way to understand and manage cyber risk.

14.1 Cybersecurity Is Part of Running the Business

Imagine a company whose leadership says: “Cybersecurity is handled by IT.”

Now imagine asking:

Who approves financial transactions?
Who hires and terminates employees?
Who decides which vendors are used?
Who signs contracts?
Who determines which customer information the company collects?
Who approves remote work?
Who purchases cloud applications?

Who decides whether employees may use artificial intelligence?
Who determines how long the company can tolerate an outage?
Who speaks to customers after an incident?
Who decides how much risk the organization is willing to accept?

Those are not purely IT decisions.

They are business decisions.

Cybersecurity touches them because each decision affects information, access, technology, or operational risk.

A finance process can either reduce or increase fraud risk. A Human Resources process can either remove or leave behind former employee access. A purchasing decision can introduce a new cloud provider. A contract can create cybersecurity obligations. A remote-work policy can expand where business information travels. An AI policy can determine whether confidential information reaches an external system.

Cybersecurity is therefore part of normal business management.

The strongest organizations recognize that reality.

14.2 Leadership Does Not Need to Know Everything

One reason executives sometimes disengage from cybersecurity is the belief that the subject is too technical.

Security professionals may discuss EDR, SIEM, XDR, SASE, ZTNA, IAM, PAM, CSPM, DLP, MDR, SOC, CVE, and CVSS.

Those acronyms are useful shorthand for technical professionals, but they can become a barrier for business leaders. The terms are explained below in plain English.

EDR - Endpoint Detection and Response

EDR is security technology that monitors laptops, desktops, and servers for suspicious activity and helps security teams investigate and stop attacks. A useful way to think about EDR is as a combination of a security camera and a response system for company computers.

SIEM - Security Information and Event Management

A SIEM collects security logs from many sources, such as firewalls, servers, cloud services, applications, and user accounts. It helps security teams search those logs and identify patterns that may indicate suspicious activity.

XDR - Extended Detection and Response

XDR expands the idea of EDR beyond individual computers. It combines information from endpoints, email, identity systems, cloud services, networks, and other security tools so an attack can be investigated across the organization rather than from one device alone.

SASE - Secure Access Service Edge

SASE is a cloud-based approach that combines networking and security so employees can securely access business resources from offices, homes, or other locations. It is designed for organizations whose users and applications no longer sit entirely inside one traditional corporate network.

ZTNA - Zero Trust Network Access

ZTNA provides users with access only to the specific applications or resources they are authorized to use. It does not automatically trust someone simply because the person connected to the company network or VPN.

IAM - Identity and Access Management

IAM refers to the systems and processes used to manage user identities and access. It answers questions such as: Who has an account? How does the person prove identity? What applications or information is the person allowed to use?

PAM - Privileged Access Management

PAM focuses on highly powerful accounts, such as system administrators, database administrators, and others who can make major changes. It helps protect, monitor, and control access to those privileged accounts.

CSPM - Cloud Security Posture Management

CSPM technology examines cloud environments for unsafe configurations, excessive permissions, publicly exposed resources, and other security weaknesses.

DLP - Data Loss Prevention

DLP refers to technology and policies designed to help prevent sensitive information from being copied, emailed, uploaded, or shared where it should not go.

MDR - Managed Detection and Response

MDR is a cybersecurity service in which an outside security team monitors systems for threats, investigates suspicious activity, and assists with response. Many MDR services operate around the clock.

SOC - Security Operations Center

A SOC is the people, processes, and technology responsible for monitoring cybersecurity activity, investigating alerts, and responding to threats. A SOC can be operated internally or by an outside provider.

CVE - Common Vulnerabilities and Exposures

CVE is a standardized identification system for publicly known cybersecurity vulnerabilities. A CVE number acts like a reference number so vendors and security professionals can be certain they are discussing the same vulnerability.

CVSS - Common Vulnerability Scoring System

CVSS is a standardized method for estimating the technical severity of a vulnerability, usually on a scale from 0 to 10. Higher scores generally indicate greater technical severity, although an organization still has to consider whether the vulnerability affects its own environment and business.

Several distinctions are useful.

EDR is primarily technology that monitors endpoints. MDR is primarily a service in which security professionals monitor and respond to security activity, often using EDR technology.

A SIEM is a technology platform that gathers and analyzes security information. A SOC is the team or operation that uses tools such as a SIEM to monitor and investigate activity.

IAM manages identities and access broadly. PAM focuses specifically on the organization's most powerful accounts.

EDR concentrates mainly on endpoints such as laptops and servers. XDR attempts to combine endpoint information with signals from email, identity, cloud systems, networks, and other sources.

A CVE tells you which vulnerability you are discussing. CVSS helps describe how technically severe that vulnerability may be.

Cybersecurity professionals use acronyms because they make technical conversations faster, but acronyms can also create an unnecessary barrier between technology teams and business leaders. A business owner does not need to memorize every abbreviation. What matters is understanding the business purpose behind the technology.

When someone recommends EDR, SIEM, PAM, DLP, or another security product, leadership should be comfortable asking:

What problem does this solve for our organization?

What risk does it reduce?

How will we know it is working?

Leadership does not need to master every acronym.

It needs to understand the question behind the acronym.

Instead of asking only, “Do we have EDR?” leadership can ask, “Can we detect and contain suspicious activity on company computers?”

Instead of asking, “Do we have PAM?” ask, “How are powerful administrator accounts protected and monitored?”

Instead of asking, “Do we have DLP?” ask, “How do we prevent sensitive information from leaving through inappropriate channels?”

Instead of asking, “Do we have ZTNA?” ask, “How do remote users receive only the access they actually need?”

Instead of asking, “What is our SIEM?” ask, “Where do our important security logs go, and who reviews them?”

Those are leadership questions.

Technical professionals can explain how the answer is implemented.

14.3 Ask What Happens If the Control Fails

Cybersecurity discussions often focus on whether a protection exists.

A stronger leadership question is: What happens if it fails?

Suppose email filtering fails and a phishing message reaches an employee.

Do we have MFA? Does finance independently verify payment changes? Can suspicious activity be detected? Can the account be disabled quickly?

Suppose endpoint protection fails and malware executes.

Can the device be isolated? Can the attacker move freely through the network? Are administrator credentials protected? Can affected systems be restored?

Suppose a cloud provider experiences an outage.

Can employees continue working? Do we have alternate processes? How long can operations tolerate the interruption?

Suppose a backup repository is compromised.

Do we have another recovery copy?

Suppose an employee ignores policy.

Can technical controls limit the damage?

This is resilience thinking.

Cybersecurity is stronger when one control is not the only thing standing between the organization and disaster.

14.4 Think in Layers

A useful cybersecurity program uses multiple layers of protection.

Consider business email compromise.

One layer is email filtering. Another is employee awareness. Another is MFA. Another is conditional access. Another is unusual-login monitoring. Another is financial verification. Another is mailbox auditing. Another is rapid incident response.

If one layer fails, another may still prevent loss.

Now consider ransomware.

One layer is patching. Another is endpoint detection. Another is limited administrator access. Another is network segmentation. Another is monitoring. Another is backup isolation. Another is recovery testing. Another is incident response.

This approach is sometimes called defense in depth.

The phrase sounds technical.

The idea is simple: Do not depend on one thing.

14.5 Know Your Crown Jewels

Not everything in the organization deserves exactly the same level of protection.

A public marketing brochure does not require the same controls as payroll records.

A temporary test workstation does not have the same business value as the accounting system.

A general user account does not create the same risk as a domain administrator account.

An organization therefore needs to know what matters most.

Critical assets include important information, critical business systems, high-impact accounts, financial processes, customer relationships, operational capabilities, intellectual property, contract obligations, and essential vendors.

Leadership should periodically ask:

If one thing were stolen, what would hurt us most?
If one system were unavailable for a week, what would stop the business?
If one account were compromised, which account could cause the greatest damage?
If one vendor failed, which dependency would disrupt us most?

Those questions reveal cybersecurity priorities.

14.6 Protect Identity First

Across modern organizations, identity has become one of the most important security boundaries.

An attacker with valid credentials may not need sophisticated malware.

The attacker may simply log in.

That is why so many lessons throughout this book return to identity.

Require MFA. Protect administrator accounts more strongly. Remove dormant access. Disable former employees quickly. Review vendor accounts. Avoid unnecessary privilege. Monitor unusual authentication. Control third-party application permissions. Protect password-reset paths.

An organization that improves identity security often reduces risk across many systems simultaneously.

14.7 Protect the Administrator More Than the Average User

Not all accounts create equal risk.

An ordinary user may have access to email, files, and selected applications.

An administrator may be able to create users, reset passwords, change security settings, disable protections, read other mailboxes, modify cloud services, access servers, install software, and destroy logs.

An administrator account is therefore a high-value target.

Organizations should know how many privileged accounts exist, who uses them, whether they are separate from everyday accounts, whether strong MFA is required, whether activities are logged, whether privileges are reviewed, whether vendors hold privileged accounts, and whether shared administrator accounts are still being used.

14.8 Make the Secure Path the Easy Path

Employees frequently choose insecure workarounds because the approved process is difficult.

An employee uses personal cloud storage because company file sharing is frustrating.

A worker sends information to a personal email account because remote access is slow.

A developer pastes code into an unapproved AI tool because the approved system cannot perform the task.

A manager shares a password because requesting separate access takes three days.

A salesperson installs an unauthorized application because procurement takes weeks.

Security teams may respond by blaming the employee.

Sometimes employee behavior absolutely needs correction.

But leadership should also ask: Why was the insecure method easier?

Good cybersecurity should remove unnecessary friction wherever practical.

14.9 Security Culture Is Created by What Leaders Do

Employees pay attention to leadership behavior.

If executives bypass security rules because they are busy, employees learn that security rules are optional.

If senior managers share passwords, employees learn that password sharing is acceptable.

If an executive becomes angry because a finance employee verifies an unusual payment request, other employees learn not to challenge suspicious transactions.

If leadership delays important security improvements year after year while telling employees that cybersecurity is a priority, employees notice the contradiction.

Security culture is not created by posters.

It is created by behavior.

14.10 Your IT Provider Is a Partner, Not the Owner of All Risk

Small and mid-sized organizations frequently rely on managed service providers.

That can be an excellent model.

A strong provider can deliver skills, monitoring, technology management, and support that would be expensive to build internally.

But a provider cannot make every business decision.

The provider cannot determine which customer information is most sensitive, what business interruption is tolerable, which contracts create obligations, how much financial risk leadership will accept, which vendors are strategically critical, whether a particular AI use is appropriate, or how the organization should communicate with customers during a crisis.

The provider can advise.

Leadership must decide.

14.11 Ask Providers for Evidence, Not Reassurance

A common cybersecurity conversation sounds like this:

“Are our backups good?” “Yes.”
“Are we secure?” “Yes.”
“Is MFA enabled?” “Yes.”
“Are we monitoring everything?” “Yes.”

Those answers provide comfort.

They do not provide evidence.

A better conversation asks:

Show us the successful restoration test.
Show us MFA coverage.
Show us which systems are monitored.
Show us the list of privileged accounts.
Show us the unresolved critical vulnerabilities.
Show us how quickly former users are removed.
Show us the incident-response escalation process.
Show us the critical vendor list.

Verification is simply good management.

14.12 Cybersecurity Budgets Should Follow Risk

Security vendors sell many useful products.

Organizations cannot buy all of them.

That means budget decisions require prioritization.

Security spending should solve important problems, not merely purchase impressive technology.

14.13 Risk Cannot Be Reduced to Zero

No organization can eliminate all cybersecurity risk.

Every useful technology creates some exposure.

Cybersecurity is not about eliminating all technology because technology creates risk.

It is about understanding risk and deciding which risks are acceptable.

14.14 Know the Difference Between Accepting Risk and Ignoring It

Risk acceptance should be deliberate.

An unsupported legacy system that is known, isolated, monitored, and scheduled for replacement represents accepted residual risk.

An unsupported system that nobody has reviewed in years is unmanaged risk.

The difference is knowledge and decision-making.

14.15 Cybersecurity Must Be Able to Survive Staff Changes

Many smaller organizations depend heavily on one knowledgeable employee.

Important cybersecurity processes should be documented sufficiently that essential capabilities survive personnel changes.

Documentation does not replace skilled employees.

It reduces organizational dependency on individual memory.

14.16 Build Cybersecurity Into Hiring

Cybersecurity begins before an employee receives a laptop.

Human Resources, management, and IT should coordinate access based on job requirements.

Good onboarding establishes secure behavior from the beginning.

14.17 Role Changes Need Security Review

Employees receive promotions, move between departments, take temporary assignments, join projects, and gain administrative responsibility.

The problem occurs when old access remains.

Periodic access review helps control access creep.

14.18 Offboarding Is a Cybersecurity Control

When an employee leaves, accounts may need to be disabled, sessions revoked, devices returned, cloud applications removed, and privileged access terminated.

Offboarding should be a defined business process.

14.19 Vendors Are Part of Your Security Environment

Every vendor relationship creates some form of dependency.

Leadership should know which vendors are critical, what information they hold, what systems they can access, what happens if they are compromised or unavailable, how they will notify the organization, and how quickly access can be removed.

Supplier risk is business risk.

14.20 Cloud Does Not Remove Responsibility

Cloud services can improve availability, scalability, collaboration, and security.

But customers remain responsible for important decisions involving identity, permissions, sharing, configuration, data handling, applications, logging, and account life cycle.

Cloud security is shared responsibility.

14.21 Artificial Intelligence Should Become a Normal Risk Discussion

As AI becomes part of ordinary business technology, AI risk should become part of ordinary management.

When reviewing a new AI system, ask what problem it solves, what information it receives, who can use it, what systems it can access, whether it can take actions, which vendor provides it, what contractual protections exist, what happens if the answer is wrong, and whether access can be revoked.

14.22 Prepare for the Incident You Hope Never Happens

The most secure organizations still prepare for incidents.

That is not pessimism.

It is responsible management.

Know who leads the response, who has authority, how to isolate systems, how to contact insurance and legal counsel, how to reach the bank, how to operate if email is unavailable, and how to restore critical systems.

Then test the plan.

14.23 Recovery Is a Leadership Capability

Technical teams restore systems.

Leadership restores the business.

Cyber recovery connects directly to business continuity.

Leadership should understand recovery priorities before an incident forces those decisions.

14.24 Measure What Helps You Decide

Do not measure something merely because the software produces a number.

Measure what helps leadership make a decision.

Useful questions include whether important accounts are protected, critical systems are patched, backups can be restored, suspicious activity can be detected, compromised systems can be isolated, vendors are reviewed, former employees are removed, and improvement projects are being completed.

14.25 Review Cybersecurity Like Other Business Risk

Cybersecurity should become part of normal leadership review.

Leadership should leave the meeting knowing what matters, what is changing, what needs attention, and who owns the next action.

14.26 A Practical Quarterly Leadership Conversation

A quarterly cybersecurity conversation might ask:

What are our top five cyber risks today?
What changed during the quarter?
What incidents or near misses occurred?
Which critical controls have been tested?
Which risks improved?
Which risks became worse?
What decisions are needed from leadership?

This meeting does not require a large cybersecurity department.

It requires disciplined questions.

14.27 Cybersecurity Should Support Growth

Good cybersecurity should help the organization grow safely.

The objective is not to stop change.

The objective is to make change safer.

14.28 Security and Convenience Must Be Balanced Deliberately

Some controls create friction.

Leadership should evaluate those tradeoffs deliberately.

Good cybersecurity is not maximum restriction.

It is appropriate control.

14.29 Cybersecurity Is Never Finished

There will always be another vulnerability, technology, vendor, employee, cloud application, attack technique, regulatory requirement, AI capability, and business change.

Businesses do not “finish” finance, Human Resources, operations, customer service, or legal compliance.

They manage them.

Cybersecurity should be viewed the same way.

14.30 Do the Basics Exceptionally Well

Many damaging incidents still succeed because basic controls were incomplete.

Passwords were weak. MFA was missing. Access was excessive. Software was outdated. Backups were untested. Email requests were not verified. Former employees still had accounts. Administrator credentials were exposed. Logs were not reviewed. Employees did not know whom to call. The organization had no incident plan.

There is enormous security value in doing ordinary things consistently well.

14.31 The Ten Questions I Would Ask Any Organization

1. What information and systems could hurt you most if they were stolen, changed, or unavailable?
2. Who has administrator access?
3. Is strong MFA required for important accounts?
4. Can you restore your most important system from backup?
5. How would you know if an attacker were already inside your environment?
6. What happens when an employee leaves?
7. Which vendors can access your systems or sensitive information?

8. Who leads the response if ransomware appears tomorrow morning?
9. What are your three most important cybersecurity improvements this year?
10. How do you know those improvements are actually happening?

Those ten questions will not replace a full assessment.

They will reveal whether cybersecurity is being managed.

14.32 A Practical Leadership Cybersecurity Checklist

Know the Business:
- We know our critical systems.
- We know our most sensitive information.
- We know our critical business processes.
- We know which vendors are essential.

Protect Identity:
- Important accounts use MFA.
- Privileged accounts receive stronger protection.
- Dormant accounts are removed.
- Employee access changes with job responsibilities.
- Access for departed users is disabled promptly.

Protect Technology:
- Managed devices are patched.
- Endpoint protection is active.
- Encryption is used appropriately.
- Internet-facing systems are reviewed.
- Unsupported systems are tracked.

Protect Information:
- Sensitive information is identified.
- Sharing is controlled.
- Cloud access is reviewed.
- AI use follows defined information-handling rules.

Protect Business Processes:
- High-value payments require verification.
- Important changes require approval.
- Security responsibilities are coordinated across departments.

Manage Vendors:
- Critical vendors are known.
- Vendor access is reviewed.
- Security responsibilities are documented.
- Vendor incidents can be escalated quickly.

Prepare for Incidents:
- Employees know how to report concerns.
- Emergency contacts are available.
- Cyber insurance procedures are understood.
- Backups have been restored successfully.
- Incident-response exercises are conducted.

Govern:
- Leadership reviews cyber risk periodically.
- Top risks are documented.
- Improvement actions have owners.
- Progress is measured.
- Risk acceptance is deliberate.

14.33 The Leadership Principle Behind the Entire Book

Every chapter in this book has approached cybersecurity from a different angle.

But the same principle appears repeatedly:

Know what matters, understand what can go wrong, put reasonable safeguards in place, verify that they work, and prepare to respond when they do not.

Cybersecurity tools will change.

The management principle will remain.

14.34 Cybersecurity Is Ultimately About Trust

Customers trust organizations with information.

Employees trust employers with personal records.

Business partners trust each other with access.

Government agencies trust contractors with protected information.

Banks trust authorized transaction processes.

Leadership trusts technology teams.

Technology teams trust vendors.

Every digital interaction depends on trust.

Cybersecurity protects that trust.

14.35 You Do Not Need to Solve Everything Today

A small organization reading this book may identify many gaps.

Do not try to solve everything at once.

Start with what matters most.

Cybersecurity improvement is cumulative.

Small improvements matter.

The important thing is consistency.

14.36 The Next Reasonable Step

Every organization is different.

But every organization can ask:

What is the next reasonable cybersecurity step for us?

The next step may not be perfect.

It does not have to be.

It has to move the organization in the right direction.

14.37 Final Leadership Questions

1. What are we protecting?
2. Which cyber risks could cause the greatest business harm?
3. Who owns cybersecurity risk?
4. Are important identities strongly protected?
5. Are critical systems monitored and patched?
6. Can we recover from a major outage or ransomware event?
7. Do we know which vendors create significant cybersecurity dependency?
8. Do employees know how to report suspicious activity?

9. Are AI and cloud services governed intentionally?
10. Have we practiced incident response?
11. What are our top cybersecurity improvement priorities?
12. How will leadership know whether those improvements are occurring?

If the answers are clear, the organization has moved beyond cybersecurity as a collection of products.

It has begun managing cybersecurity as a business discipline.

14.38 Final TRG Practical Takeaway

Cybersecurity does not need to be mysterious.

Start with the business.

Know what matters.

Protect identity.

Limit privilege.

Keep systems maintained.

Protect sensitive information.

Control financial processes.

Understand vendors.

Use cloud and AI deliberately.

Monitor important activity.

Test backups.

Prepare for incidents.

Measure improvement.

Ask questions.

Verify answers.

Repeat.

That is cybersecurity without the jargon.

14.39 A Final Word to the Reader

If you remember only one thing from this book, remember this:

Cybersecurity is not a technology you buy. It is a responsibility you manage.

Technology will help. Good people will help. Strong providers will help. Policies will help. Training will help. Insurance will help.

But leadership still has to decide what matters, what risk is acceptable, what needs improvement, and whether those improvements are actually occurring.

You do not need to become a cybersecurity expert.

You do need to remain engaged.

Ask questions. Expect clear answers. Do not let jargon end the conversation. Do not assume a product automatically solved the problem. Do not assume a quiet year means a secure year. Do not wait for an incident to discover what matters.

The goal is not perfection.

The goal is preparation.

The goal is resilience.

The goal is informed action.

Conclusion

The world will continue to change.

Technology will continue to change.

Cyber threats will continue to change.

Your organization will change too.

Cybersecurity must change with it.

But the basic discipline remains simple:

Know. Protect. Verify. Prepare. Improve.

Know what matters.

Protect it reasonably.

Verify that protections work.

Prepare for failure.

Improve continuously.

Then repeat.

That is how organizations manage cybersecurity risk.

That is how leaders protect businesses, customers, employees, information, operations, and trust.

And that is the purpose of Cybersecurity Without the Jargon.

A Practical Guide to Protecting Your Business, Your Data, and Your People

Authored by Dr. Charles Edwards